

Guidance-Enriched Visual Analytics

PhD THESIS

submitted in partial fulfillment of the requirements for the degree of

Doctor of Technical Sciences

within the

Vienna PhD School of Informatics

by

Davide Ceneda, MSc.

Registration Number 1528046

to the Faculty of Informatics

at the TU Wien

Advisor: Univ.Prof.in Dr.in rer.soc.oec. Silvia Miksch

Second advisor: Priv. Doz. Dr.-Ing. habil. Christian Tominski

External reviewers:

Prof. Christopher Collins, Ontario Tech University, Canada

Prof. Tim Dwyer, Monash University, Australia

Vienna, 20th August, 2020

Davide Ceneda

Silvia Miksch



Die approbierte gedruckte Originalversion dieser Dissertation ist an der TU Wien Bibliothek verfügbar.
The approved original version of this doctoral thesis is available in print at TU Wien Bibliothek.

Declaration of Authorship

Davide Ceneda, MSc.

I hereby declare that I have written this Doctoral Thesis independently, that I have completely specified the utilized sources and resources and that I have definitely marked all parts of the work - including tables, maps and figures - which belong to other works or to the internet, literally or extracted, by referencing the source as borrowed.

Vienna, 20th August, 2020

Davide Ceneda



Die approbierte gedruckte Originalversion dieser Dissertation ist an der TU Wien Bibliothek verfügbar.
The approved original version of this doctoral thesis is available in print at TU Wien Bibliothek.

Acknowledgements

Thanks...

- To Federica, my family, my parents, and brothers for the all the love and support. ps - You owe me an aperitif, now.
- To Silvia for giving me the opportunity to come to Vienna and do research.
- To the whole CVASt group for the great company and support, you were more than colleagues: Theresia, Christian, Markus, Roger, Paolo, Alessio, Victor, Velitchko, and Nikolaus. A special thanks goes to Theresia, Markus, and Christian, who had the patience to deal with me every single day of the last five years.
- To all my co-authors, advisors, reviewers, and fellow scientists involved in the works cited in this thesis, for sharing your time and knowledge with me.

Finally, the research described in this dissertation was supported and funded by the Austrian Science Fund (FWF) with the grant #P31419-N31



Die approbierte gedruckte Originalversion dieser Dissertation ist an der TU Wien Bibliothek verfügbar.
The approved original version of this doctoral thesis is available in print at TU Wien Bibliothek.

Kurzfassung

Stetig größer werdende Datenmengen machen es immer notwendiger, daraus nützliche Informationen abzuleiten. Die Bedeutung erfolgreicher Datenanalyse ist offensichtlich: Wir verwenden Datenanalyse häufig unbewusst, wenn wir Produkte vergleichen oder nach Unterkünften für unseren nächsten Urlaub suchen. Für ein erfolgreiches Unternehmen ist Datenanalyse sogar noch wichtiger, indem sie die Grundlage für viele wichtige Entscheidungen bildet. Informationsvisualisierung und Visual Analytics (VA) sind zwei Bereiche der Datenanalyse, die auf Basis visueller Darstellungen von Daten verschiedenste analytische Aufgaben unterstützen. Obwohl die Verwendung visueller Datenanalysemethoden weit verbreitet ist, ist der Erkenntnisgewinn ein schwieriger Prozess. Probleme können in jeder Phase der Analyse auftreten: Typische Herausforderungen entstehen bei der Vorbereitung der Daten für die Analyse oder der Wahl geeigneter Analysemodelle sowie visueller Darstellungen für eine bestimmte Aufgabe. Benutzer*innen stoßen oft auf Schwierigkeiten, wenn sie Daten untersuchen und Ergebnisse im Kontext eines bestimmten Fachgebiets interpretieren, Ergebnisse strukturieren, um daraus Erkenntnisse abzuleiten, Arbeitshypothesen beweisen und letztendlich neues Wissen generieren. In der Folge sind die Benutzer*innen bei Problemen mitunter überfordert, die Analyse gerät ins Stocken und die Leistungsfähigkeit der visuellen Analyseansätze wird gemindert bzw. die Qualität der daraus resultierenden Ergebnisse beeinträchtigt.

Aus den genannten Gründen ist es wichtig, die Benutzer*innen durch *Guidance* zu unterstützen. Ziel von *Guidance* ist es, eine effektive Nutzung der Analysewerkzeuge zu fördern und den Anwender*innen zu helfen, mögliche Probleme während der Analyse zu bewältigen. Obwohl in der Vergangenheit gerade die Unterstützung der Benutzer*innen als eines der Hauptziele der visuellen Datenanalyse galt, wurde noch nie systematisch untersucht, wie dieses Ziel tatsächlich erreicht werden könnte. Aufbauend auf frühere Forschungsarbeiten und unter der Verwendung einer nutzer*innen-zentrierten Methodik, werden in der vorliegende Arbeit *Guidance*-Ansätze umfassend charakterisiert. Ein besonderer Schwerpunkt liegt auf *Guidance* in den VA. In einem ersten Schritt (1) wird *Guidance* definiert. Danach werden die Hauptaspekte des Guiding-Prozesses zusammengefasst. Dabei wird gleichzeitig veranschaulicht, was *Guidance* überhaupt ausmacht. Weiters (2) werden die Auswirkungen verschiedener Arten von *Guidance* beschrieben, und zwar im Speziellen, wenn analytische Aufgaben von Nutzer*innen mit unterschiedlichen Fachkenntnissen durchgeführt werden. Es wird beleuchtet, wie *Guidance* die Umsetzung der Analyse beeinflusst bzw. wie die Nutzer*innen darauf reagieren. Schließlich (3) wird

skizziert, wie *Guidance* mit einem schrittweisen Vorgehen in VA-Ansätze integriert werden kann, d.h. direkt am Anfang der Entwicklung von VA-Tools. Obwohl noch viele Herausforderungen vor uns liegen, stellen die beschriebenen Ergebnisse erste Belege für den Wert von *Guidance* in den VA sowie dessen positive Auswirkungen auf die Nutzer*innen dar. Gleichzeitig wird ein Grundstein für die Entwicklung inspirierender Forschungsprototypen gelegt, die gezielt *Guidance* miteinbeziehen (Guidance-Enriched VA).

Abstract

As we continue producing data, the need to extract useful information from it has become essential. The implications of pursuing successful data analysis are before our eyes. We use it, although often unconsciously, when we compare products or check the availability of rooms for our next vacation. Data analysis is even more important for successful business, as it provides the foundation for solid decision making. Information visualization and visual analytics (VA) are two prolific branches of data analysis exploiting external means—namely, visualizations—to support the execution of analytical tasks. However, despite the vast adoption of such techniques, the issue with visual data analysis is that gaining insights is often quite a challenging process. Issues may arise at any phase of the analysis: Typical challenges include preparing data for the analysis or choosing appropriate analytical models and visual means for a given task. Users may encounter difficulties when exploring data and interpreting findings in light of domain knowledge, when organizing findings consistently into insights, when proving working hypotheses, and when generating new knowledge. The consequence is that, when issues arise, users are typically overwhelmed, the analysis stalls, and the efficacy of VA approaches is reduced; in addition, the quality of the resulting results is impaired.

For these reasons, it is important to support users with guidance. The aim of guidance is to foster an effective use of analysis tools and help users overcome any possible issue that might occur during the analysis. Historically, supporting users has been one of the very main goals of visual data analysis. However, how to reach this goal has never been researched in a systematic way. In this thesis, building on previous research and utilizing a user-centered methodology, we describe a thorough characterization of guidance approaches. Particular emphasis is given to guidance in VA. The main contributions of this thesis are: (1) we characterize guidance and summarize the main aspects of the process of guiding, also illustrating what constitutes guidance in the first place; (2) we report the effects of providing different types of guidance to users performing analytical tasks with different levels of expertise, shedding light on how guidance influences the way the analysis is conducted and how users react to it; and (3) we outline how guidance can be integrated with a step-by-step procedure in VA approaches—that is, directly at the point of designing VA tools. Although many open challenges still lie ahead, the results described in this thesis represent an initial demonstration of the value of guidance in VA and its positive effects on users while preparing a solid foundation for the development of inspiring guidance-enriched VA prototypes.



Die approbierte gedruckte Originalversion dieser Dissertation ist an der TU Wien Bibliothek verfügbar.
The approved original version of this doctoral thesis is available in print at TU Wien Bibliothek.

Contents

Kurzfassung	vii
Abstract	ix
Contents	xi
I Overview	1
1 Guidance in Visual Analytics	3
1.1 Motivation and Problem Description	3
1.2 Objectives and Research Questions	4
1.3 Conceptual Grounding and Methodology	6
1.4 Defining Guidance	13
1.5 A Review of Guidance in Visual Analytics	25
1.6 Steering the Guidance Process	34
1.7 The Effects of Guidance	41
1.8 Designing Effective Guidance	54
1.9 Scientific Contributions and Future Perspectives	66
1.10 Publications	70
Bibliography	73
II Papers	85
2 Characterizing Guidance in Visual Analytics	87
2.1 Abstract	87
2.2 Introduction	88
2.3 Guidance: Terminology and General Concepts	89
2.4 A Review of Guidance in Visual Analytics	96
2.5 Discussion and Future Work	106
2.6 Conclusion	109
	xi

Bibliography	111
3 You Get by with a Little Help: The Effects of Variable Guidance Degrees on Performance and Mental State	117
3.1 Abstract	117
3.2 Introduction	118
3.3 Related Work	119
3.4 Aims and Terminology	121
3.5 Hypotheses	122
3.6 Study Design	124
3.7 Results	134
3.8 Discussion and Future Work	142
3.9 Conclusion	145
Bibliography	151
4 Guide Me in Analysis: A Framework for Guidance Designers	155
4.1 Abstract	155
4.2 Introduction	156
4.3 Related Work	157
4.4 Motivating Example	158
4.5 A Framework for Guidance Designers	161
4.6 Designing Guidance: Three Scenarios	174
4.7 Discussion	189
4.8 Conclusion and Future Challenges	190
Bibliography	193
List of Figures	197
List of Tables	205

Part I

Overview



Die approbierte gedruckte Originalversion dieser Dissertation ist an der TU Wien Bibliothek verfügbar.
The approved original version of this doctoral thesis is available in print at TU Wien Bibliothek.

Guidance in Visual Analytics

1.1 Motivation and Problem Description

This dissertation presents and describes the process of characterizing, designing, and evaluating guidance approaches to support analysts pursuing Visual Analytics (VA) tasks.

Before getting to the details, we start by discussing two basic questions: (1) What is guidance? and (2) Why do we need guidance in the first place?

The concept of guidance simply refers to the act of helping somebody reach a goal. It is also well known that we experience guidance since birth. First supported by our parents and then supervised and directed by teachers, we spend our entire childhood and part of our adult life learning how to trace our path in life. The reason is pretty obvious: No one is born with the experience and the knowledge to tackle all the challenges s/he will encounter in life. The need for guidance stems from this contrast between *what we need to know* to pursue a given goal and *what we actually know* or are able to do at a specific point in time. We will refer to this discrepancy as the *knowledge gap*, which is what, generally speaking, guidance aims to fill.

Examples of guidance occur in many aspects of our lives. One that we experience almost daily (and that shares similarities with data analysis as well) is driving cars. When we drive, the car itself provides us with guidance. The system that assists us is typically called a navigation system. Thanks to its support, we are able to reach specific places.

The assistance we receive is so specific that the system is able to suggest to us when and where to turn. If the guidance system is sufficiently advanced, it can also inform us about points of interest or weather conditions along the route, suggest to us where to stop to buy some food, and make some deviations to avoid traffic.

Therefore, we can already answer our initial questions: (1) Guidance refers to processes and procedures to support somebody completing a given task. (2) The reason why we

need guidance is that there is a knowledge gap (e.g., we do not know how to reach a specific place). Finally, as we have seen from the examples described, guidance is a complex process: Many intrinsic characteristics affect the kind of support these guidance systems can provide, such as the different kinds of assistance provided by the car.

Returning to the topic of this dissertation, the same observations we made for the car, and for guidance in general, also apply to visual data analysis and, specifically, VA. Performing VA is generally considered a complex task that may be hindered by challenges during any of its phases [Sac+14b]. Analysts should know how to visualize the data in the first place, find appropriate computational models, explore the data, manage the insights, and gain new knowledge. If the user is not able to complete any of these operations effectively, the analysis will typically stall. Hence, analysts could benefit from receiving guidance to overcome such issues and complete the analysis successfully. From this last statement stems the motivation for the work we pursued in recent years, as the first lines of this section summarized: We investigated the fundamental characteristics of the process of guiding users in VA. We studied what effects the provision of guidance could have on users solving VA tasks. Finally, we formalized a methodology to design effective guidance in VA. Our work is summarized in the following cumulative thesis of three journal articles (see Section 1.10 for a full list of articles included in this thesis).

1.2 Objectives and Research Questions

The aim of this dissertation is to answer the following main research question:

“How can we devise guidance methods for supporting users performing Visual Analytics tasks?”

This main research question can be further articulated into a number of sub-questions:

- S1:** Is it possible to devise a general framework and a common guidance definition embodying the current state-of-the-art approaches and literature?
- S2:** What are the benefits (if any), and in general what are the effects of using guidance during visual analytics?
- S3:** How is it possible to design effective guidance to support users throughout the visual analytics process?

Visual Analytics can help make the visual analysis of complex datasets a success [Sac+14b]. However, usually the complexity of the data is directly reflected by the VA tools that aim to support the analysis of that data [Liu+18]. The result is that sometimes the analysis is hampered due to issues arising during the process. Guidance can make this situation significantly better.

Scientific Scope. This dissertation is structured in two main parts. Part I provides an overview of the topics we covered and provides a "plot" that unifies our findings under a common story line. Part II is mainly constituted by the three journal articles we published on the topic of guidance, which compromise the structure of this dissertation.

As often happens in science, when we first approached the fascinating research of guidance methods, we conducted a detailed analysis of the literature. We looked at multiple works and analyzed how different authors tried to improve VA by solving specific issues arising during the analysis process. Based on such initial work, we could summarize the fundamental characteristics of guidance in VA and condense them in a semi-formal definition. In the upcoming sections, we summarize such analytical work and provide additional background notions necessary to understand this thesis. In Section 1.4, we introduce the definition of guidance in VA and describe its main aspects. This work is also presented in its entirety in Chapter 2 which is also our first publication on the topic [Cen+17b].

Our initial definition focuses primarily on the system-side of guidance, meaning that, when we started, we mainly considered analytical systems providing guidance to the user. However, as we made progress in our research, we extended our initial characterization to the user-side of the guidance process as well. In other words, we considered not only the guidance provided by the system to the user, but also the guidance the user provides to the system in return. This kind of guidance can be exploited to steer the course of the analysis and is typically mediated by the user's interaction with the interface. This refinement allowed us also to rule out what does not constitute guidance in the first place. The outcome of this research is mainly summarized in Section 1.6 and is partly based on a state-of-the-art report we published that summarizes the literature on guidance provided by either the system or users [CGM19a] of the last 30 years. We complement our discussion with decision-tree based methodology to decide if an approach constitutes guidance in the first place [Cen+18a].

Following our guidance characterization, we investigated the effects of providing guidance in VA. In particular, we tried to capture how users felt when supported with guidance in contrast to how they behaved without any form of support. In other words, we investigated whether the analysis was more pleasant and profitable for them with or without guidance [Cen+18b; CGM19b]. The outcome of this research is described in Section 1.7. The journal paper we published on the topic is entirely reported in Chapter 3.

Finally, in a more recent phase of the PhD program, we investigated the meaning of *effective* guidance in VA and listed what qualitative requirements are necessary in a VA tool to obtain effective guidance support. We structured our findings into a framework to support effective guidance design. We discuss the main questions to be answered during the development of guidance and possible problems that might arise during this process, together with a discussion of possible mitigation strategies to counteract them. The outcome of this research is summarized in Section 1.8, and it is reported in its entirety in Chapter 4.

1.3 Conceptual Grounding and Methodology

Before getting into the details of this dissertation, we briefly discuss some basic notions and concepts that will be useful for contextualizing guidance in VA. In the following, we introduce information visualization and visual analytics itself and describe the research methodology we used throughout the PhD program.

1.3.1 A Need for Visual Data Analysis

In the past decades, following the birth and growth of the World Wide Web, we witnessed huge advancements and the widespread adoption of digital technologies. The internet, smart devices, and social networks are all direct products of such progress.

Today, digital technologies are pervasive, and their use has become an integral part of our life. We use them at work, as many companies have transitioned to a fully digital workflow. At home, we might ask our virtual assistant to set the thermostat temperature. We use it for travelling when we compare prices and buy tickets. We use it to buy groceries. The examples go on and on and demonstrate that digital technologies are pervasive in every aspect of our lives. Among the many consequences and implications of these technological advancements, one worth mentioning for the scope of this thesis is that — thanks to such advancements — we all have an unprecedented, 24/7, facilitated, and immediate access to unlimited content and data.

On the one hand, it is clear that the access to unlimited content has changed our lives and, in many aspects, made them better. The availability of all this information is useful, for instance, for making informed decisions; it has also allowed the development of new markets and products. However, as with many things, benefits come at a price. Usually, information is not immediately available, but hindered in the data, which needs to be processed in conscious ways to be of some use and value. The dimensionality of data also constitutes a challenge: Enormous amounts of data can be overwhelming for a human mind. It is well known that we as humans can process only a limited amount of data [Mil56]. Data dimensionality poses challenges for computational systems as well [LJ12]. In other words, before making any decision, insights must be gained and the data processed and explored, which is a challenging task. Finally, the personal abilities of the analyst also influence the outcome of the analysis.

Data analysis can be challenging. Hence, in the past the research community assisted with the proliferation of analytical methods that try to make the best of the data. Data analysis can be seen as a process in which at least two agents (of which at least one is typically human and one is a computer) join their efforts to achieve a common analytical goal [Ter95; Tuk77].

We can condense the goal of data analysis into the generic aim of gaining insights. Consequently, the research in this area focused on either improving the analysis for the users so that they can more easily reach those findings or developing efficient algorithmic solutions to extract such findings in a more automatic way.

1.3.2 Information Visualization

A typical user-centered solution to data analysis is information visualization (InfoVis). InfoVis is commonly known as an effective way to make sense of complex data. The aim of InfoVis is to define visual artifacts, that can be visualized, for instance, on the monitor of a computer, representing any abstract concept or process, or generally, data [Heg04]. Using visualizations has the benefit of alleviating the burden of the analysis on the human working memory and reversing it to an external medium—in this case, visualization. In other words, thanks to visualizations, users can reallocate their cognitive resources, reason about the data in an easier way, and concentrate better on the resolution of analytical tasks.

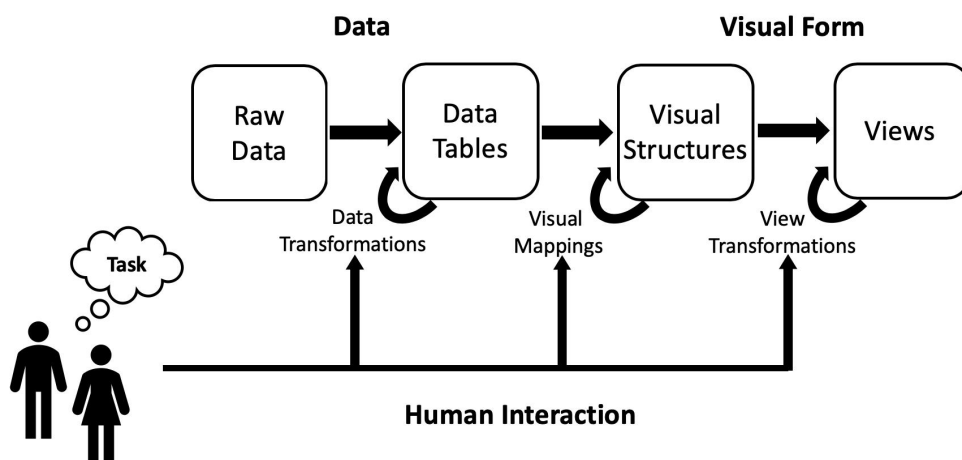


Figure 1.1: The Information Visualization Reference Model by Card et al. describes how data can be transformed into an interactive visualization. Reproduced after [CMS99]

The InfoVis reference model was first introduced by Card, Mackinlay, and Schneiderman in 1999 [CMS99]. The model describes how (abstract) data can be transformed into a (interactive) visual form to solve a given task.

InfoVis is a step-wise process. In the initial phase, the data has to be transformed into a computer-readable format. Additional operations include the annotations, preprocessing, and quality assessment to make it ready for the analysis. This first step is typically named data transformations. In the following step, the visual appearance of the data has to be chosen. The particular aim of this phase is to decide the visual marks and the retinal properties that should be associated with the data dimensions and be used to visualize the data on the screen. This is referred to as the visual mappings step. In the final step, called view transformations, the different visualizations are combined together, arranged, and transformed to make the information necessary for the task readily available. The role of interaction, as can be seen in Figure 1.1, is central to and

pervasive of the whole process. The user/analyst can influence the outcome of each phase of the reference model.

Throughout the years, non-visual methods for data analysis have also been developed. In parallel to the spreading of digital technologies, the society have assisted in many technical improvements that brought us faster memories, larger storage space, and increased computational power, which further allowed their use for analysis purposes. In this regard, data mining [FPS+96] is probably one of the most known machine-centered methods. Compared to InfoVis, which is mainly focused on studying how to visualize data and allow for interactive analysis, data mining can be seen as an algorithmic solution — a black box producing data models and reports that have to be read and interpreted by the analyst.

1.3.3 Visual Analytics

In addition to InfoVis and data mining, during the first years of this century, the research community assisted in the development of a third way to enable data analysis by combining the strengths of visualizations and computational modelling. This new analysis paradigm was given the name of visual analytics. VA aims to combine the power of modern computers with the cognitive abilities of the human mind [Kei+08a] to gain insights and generate knowledge. In other words, the end goal of VA is to enable successful data analysis through a so-called mixed-initiative process, in which the system and the analyst collaborate and share a "dialogue"¹ to complete analytical tasks.

VA was initially characterized by Thomas and Cook [TC05]. Following their work, Keim et al. [Kei+10] presented the reference model of VA, as depicted in Figure 1.2. As it can be seen, the upper part of the model resembles the typical InfoVis pipeline (see Figure 1.1 for a comparison). In other words, the traditional visualization pipeline is combined with the power of data modelling, represented in the diagram by the models process. In addition to the original InfoVis model, visualizations can be used to support model-building activities or visualize an already existing model of the data. Thanks to the additional possibilities offered by computational models and their interconnection with visualizations, the user can achieve results in an efficient way.

A few years later, Sacha et al. [Sac+14b] extended the VA reference model and described how new knowledge can be generated. The new model, shown in Figure 1.3, introduces three main analysis phases. The *Exploration Loop* describes how findings are extracted from the data thanks to exploration activities. Next, in the *Verification Loop*, the findings are grouped together to prove or disprove working hypotheses and form insights. Finally, in the *Knowledge Generation Loop*, the insights are condensed into new knowledge.

¹The roots of the word dialogue come from the Greek words dia and logos. Dia mean "through"; logos translates to "word" or "meaning". In essence, a dialogue is a flow of meaning. The research and improvement of the human-computer interface is at the core of Human-computer Interaction and, hence, VA [GS86]




Figure 1.2: The VA reference model by Keim et al. aims at combining the strengths of computational systems and human perceptual abilities to make the best of visual data analysis. (Figure taken from [Kei+08b], ©2008 Springer. Used with Permission)

Strictly connected to the knowledge-generation model is the sense-making process described by Pirolli and Card [PC05]. They described the process of making sense of the data through a number of operations, including information gathering, transformation of the information into a schema, generation of insights, and finally new knowledge. The first half of the process is called the Foraging Loop while the second half is the Sense-making Loop. It is easy to associate and compare this model to the three phases of VA described by Sacha et al. [Sac+14b].

Despite their vast adoption and effectiveness, one of the open problems with VA is that the mixed-initiative process VA aims to enable is rarely realized. Typically, the whole burden of the analysis is reversed on the user in its entirety. The user/analyst is required to possess enough expertise and skills to pursue the analysis while exploiting the set of features offered by the tool. What happens is that the user is typically overwhelmed by the number of visualizations he/she can choose, the analytical methods, or the analysis strategies to follow. The consequence is that the real effectiveness of VA results is limited; thus, the analysis stalls, and knowledge is hardly generated. This is where the adoption of guidance could provide great benefits.

1.3.4 Methodology

When designing (as well as when evaluating) VA approaches, appropriate considerations should be made about the context in which the analysis is pursued.

In recent years, in addition to characterizing guidance approaches from the theoretical point of view, we also put them into practice, which required us to design, implement, and later evaluate specific VA approaches enriched with guidance. In this regard, we approached this challenge following a standard user-centered methodology.

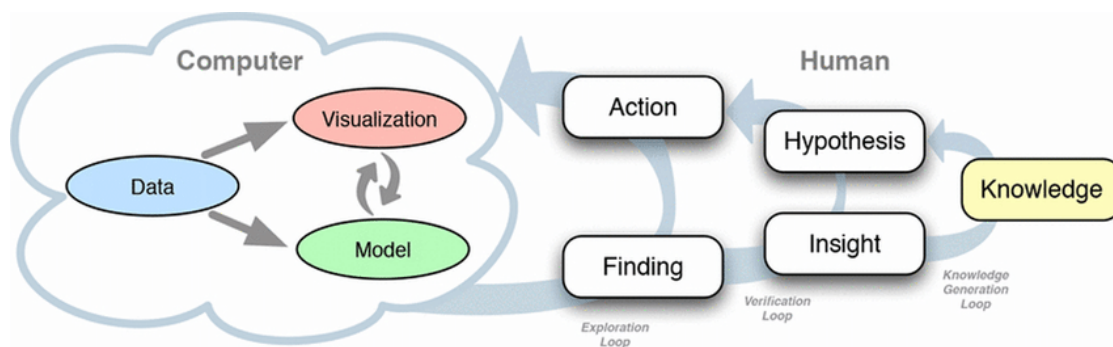


Figure 1.3: The Knowledge Generation model by Sacha et al. describes how knowledge is generated thanks to VA. (Figure taken from [Sac+14b], ©2014 IEEE. Used with Permission)

The design of VA approaches is a process that is entirely developed and constructed around the user. A phrase commonly heard in our field is *"the user in the loop"* [Hor99b], which aims to highlight the importance of the interactive analysis and stimulate the design of approaches that will occasionally ask and exploit users' feedback to steer the analysis process. More recently, this concept evolved into *"the user is the loop"*, where the focus is pushed toward recognizing users' work processes and seamlessly fitting analytics into this existing interactive process [End+14].

Data-Users-Tasks. When designing VA and visualizations in general, it is good to reason about (1) what data we want to analyze, (2) who are the users who will use the tool, and (3) what are the tasks the users wish to solve (see Figure 1.4). Miksch and Aigner [MA14] stated that considering these aspects will largely determine which visual representations, analytical means, and interaction methods are suitable in a given analysis scenario (see the design triangle in Figure 1.4). They also complement the design methodology with three main qualitative requirements to guide the choice among alternative design options. According to the methodology, different visual design alternatives should be rated according to their expressiveness, which refers to the ability of the visualizations to show the exact information contained in the data [Mac86]. Effectiveness is a further quality that should be enforced; it primarily considers not only the degree to which visualizations address the cognitive capabilities of the human visual system, but also the tasks at hand, the application background, and other context-related information to obtain intuitively recognizable and interpretable visual representations [Mac86]. Finally, the last qualitative criterion regards the appropriateness of multiple VA solutions. This metric involves the consideration of a cost-value ratio in order to assess the benefit of the visualization process with respect to achieving a given task [Wal+18; Van06].

The Nested Model. Munzner's Nested Model is a further example of user-centered design methodology typically applied in VA scenarios [Mun09] (see Figure 1.5). Munzner describes a practical framework that illustrates how it is possible to design and implement

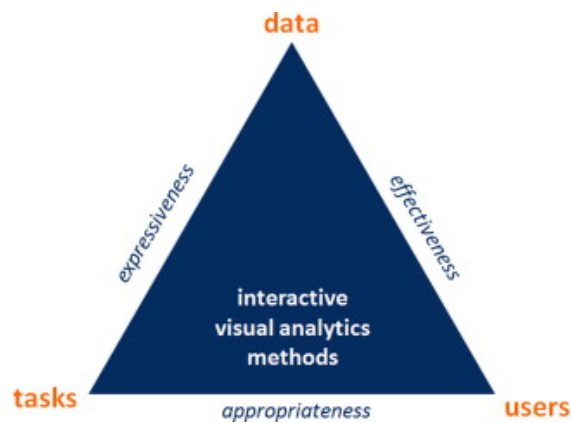


Figure 1.4: The design-triangle by Aigner and Miksch describes the fundamental aspects influencing the design of VA solutions. (Figure taken from [MA14], ©2014 Elsevier. Used with Permission)

visual representations of data starting from the analysis of prerequisites. The design model is structured into four nested steps, the first of which aims to characterize the domain problem. In the second step, the design maps the problem characteristics to abstract data types and operations. In the third step, actual visual encodings and interaction techniques are chosen. Finally, in the last phase, the algorithms to implement the selected visual representations and interaction dynamics are selected.

This framework is fairly useful in that it not only supports the design, but also allows the constant evaluation of the outcomes of each design phase while suggesting countermeasures to common pitfalls. For instance, issues may occur if the problem is not correctly characterized or a proper encoding is selected. To counteract this issue, the aforementioned methodology suggests to perform users' interviews and to choose visual marks according to well-known aesthetic criteria [Mac86].




Figure 1.5: The nested model for visualization design and evaluation described by Tamara Munzner. (Figure taken from [Mun09], ©2009 IEEE. Used with Permission)

Evaluation Methodologies. A critical step during the development of VA approaches is their constant evaluation. Typically, design and evaluation proceed in alternating

cycles, similar to any software development process. The evaluation itself is pursued continuously, in an iterative fashion, while the visualization is developed in order to identify problems in early stages. Evaluating VA is considered a challenging task as it involves dealing with users and understanding how VA could potentially affect their analysis workflow.

Similar to what happens with VA approaches, the real effectiveness of guidance should also be evaluated with a user-centered methodology. Theoretically, the evaluation of guidance adds complexity on top of the evaluation of normal visualizations. When evaluating guidance, we should consider not only the goodness of visual representations, but also their intertwining with the guidance process in order to determine how, and to what measure, the different parts contribute to the real effectiveness of the VA tool. To the best of our knowledge, no specific evaluation methodology exists that describes how guidance methods should be evaluated. However, as guidance is strictly integrated into the visualizations, we can assume that the general criteria applied to the evaluation of "normal" VA approaches still hold true for guidance. The evaluation of guidance methods constitutes a formidable challenge for the development and adoption of guidance in VA.

Lam et al. [Lam+11] conducted a thorough review of evaluation methodologies in visual data analysis, considering evaluation methodologies whose scope goes beyond the evaluation of usability and visualization expressiveness. They categorized evaluation approaches into two groups: process and visualization. The first group includes all evaluation approaches aimed at understanding the analysis process in its entirety and the overall role of visualizations. Thus, the goal is to examine the visualization environment as a whole. The second group includes all the approaches that dig into the details of the specific visual representations used in order to compare them among each other or understand how they perform compared to similar techniques.

These two macro-groups are further subdivided into seven evaluation categories:

1. *Evaluating Environments and Work Practices* focuses on design requirements and on understanding the work or practices enacted by a given group of people.
2. *Evaluating Visual Data Analysis and Reasoning* deals with evaluating if and how visualizations contribute to the generation of new knowledge in a specific domain.
3. *Evaluating Communication through Visualization* focuses on testing the capacity of visual approaches to support communication; it is typical in learning environments and for evaluating teaching efforts.
4. *Evaluating Collaborative Data Analysis* techniques studies the fitness of visual approaches to support collaboration.
5. *Evaluating User Performance* approaches studies if and how specific features of a visualization tool affect user performance objectively and in a measurable way.
6. *Evaluating User Experience* studies people's subjective responses and opinions to visualization environments.
7. *Evaluating Visualization Algorithms* is useful for understanding how visualization performs in a given circumstance.

In this dissertation we mainly evaluated our results and the implemented guidance solutions focusing on measurable performance improvements due to guidance. We also evaluated users' experience when receiving guidance and tested how much guidance contributed to generating new knowledge from the data. These evaluations correspond to the categories (2),(5), and (6).

In particular, in the first work, we investigated how different types of guidance affect users with different knowledge and expertise when solving domain-dependent or exploration tasks [CGM19b]. We applied user-centered methodology to test not only whether guidance influenced the performance (mainly correctness of the answers and time to task completion) of such users, but also how guidance stimulated the growth or contributed to attenuating feelings of frustration due to the user experiencing issues during the analysis. In the following work, we focused on a specific type of guidance — namely, orienting guidance [Cen+18b] — and tested how, thanks to the guidance support, the user was able to gain insights more easily compared to a normal exploratory analysis. This also allowed us to study how guidance can affect and completely change the resolution strategies of different users [Cen+18b].

1.4 Defining Guidance

Outline *In order to answer S1 – "Is it possible to devise a general framework and a guidance definition?" – we performed a comprehensive literature review to understand the different nuances of the term guidance, not only in VA but also in related fields. From a practical point of view, we also analyzed how different types of guidance have been implemented in the literature. We condensed all our findings into a definition and a conceptual model of guidance. We describe the working principles shared by any guidance approach [Cen+17b]*

Achieving effective system–user integration as required by VA approaches is still a largely unresolved task. This happens mostly because VA requires effectively combining multiple aspects of the data analysis process, such as a good understanding of the tasks to be supported, the set of computational methods needed, the knowledge users possess before starting the analysis, and the visual means to be utilized, just to name a few. In past years, great efforts have been made in this direction. Scientists have sought to tackle selected aspects of the knowledge generation model [Sac+14b] (see Figure 1.3). An interesting point of view on VA approaches is provided by Bertini and Lalanne [BL09], who described scientists' attempts to enhance visualizations with computational methods and, conversely, support mining approaches with visual means. However, they stated that an effective integration, in which the affordances of human and computer are balanced and effectively combined, has not yet been accomplished. In particular, they claim that opening the black box of data-mining techniques, which should allow users to steer the process, is far from being solved.

In this dissertation, we propose guidance as a promising attempt to enable a better collaboration of the human and the computer.

The goal of guidance is to facilitate their dialogue and, hence, improve the quality of their collaboration. However, how such a goal can be achieved has long been open to debate. The problem is that the process of guiding users has many facets and nuances. There exist approaches providing different types of assistance, encoded and provided in a wide variety of ways and for multiple purposes.

In the past, such approaches were not even labeled directly as guidance, but rather by "support," "assistance," or "recommendations". This contributed to the confusion on the topic. However, if we look at the details of the single techniques it is possible to identify common aspects of guidance. Schulz et al. [Sch+13] were the first to attempt bringing together the nuances of such approaches, proposing the umbrella term "guidance" to group them. We built on and expanded their initial work by proposing a general characterization of guidance approaches in VA.

Before presenting our definition of guidance in VA, however, we make a small step back and describe the process through which we went to arrive at it. We start reviewing how the term guidance is used in general in VA and related fields.

A first reference to guidance can be found in fields closely related to VA. In Human-computer Interaction (HCI), Smith and Mosier [SM86] emphasized the importance of guidance, which they defined as a "pervasive and integral part of interface design that contributes significantly to effective system operation." Guidance is also cited among their guidelines for designing visual interfaces. Dix et al. [Dix+04] stressed the need to incorporate guidance in visual analysis interfaces. Their discussion started from the consideration that, among the users of a certain analysis tool, there will inevitably be someone who will have issues using it, which happens because each user has a personal variable degree of knowledge that, in some cases, does not match the knowledge required to operate the analysis tool. Thus, according to them, guidance is essential to fill the knowledge gap.

Still in HCI, Engels [Eng96] analyzed guidance discerning its main components and dimensions. He affirmed that guidance techniques are always characterized by a (1) "what" that clarifies what the problem to be solved is. Engels further stated that the "what" should be decomposed into an "initial state," characterizing the state at the beginning of the analysis, and a "goal state" that has to be reached. The second component of guidance is (2) the "how" that aims to solve the discrepancies between the initial and goal states. In other words, the "how" describes the functioning mechanisms of guidance.

Finally, in the visualization literature, we already introduced the early work on guidance by Schulz et al. [Sch+13]. In their thinking, guidance refers to methods that have the goal of providing dynamic support to users, such as guiding data exploration or assisting users when choosing visual mappings for presenting analysis results.

From all the different interpretations in the various fields, we derived the following definition of guidance in the context of VA:

Guidance is a computer-assisted *process* that aims to actively resolve a *knowledge gap* encountered by users during an *interactive* visual analytics session [Cen+17b, p.2].

Let's focus on the three emphasized phrases in the definition: Guidance is meant to be a *dynamic process* that aims to support users solving particular tasks. The resolution of tasks is being hindered by a *knowledge gap*, and it is the goal of guidance to solve. The third word underscores the *interactive* nature of guidance in that it is a reaction to users' actions.

We can see how our definition is strongly grounded in HCI. In particular, the definition recalls Engels's statement that the solution of any task can be decomposed into a series of actions that lead to the completion of the analysis [Eng96]. Building on this consideration, the guidance process provides support for at least one of these actions in situations where a user is unable to identify, judge, or execute the action itself.

Finally, it is important to note that the definition focuses on the system-side of guidance in that it mainly considers systems guiding human actors [Hor99a]. We will investigate the user perspective of guidance later in Section 1.6.

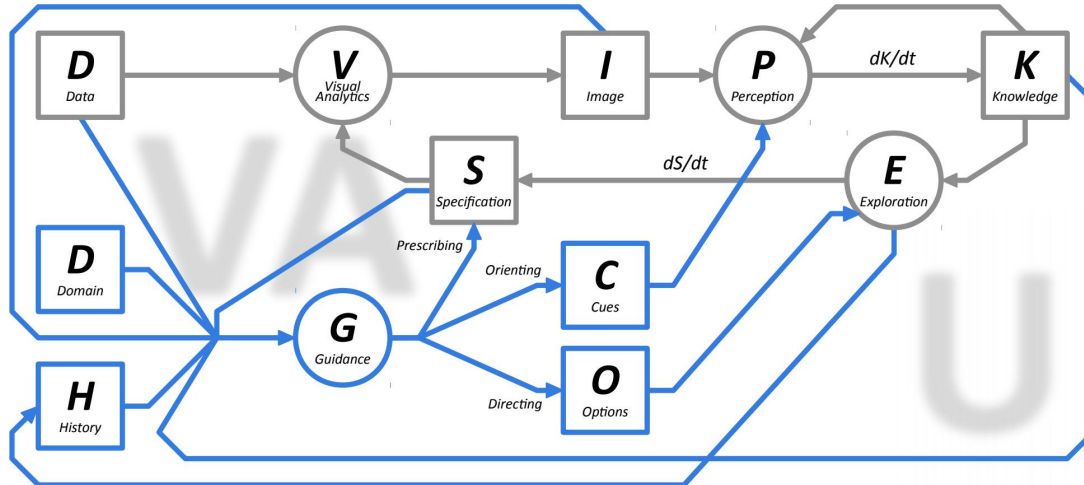


Figure 1.6: The conceptual model of guidance in VA, discussed in Chapter 2. The model was created extending the van Wijk's [Wij06] model (in gray) for visualization with guidance-related blocks (blue). As in the original model, the system aspects of guidance are visualized on the left of the scheme while user aspects (U) are on the right. Figure reproduced with permission after [Cen+17a]

1.4.1 Conceptual Model of Guidance

In addition to the definition, we structured our findings in a reference model of guidance. Similar to the VA model [Kei+08b] or the visualization model [CMS99], our model aims to show the fundamental mechanisms of guidance in relation to the visualization process it seeks to assist.

We based our model and its description on the visualization model by van Wijk [Van06], which we adapted and extended to include guidance. In doing so, we adopted the same formalism he used in his original paper. Boxes represent sources of input for generating guidance or the output produced by active guidance processes, while processes are represented by circles. Thus, boxes represent artifacts and circles represent functions (see Figure 1.6). Initially, van Wijk's model was created to represent only the visualization process. Consequently, it does not directly convey the special characteristics of VA. In this, it is more similar to the model by Pirolli and Card [CMS99] than the VA model by Keim and Sacha [Kei+08b; Sac+14b]. To show that guidance can be applied to VA too, we expanded the visualization model to include the analytical processes. We did not change the appearance of the model, though. We only varied the semantics of the "specifications." In our scheme, the [S] block represents not only the specifications of the visualization, but also those of the algorithms/methods to analyze the dataset and their parameters.

As in the original model, gray boxes and circles represent an abstraction of the VA process in its entirety. To that, we attached new guidance components, shown in blue, to represent the specific aspects of guidance and their relations. For additional details, refer to Section 2.4, which presents a detailed analysis of the different model components.

1.4.2 Guidance Characteristics

From the discussion thus far, we can clearly state what the main characteristics of guidance are: (1) *what we refer to as the knowledge gap*, which represents what the guidance process aims to solve. As with all processes, guidance can also be characterized by (2) *an input and an output*. The input consists of a list of resources the process could exploit to generate the necessary assistance. The guidance process also has an output that can be thought of as the computed answer to the user's knowledge gap together with a set of visual means to communicate the answer itself. The guidance can be further characterized by (3) *a degree*, indicating the amount of assistance provided by the answer. Taken together, these factors constitute the main characteristics shared by all guidance approaches. We will discuss them briefly, also making short references to the literature (see Figure 1.7).

(1) Knowledge Gap

The existence of a knowledge gap is the reason why guidance is needed. It is related to the question: "What does the user need to know to complete the analysis?"

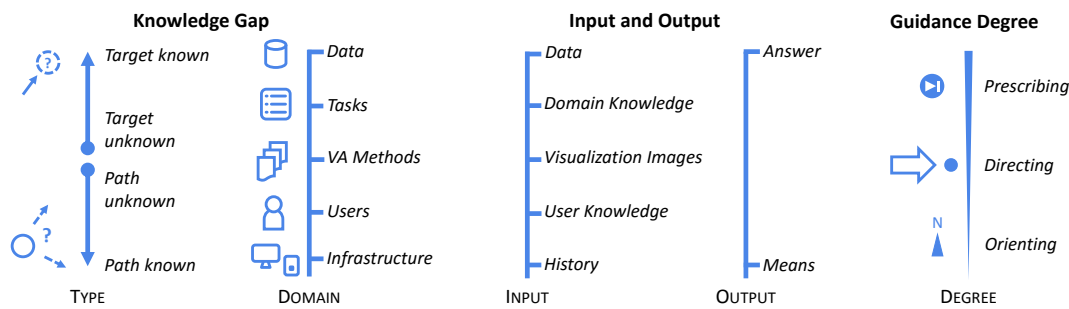


Figure 1.7: Guidance can be characterized in terms of the main aspects: knowledge gap, input and output, as well as guidance degree. (Figure taken from [Cen+17b], ©2017 IEEE. Used with Permission)

The knowledge gap stems from a discrepancy between what a user can do and what he/she needs to do to solve a task. If the user is able to perform the tasks, then guidance is clearly not needed. However, what typically happens is that problems arise at different stages of the analysis.

If we observe specific guidance techniques, we can clearly see that the knowledge gaps can be classified according to the following analytical *domains* to which they pertain.

Data. A group of literature approaches describe mechanisms to cope with problems that arise directly from the management of the data. For instance, Kandel et al. [Kan+11b] described how the user can be assisted in choosing appropriate methods to preprocess the data before the analysis. May et al. [May+11] developed an approach called SmartStripes that aims to provide guidance to support the selection of data features. Gratzl et al. [Gra+14] also described how the user can be supported and guided through a data exploration session.

Tasks. Other approaches aim to support directly the resolution of a specific task. For instance, Streit et al. [Str+12] described an approach that is able to support the resolution of predefined tasks. Based on a detailed and structured analysis of the analysis scenario, their guidance approach is able to detect the current state of the analysis and subsequently suggest to the user what to do next at each phase of the analysis.

Visual Analytics Methods. The peculiarity of VA, which also constitutes its strength, lies in the integration of visualizations and analytical methods. It is typical for guidance techniques falling into this category to support the selection of appropriate algorithms to model the data and their parameters. For instance, Choo et al. [Cho+10] described a system that interactively helps the user classify data. A similar approach is presented by Bernard et al. [Ber+17] to guide the analyst through the selection of appropriate data labeling.

Users and Infrastructure. Finally, for the sake of completeness, guidance can also be provided by the analytical system not only to reach an analytical goal itself, but also for selecting the correct analyst for the job (for instance, one with specific expertise) and

the computational hardware on which to run the analysis (for instance, based on some load and computing power requirements).

A more generic way of looking at the knowledge gap, is focusing on its type. According to Silver [Sil91], the user may need to be supported with guidance either to identify the target of the analysis or to define a set of operations and action to reach it. In this case, we can talk about the following approaches.

Target Known/Unknown. In this first case, the knowledge gap relates to finding the target of the analysis. This target can be any goal and/or subgoal the user may have in mind [BH19]. For instance, a typical case is an exploratory analysis, which typically starts with no clear hypotheses in mind. The goal of guidance is to help the analyst delineate the analysis target and "discover the unexpected". Fujishiro et al. [Fuj+97] described an approach that exploits a knowledge base to suggest to the user what visualization and visual encoding (of the target) are best suited to solve a certain task. Bernstein et al. [BPH05] presented a guidance technique that aims to provide guidance in the selection of appropriate data-mining algorithms (the target). The result is achieved by matching a list of available models with the characteristics of the user's task at hand.

Path Known/Unknown. In other approaches, the knowledge gap is related to the execution of the conceived plan, such as by guiding the determination of a structured sequence of operations that could lead the user to prove or disprove his/her own working hypotheses. This could also include the choice of parameters for each step. An approach that addresses a path unknown problem is the scented widgets by Willet et al. [WHA07]. This guidance technique aims to suggest the operations a user could do to reach a certain target. In addition, May et al. [May+12b] focused on the problem of guiding the user toward interesting regions of huge graphs. In particular, they used visual glyphs to indicate in which direction such regions are located and the shortest path to reach them.

(2) Input and Output

Being a process, guidance can be characterized by an input and an output. The input is related to the following question: *What is the basis upon which we can build guidance?* In other words, the input represents the sources used by the guidance process to generate the assistance.

In general, all kinds of data, artifacts, and knowledge types can be exploited to generate guidance. Once a knowledge gap is detected, the guidance system should find the appropriate matching piece of knowledge or information that, once provided to the user, will allow him/her to continue the analysis. In general, we identify the following sources of input.

Data. A typical source for guidance is the data itself. The analysis system can extract statistics from it, process it with modeling algorithms, and subsequently provide this information to the analyst for guidance, but also suggest where interesting subsets lie. Lex et al. [Lex+12] designed a system that generates guidance out of the data. The algorithms

behind the visual environment they developed can detect and calculate relationships between multiple datasets and subsequently present them for the convenience of the user. The very same strategy is also at the base of the mechanism we used to provide guidance during a user study we designed to evaluate the effects of multiple degrees of guidance on the user. In particular, we were able to extract from the data some statistics that allowed us to suggest to the analyst what to look for during data exploration. Section 1.7.1 provides the details.

Domain Knowledge. A further source of input is domain knowledge. Typically, before being used, this kind of knowledge requires an elaboration phase, called externalization, through which the knowledge (e.g., from experts) is represented and stored in a computer-readable format. For example, Gotz et al. [GW09], in their approach, were able to interpret the actions performed by users on a certain dataset and understand their intent. Thanks to this information, which is obtained by matching interaction patterns with a knowledge base of known tasks, the system is able to suggest modifications to the visual environment that could improve the user’s performance.

Visualizations. Visualizations and, in particular, the information about the visual encoding can be used as a source of information to generate guidance. In the literature, we did not find many approaches in this category. Wang et al. [Wan+16] designed an approach providing guidance exploiting information about the visual layout of a graph. In a typical graph drawing, in order for the layout to be considered good, it should satisfy particular aesthetic criteria, like a minimum number of edge crossings or a sufficient distance among the nodes. In fact, this increases the comprehension of the underlying data. The system designed by Wang et al. generates guidance making use of a derived measure considering the degree of the ambiguity of the drawing. In particular, the visualization system is able to suggest to the analyst what (cluttered) part of the drawing needs special attention.

User Knowledge. The user knowledge can be exploited to produce guidance. The approach by Gotz and Wen [Got+10] makes use of information about a user’s actions to infer the user’s intent and generate the necessary assistance. In general, the user’s knowledge serves as the main mechanism to steer and direct the analysis along the analysis path desired by the user. We will discuss this in detail in Section 1.6, where we provide a detailed analysis of feedback mechanisms.

History. A final source of information can be obtained through the analysis of the user’s action history. Derthrick and Roth [DR00], for instance, draw the analysis paths already covered by the analyst so that they are able to have a comprehensive view of what is missing to complete the exploration of the data. Horvitz et al. designed Clippy [Hor+98], a visual assistant for commercial analysis tools that is able to suggest how the user should continue the analysis by analyzing and modeling the user’s behavior.

On the other end of the guidance generation process is the output. The output answers two questions: *What is the answer to the user’s knowledge gap? How could the guidance*

be conveyed to the user? The output of the guidance generation process is defined in terms of an answer to a specific knowledge gap and the means by which this answer is conveyed to the user.

The Answer. The answer constitutes the solution, or partial solution, to a knowledge gap. In mathematical terms it can be defined as the output of the guidance function:

$$\textit{guidance}(\textit{gap}, \textit{input}) \rightarrow \textit{answer}$$

We can imagine the guidance process as an iterative loop. Multiple iterations of guidance convey, in each round, multiple guidance hints to the user, leading to a variable amount of knowledge. Through iteration after iteration, this will (hopefully) solve the knowledge gap.

We can identify two types of answers. In the first case, the answer is meant to address the knowledge gap directly. For instance, Bernard et al. [Ber+17] designed a guidance tool that provides guidance in directly tackling the user's knowledge gap. Specifically, the user is required to label a certain dataset, and the guidance system automatically suggests a possible labelling. It is easy to see the direct connection between the issue experienced by the user and the solution provided by the system. Similarly, Gratzl et al. [Gra+14] developed a system that highlights connections among multiple datasets to support the user seeking them. Finally, Ankerst [AEK00] supported the user in building a data clustering by suggesting what data should be added to different clusters of the dataset.

Indirect Indirect answers are possible too. These kinds of approaches do not provide a direct answer to the encountered problem, but rather suggest alternative options that might solve the given knowledge gap, but indirectly. For instance, the approach designed by Fujishiro et al. [Fuj+97] tries to support the exploration of a dataset. However, the system does not support the task, for instance suggesting to the user what data to explore or how to explore it. Instead, in an indirect way, it suggests possible visualizations that could support the analysis of that given data.

The (Visual) Means. A factor that greatly influences the effectiveness of the guidance answer and its ability to solve the knowledge gap is how the answer itself is conveyed to the user. The use of visual means is typical. However, we do not exclude that other sources could be used for the scope, like haptic feedback or sounds[McC+18].

As we will see in the next section, the type of visual means is typically connected with the type of support and degree of assistance required during the analysis. However, in general, any kind of visual encoding can be used. For instance, May et al. [May+12a] described an approach utilizing visual glyphs that aim to show where interesting sections of a graph are located (see Figure 1.8d). In addition, the approach by Gratzl et al. [Gra+14] similarly presents the user with glyphs representing the connections among heterogeneous datasets. Changing color is another expedient that can be used to attract the attention

of the analyst where needed. For example, Heer et al. [HB05] developed a tool that uses color codes to highlight communities in large social networks, supporting efficient data exploration.

(3) Guidance Degrees

A final aspect characterizing the guidance process is how much assistance the system is able to provide or, inversely, how much freedom users retain during the analysis. According to the degree of support required, the user's actions are restricted or directed to prescribe a certain line of enquiry. In addition to the absence of guidance, we can identify in total three guidance degrees: prescribing, directing, and orienting guidance [Cen+17b; Cen+18a].

The first two degrees, prescribing and directing, are meant to provide a high level of assistance, with the former focused on the choice of the single most appropriate way to achieve the results and the latter describing the mechanisms to calculate a wide set of options to continue the analysis. In simpler words, prescribing guidance and directing guidance are focused on the provision of suggestions. Orienting guidance, on the other hand, plays its role at a lower and more subtle level, exploiting the user's perceptual abilities to provide him/her with a set of visual hints to push the analysis forward. In other words, orienting guidance is more focused on providing means for the user to understand the answer to the problem at hand instead of providing ready-made answers.

Prescribing. One way of providing guidance, which is also the strongest type of assistance possible, is pushing the user along a promising analysis path. The system, given the analysis context (i.e., the current data and tasks), computes the best way, including in terms of efficiency, to reach directly a satisfactory conclusion of the analysis [Cen+18a].

Horvitz et al. [Hor+98] described the design and implementation of a system that provides assistance to software users. The system exploits Bayesian user modeling to transform interaction into useful hints related to the user's intentions. In particular, the system is able to infer the different phases of the analysis, the tasks, and the user's needs and subsequently provide suggestions to continue the analysis and pursue a (inferred) task. Chen and Scott [CBY10] automatically calculated annotations of data snippets selected by the user. The user can directly modify the annotation, which again affects the generation of future annotations. Ip et al. [IV11] guided the user through the visualization of large images by calculating and providing a step-by-step exploration of the most promising and interesting views.

Directing. At a lower level of support we find directing approaches. The name suggests their main purpose: directing the analysis. In fact, approaches that provide this guidance degree aim to solve the user's knowledge gap by presenting a set of alternative options to continue the analysis. The given suggestions/options could differ in terms of quality and costs for different paths, leading to the same result or, in terms of interest for paths, leading to similar or new targets. When compared to prescribing guidance, the options

provided by this guidance degree are higher in number and differ in quality. For instance, an analytical system, based on an interestingness indicator, may suggest to the user a set of alternative data cases that may be useful for the analysis or provide a set of alternative interaction steps. Although it is clear that the user's freedom is greater (given the higher number of options provided), this guidance degree may also introduce a certain level of uncertainty and provide suggestions not directly related to the tasks in focus.

Directing guidance is also strongly tied to recommender systems. Recommendations assume different forms according to the different goals for which they are created and the analysis step they should support. To support data transformation, the system simply suggests to the user the most suitable functions to modify the data [Kan+11b], clean and polish the data [Kan+12], or to support feature selection for data profiling [May+11]. The same happens for data visualization, where directing guidance translates into systems that provide suggestions for different visualization alternatives, usually ordered on the basis of specific perceptual characteristics [Won+16; BGV16; Fuj+97]. To support data modeling, directing guidance usually provides the user with different algorithms and parameters [DFB11; AEK00]. However, this guidance degree assumes a particular interest when it allows the steering of the whole exploration process, pointing users to interesting findings [Joh+05; May+12b; Str+12].

Orienting. Maintaining the user's mental map and orientation is a fundamental goal of any visualization tool. This importance has been recognized in various studies [PHG07; AP13]. With the term orientation, we refer to the structural cognitive information a user creates internally by observing an image, which represents the user's underlying understanding of the information. Hence, sustaining context comprehension and improving the user's orientation during the analysis influence the user's perception of the dataset and the tasks. Orienting guidance can be provided to the user by exploiting low level information extracted from the dataset, in which case this information is mapped to basic perceptual properties to guide the user, or by exploiting users' interaction to support the analysis by providing suggestions (see Figure 1.8). Unlike directing guidance, the suggestions do not have a clear order or priority. In the following discussion, we distinguish the different orienting approaches according to the visual properties used to support guidance and the kind of suggestions they provide. The two groups are not mutually exclusive, but refer to different aspects of the same process.

Highlighting and removals. These approaches play with the preattentive skills of human perception to provide guidance. Contrasting the color hue and intensity of important elements with those of the surroundings allows the users to quickly and preattentively identify them, without the need for a longer sequential search [Maz09; Sch+17].

Vizster [HB05] signaled the presence of communities in social networks using color changes. Nodes and links among such communities are highlighted, and color is used to encode distance. To guide feature subset selection, May et al. [May+11] color coded interesting data columns. Color change is also used to relate selected data features to the overall quality measures, in this way highlighting causal relationships. Ankerst et al. [AEK00]

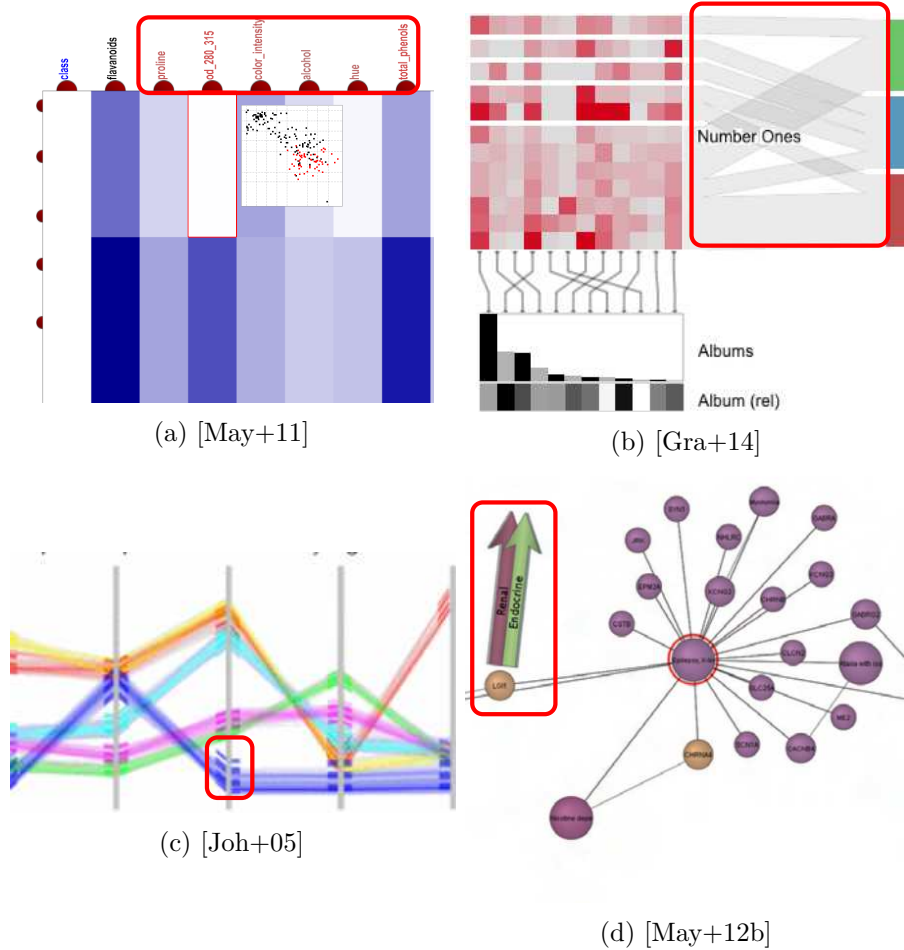


Figure 1.8: *Orienting guidance* can be provided by encoding the information the user needs by using different *perceptual properties*. Highlighting (a) can be used to make visible information that could have some interest for the user, for instance data columns (highlighted in red). Changing the layout and using forms for different data subsets may stimulate the user to explore them. (b) Different data subsets are connected with visual ribbons (highlighted in red) to signal interesting relationships. Motion can be a way to convey guidance too. (c) Motion is used to signal that the data under analysis has a specific characteristic. The reader can imagine the blue lines highlighted in Figure (c) moving up and down. (d) In general, the end-goal of orienting guidance is provide suitable *suggestions* to proceed the analysis. The system suggests interesting graph regions to explore next (the arrow highlighted in red). (Figure taken from [CGM19a], ©2019 Wiley. Used with Permission)

presented a visual technique for building decision trees. Possible split attributes and split points are highlighted to steer the building process. Similar techniques adopt highlighting for classification. Data points for which a given label has not yet been assigned, or for

which the classification is uncertain, are presented in a different color [GRM10; MW10].

Layout and Form. The 2D position, spatial grouping, and marks are properties that our eyes perceive faster and that attract our attention [Tre85]. Guidance approaches use, for instance, links to signal relations or the positioning of elements to suggest to the user that a (hopefully) better layout can be obtained. Usually this is achieved by means of (user-defined) metrics that express the user's intentions and goals. Closeness is often used to signal the belonging of a certain point to a cluster [Mül+08]. However, when uncertainty is involved, it still might not be obvious to the user which cluster to choose. To support this task, Choo et al. [Cho+10] visualized links among the data elements and the most appropriate cluster. Similarly, but at a higher level of abstraction, Stratomex and Domino [Lex+12; Gra+14] offered two approaches that present the users with relationships between different datasets using glyphs (e.g., ribbons). Building a story is important to flawlessly compose connections among facts and events. Thus, guiding story-building activities is the aim of the approach developed by Hossain et al. [Hos+11]. In particular, relationships among documents are shown for the user's convenience.

Motion. Flicker and motion are also important preattentive visual features. They are very useful to quickly attract the user's attention. This is why they are frequently used in our daily life, such as in commercials and traffic lights. However, we could find just one approach exploiting such visual properties for guidance. Johansson et al. [Joh+05] used animation and textures to direct the user's attention to important characteristics of the data. Aiming to show clusters in a high density parallel coordinates plot, they animated different lines with differing phase velocities to emphasize the skewness or variance of data clusters.

Suggestions. In general, providing orienting guidance is achieved by presenting suggestions. The system utilizes a complex combination of the expedients described above (highlighting, glyphs, etc.) to provide analytical options so that the user can proceed toward her/his goal. Unlike directing guidance, these suggestions have equal weights (i.e., are not sorted according to importance); therefore, the resulting guidance degree is considered lower.

Using a flexible degree-of-interest function, the system developed by May et al. [May+12b] can produce recommendations regarding data subsets that are worth investigating. In particular, the system supports orientation by pointing the user to graph regions outside the active exploration area by means of visual glyphs as well as a possible shortest path to reach that region. Luboschik et al. [Lub+12] facilitated the exploration of multiscale data. The approach points the user to scales and regions within the data that exhibit specific behaviors of interest so that the user does not need an exhaustive search. The system aggregates the data of the finest granularity into more coarse-grained data. Consecutive data scales are compared, with the important data characteristics of the fine-grained data being preserved and visualized in the overview so that the user can have a view of the most important characteristics without the need for panning and zooming operations. Jiang et al. [JN15] developed a method to support the creation of queries. The system calculates the relevance of the query with respect to what a user is doing (i.e., the interaction) and suggests appropriate parameters accordingly.

Minor Guidance. We dedicate a separate part to the description of those borderline approaches that comprise very small elements of guidance. Usually, these approaches cannot be considered regular guidance techniques, but still they utilize a design rationale that is interesting from a guidance perspective. These approaches are signaled by an empty circle in Table 1.1.

Undo/redo of actions and history of visualizations are common practice [Men+98; DR01; KNS04]. Usually, such approaches cannot be considered as guidance approaches as they present almost static visualizations of the interaction history. Guidance, on the other hand, is a dynamic process focused on the future of the analysis. In this context, the approach proposed by Sarvghad and Tory [ST15] differs from a common history function in that the approach they propose is able to relate the exploration history with data dimensions and enable users to see which data dimensions have been explored in the past and in which combinations. Hence, it promotes a more complete exploration.

Usually, a static visualization of a model is not considered guidance. The approach by Krause et al. [KPN16], however, is different in that it provides guidance to explore the output of predicting algorithms by showing relationships between the output model and the data features that influenced it. Users can not only understand why certain results are predicted, but also see how the predictive model responds to modifications of the data itself, which also facilitates parameter refinement.

1.5 A Review of Guidance in Visual Analytics

Outline *In this section, we illustrate instances of guidance. However, in contrast to the previous section, we do not focus directly on describing guidance characteristics. Instead, we focus on the mechanisms for generating guidance. We analyze how guidance is actually produced in different scenarios. Table 1.1 summarizes the main findings of our literature review. This section cover parts previously published in [CGM19a] ©2020 Wiley and sons. Reused with Permission.*

If we look at the literature, no single guidance approach covers all aspects discussed in the previous section. Instead, what we can see are approaches that express single characteristics of guidance. Most focus on solving single knowledge gaps, like choosing appropriate visualizations or supporting a profitable data exploration. How guidance approaches are able to reach their scope has never been studied in a systematic way. In other words, the guidance generation function (see [G] in Figure 1.6) is still a black box today. In an effort to shed light on the dynamics of guidance and provide an additional perspective on guidance approaches, we analyze the state-of-the-art with an emphasis on how guidance achieves its goals.

How can the knowledge gap be bridged?

The ultimate goal of the data analysis is making sense of data by means of exploration, verifying hypotheses, and generating new knowledge [Sac+14b]. As previously mentioned, issues arise when the user's knowledge does not match the knowledge required to fulfill the analysis goal. How the knowledge gap of the user can be solved depends on the analysis scenario at hand.

In an effort to provide an initial, albeit partial, answer to this tough challenge, we describe and mark the steps of the analysis process on which a computational system acts to provide guidance to the user. In other words, we provide a detailed view of the very mechanisms that regulate the G process introduced earlier (see Figure 1.6). In particular, we describe *how* a guidance system reaches its goal of addressing the knowledge gap.

With this goal in mind, we build upon the dimensions of the knowledge-generation process (see Figure 1.3) and characterize what we define as the analysis objective—that is, what the guidance process aims to support in order to address a given knowledge gap.

In relation to each step of the VA process, we list a set of objectives that guidance may help reach. Note that these categories are different from the one we used for the knowledge gap. Our focus is on how guidance is produced and how the knowledge gap is addressed.

- *Data*: The user may need help manipulating the data. This includes all the preprocessing procedures before the data is visualized or analyzed.
- *Visualization*: The user may need help visualizing the data or refining an existing visualization. This includes finding an appropriate visual mapping of the data.
- *Model*: The user may need help creating a model of the data or refining an existing model. This includes finding a correct model of the data or appropriate parameter settings.
- *Exploration*: Although high level activities, such as exploration of a dataset, are mostly a human prerogative, the system may still be able to provide support. This includes supporting the interaction with the system itself as well as with activities, like the definition of an analytical goal or the discovery of new *findings*.
- *Verification and Knowledge Generation*: Other high level activities a system might support are verification and knowledge generation. This includes the provision of means for collecting findings from the exploration phase and connecting them with each other in order to foster complex insights, prove or disprove hypotheses and, thus, generate new knowledge.

Let's look at some literature on approaches in detail.

Data. Approaches falling in this category provide guidance to preprocessing operations that operate directly on the data, such as data wrangling and data cleansing [Kan+11a]. Although the literature covering this first elaboration step is vast, just a few works contemplate guidance. This is usually achieved by means of recommendations and the prediction of appropriate algorithms, parameters, and visualizations [HHK15]. Most of these works stem from the initial ideas of Kandel et al. [Kan+11b] and are nowadays pursued in the context of Trifacta [Tri], a company offering commercial services for data wrangling.

Data Wrangling and Cleansing. On the subject of data wrangling, Kandel et al. proposed Wrangler [Kan+11b]), a visual interactive tool to support data transformation. In addition to the visual design considerations, some aspects of their tool relate to guidance. In particular, Wrangler is able to guide the selection of appropriate data transformations based on the data type and by matching the current data with a shared database of data transformations. Furthermore, in line with the guidance objectives, the provided transformations are not executed automatically, but the user is left with the possibility to modify them according to the specific scenario. These modifications are then used to adapt the generation of future recommendations. On the other hand, Kandel et al. [Kan+12] supported data cleansing. Data cleansing is often a semi-automatic activity, usually based on algorithms exploiting different metrics for determining data quality problems and user actions to consider the different quality issues in the right context. To support this task, Kandel et al. focused on providing suggestions of proper visualizations to compare quality metrics and the corresponding data values. Finally, May et al. [May+11] presented Smart-Stripes, a tool that provides users with the possibility to steer the process of feature selection. Feature subset selection is usually done before the data analysis to extract the most important features from a large multidimensional dataset. In their work, automated methods and user interaction are intertwined to open the algorithmic black box and provide the user with an informative overview of the most interesting features. This is achieved by decomposing the different measures characterizing the data features and relating them to precise data subsets, showing the user the overall influence of precise portions of data on the overall measures and, thus, on the resulting features.

Data Preprocessing. Whereas the previous approaches focus on single tasks, Bernard et al. [Ber+12a] focused on the whole data preprocessing process by providing guidance to compose the steps and procedures necessary to have usable data for analysis (see Figure 1.9). In particular, they aimed to integrate domain knowledge and metrics to guide the imputation of parameters and the selection of appropriate values for the single processing steps. This is achieved by showing promising parameters as well as the possible effects of the single choices on the overall result. The work by Heer et al. [HHK15] also focused on the overall preprocessing procedures and constitutes a good starting point for a better human–computer collaboration in this area. In fact, they did not propose a solution to a specific problem in data transformation, but instead proposed a framework for supporting and guiding the user throughout the process. Their idea is to have a so-called predictive interaction, in which the system suggests possible next steps and the

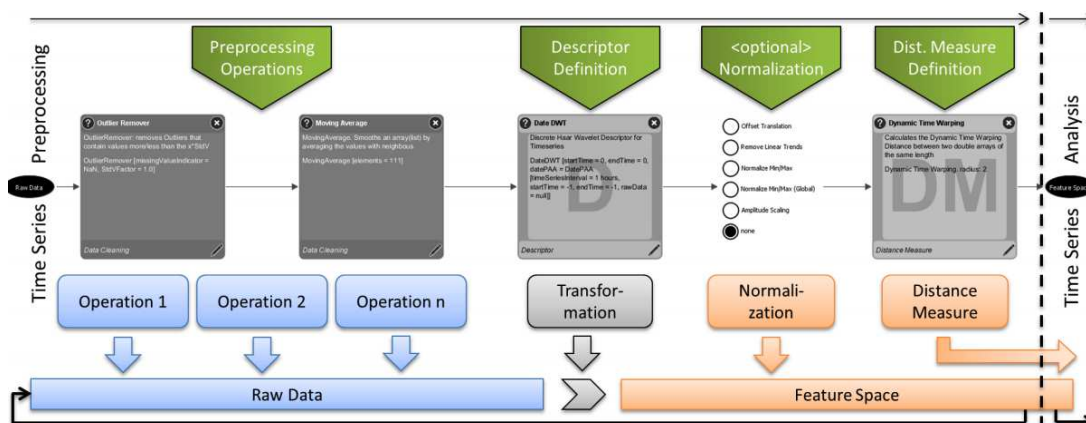


Figure 1.9: Jürgen Bernard [Ber+12b] provide *directing* guidance to the *data preprocessing* step. The approach aims to support the user in choosing the parameters of each preprocessing step. User guidance is derived from user’s *direct* interaction with the tool. This *feedback* is also used to refine the system guidance. (Figure taken from [CGM19a], ©2019 Wiley. Used with Permission)

user selects features (that will influence the generation of future suggestions) and chooses among the system’s suggestions.

Visualization. In this category we outline tools and approaches that aim to support either the mapping of data to visual forms or the visualization of data models. However, just a few guidance techniques are devoted to support the latter scenario.

Visual Mapping. Fujishiro et al. [Fuj+97] developed GADGET, a system that presents the user with possible additions to existing visualizations and complete visual mappings for the user’s convenience. The suggestions are based on the data and on task descriptors as well as on the similarity of the current visual mapping to a database of example mappings. The interaction and the choices of other users indirectly influence the provision of new suggestions. Bertini et al. [BS06] designed an approach to support the visualization of over-plotted areas and improve the overall image quality. This result is achieved through the use of algorithms and metrics that measure quality of the drawing. Through these metrics, the visualization is modified and the over-plotted areas are sampled and unnecessary data features are removed. Koop et al. [Koo+08] presented VisComplete, a system that aids users in the process of creating visualizations by using a database of previously created visualization pipelines. The system learns common design paths and, according to the current user’s input, suggests visual additions. Gotz et al. [Got+10] described a behavior-driven system suggesting a set of visualizations to the user that should be effective for a given inferred analytical task. This work is based on a previous study in which the authors demonstrated the relationship between tasks and user interactions [GW09]. In a similar manner, O’Donovan et al. [OAH15] presented DesignScape, a system proposing

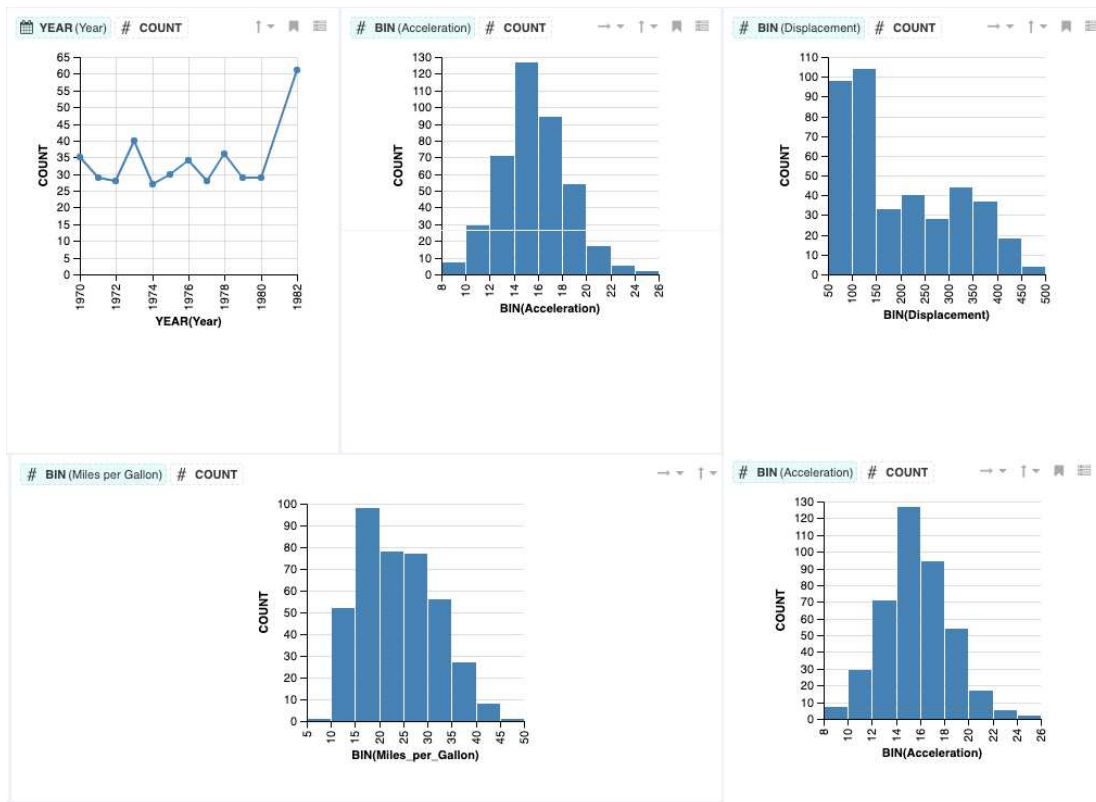


Figure 1.10: Wongsuphasawat et al. [Won+16] provide *directing* guidance to the *visual mapping* step. The approach aims to provide the user with recommendations of suitable visual mappings. Recommendations are ranked based on some objective visual criteria as well as thanks to user's direct feedback (Figure taken from [CGM19a], ©2019 Wiley. Used with Permission)

a set of ordered suggestions to improve the current visual design. Two distinct types of suggestions are available: refinement suggestions, which improve the current design, and brainstorming suggestions, which change the style. Bouali et al. [BGV16] designed a system providing the user with suggestions of proper visual mappings. The user can choose and select the most promising one and provide weights of the most appropriate data columns to be included in the final visualization. Based on these interactions, the guidance algorithm refines the recommendations. Finally, Wongsuphasawat et al. [Won+16] described Voyager, a system featuring a recommendation engine capable of suggesting effective visual mapping, considering both the current data selection as well as expressiveness criteria (see Figure 1.10).

Model Visualization. Zheng et al. [ZAM11] presented a tool that allows the user to visualize the model of a given 3D object. The system suggests informative views based on the results of a clustering algorithm. The views are sorted, and the choice of a suggestion

triggers the recalculation of the rendered scene and the calculation of new suggestions. Ankerst et al. [AEK00] proposed some useful hints for model visualization. Although mainly focused on model building, the system they described supports the visualization of the model, too. The system guides the users proposing them to change visualizations and by offering a look-ahead function to understand how the model will look like in the future.

Models. Approaches in this category deal with supporting the creation and optimization of data models, which is usually achieved by providing the users with the most promising algorithms and guiding the selection of proper parameters.

Data Mining. Most of the literature on this topic is built around classification and clustering algorithms augmented with visual means. They differ mainly in terms of the kind of algorithms involved, the application scenario, and the use of specific visual means to achieve their goal. Bernard et al. [Ber+17] devised a tool for supervising the labeling of human motion data by providing suggestions of viable candidates for labeling (see Figure 1.11).

Choo et al. [Cho+10] described a system that interactively helps the user classify data. Through the use of dimension reduction algorithms, clusters are visualized as scatter plots. Subsequently, the user is able to modify the initial classification thanks to similarity and distance metrics, showing how the initial clusters relate among each other. Through such metrics, the user is enabled to modify and steer the classification process. The results of this interaction cycle will then be incorporated and used for the calculation of the next clusters. In a similar way, Garg et al. [GRM10] used hidden Markov models to segment data. By providing the user with a semantic interpretation of the clusters, the user is then able to refine the initial segments. Migut et al. [MW10; MGW11] applied a similar methodology to a risk-assessment scenario. Patients with psychiatric diseases are classified by intertwining algorithms and user feedback. Drucker et al. [DFB11] designed a system supporting the creation of a data model by proposing the most prominent elements to be added to the different clusters. The recommendations are based on a machine learning model that adapts over time, making the suggestions dynamic.

Ankerst et al. [AEK00] proposed a supervised tool for building decision trees combining the computational power of the system and the knowledge of the user. This is one of the few examples where the model-building procedure (i.e., the black box) is opened and split into separate steps to allow for the finetuning of operations. The user is also able to steer the process and intervene at each elaboration step. In the same vein, Endert et al. [EFN12] allowed the user to finetune and steer model-building activities by representing data on the 2D plane. The system supports the user by searching for similar related data entities displayed and positioned together. The user can directly move data points around to alter the clusters and influence the discovery of similar data.

Parameter Refinement Müller et al. [Mül+08] present Morpheus, a tool supporting the visualization and interactive exploration of subspace clusterings. In particular, it

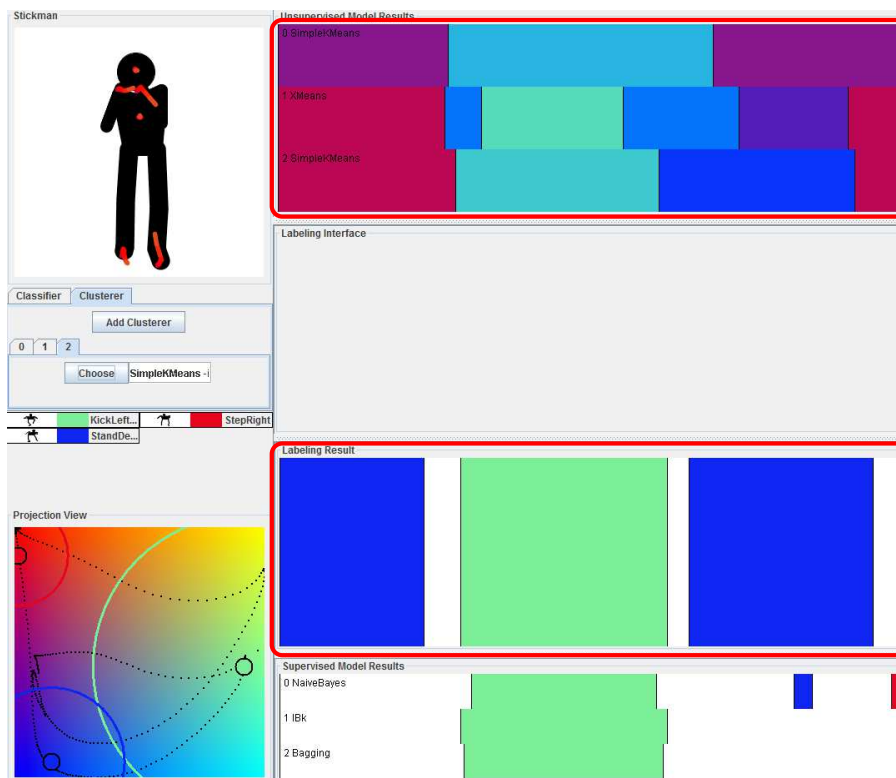


Figure 1.11: Bernard et al. [Ber+17] provide *orienting* guidance to *model building* activities. Through the use of unsupervised (top red box) and supervised (bottom red box) methodologies, the approach is able to provide the user with suitable suggestions of labels for different captures of human motion data. (Figure taken from [CGM19a], ©2019 Wiley. Used with Permission)

helps users interactively choose the best parameter setting. By presenting different visualizations representing the results of different parameter choices, users may choose the best ones that generate the desired subspace clustering. Dörk et al. [DLB13], using different metrics emphasize parameters that will lead to relevant visualizations. To foster a better parameter selection, Jeong et al. [Jeo+09] described a tool that shows the relationships between the data and the output of a principal component analysis algorithm. This is achieved by highlighting the effects that each data column has on the final result.

Exploration. Exploratory analysis combines the previous analysis steps (i.e., data preprocessing, data visualization, data models, and algorithms) to achieve a higher goal—usually a complex task. The primary aim of providing guidance during data exploration is to support the user’s interactions and the discovery of new findings. A finding is an interesting observation made by an analyst [Sac+14b], and it usually refers to the disclosure of a pattern or some interesting data subset.

Findings. The majority of the works providing exploration guidance achieve their goal by pointing the user to interesting data and data structures. Heer and Boyd [HB05] described Vizster, a visualization system supporting the exploration and navigation of large social networks with the goal of finding communities. The system identifies and highlights such communities and provides the user with the means to browse them. Johansson et al. [Joh+05] showed clusters using parallel coordinates and applied a number of visual techniques (i.e., highlighting, grouping, coloring, applying textures) that support an efficient analysis of the structure within these clusters.

To support exploration of weather phenomena, Steed et al. [Ste+09] designed a system that displays the correlation among environmental variables and the underlying data so that users can understand them better while also being able to use them more efficiently to predict weather events such as hurricanes. Adler et al. [Adl+10] supported visual navigation in surgical operations by augmenting the visualization environment with patient-specific anatomical data. The user/surgeon is enabled to set and change the most appropriate visual target as the exploration evolves. By using a flexible degree of interest function, Alsakran et al. [Als+11] showed interesting relationships among a set of streaming textual data. The user is allowed, at any time, to modify the interest measure and influence the layout of incoming textual nodes. Ip et al. [IV11] presented a system that helps the user identify salient patterns and interesting areas in very large images (e.g., landscapes). This is achieved by means of a saliency measure that serves to identify interesting areas for user exploration. Domino [Gra+14] supported the user's exploration of multiple datasets by providing hints to arrange, combine, and extract subsets of different datasets. In contrast, StratomeX [Lex+12] is focused on the exploration of relationships in cancer sub-type datasets. The system displays ribbons between data columns to highlight relationships among data features. Scorpion [WM13] is a tool that supports the exploration of data outliers by pointing users to the possible data tuples from which these outliers originated. Finally, Bernard et al. [Ber+12a] developed a system that is able to highlight interesting relationships among data subsets, thereby helping the user gain a better overview of a given dataset and its internal relationships. Gladish et al. [GST13] developed an approach that, by using a flexible degree-of-interest measure, is able to show interesting data regions to explore during the analysis of hierarchical data (see Figure 1.12).

Actions. A few works deal with guiding the exploration by directly supporting the interaction. This is usually translated into suggesting the next actions to take. One of the first attempts in this direction is the systematic yet flexible system by Perer and Shneiderman [PS08]. Unexplored states are shown to the user so he/she can systematically explore the entire dataset. In the context of personalized learning, Krishnamoorthy et al. [KB06] presented the user with a set of personalized suggestions about what documents to read next. Streit et al. [Str+12] designed a model to steer exploratory analysis. Based on data features and task descriptors, the system shows the users the data to explore and the next (alternative) steps to take in order to complete a given task. The approach developed by May et al. [May+12b] presents signposts pointing at interesting regions of

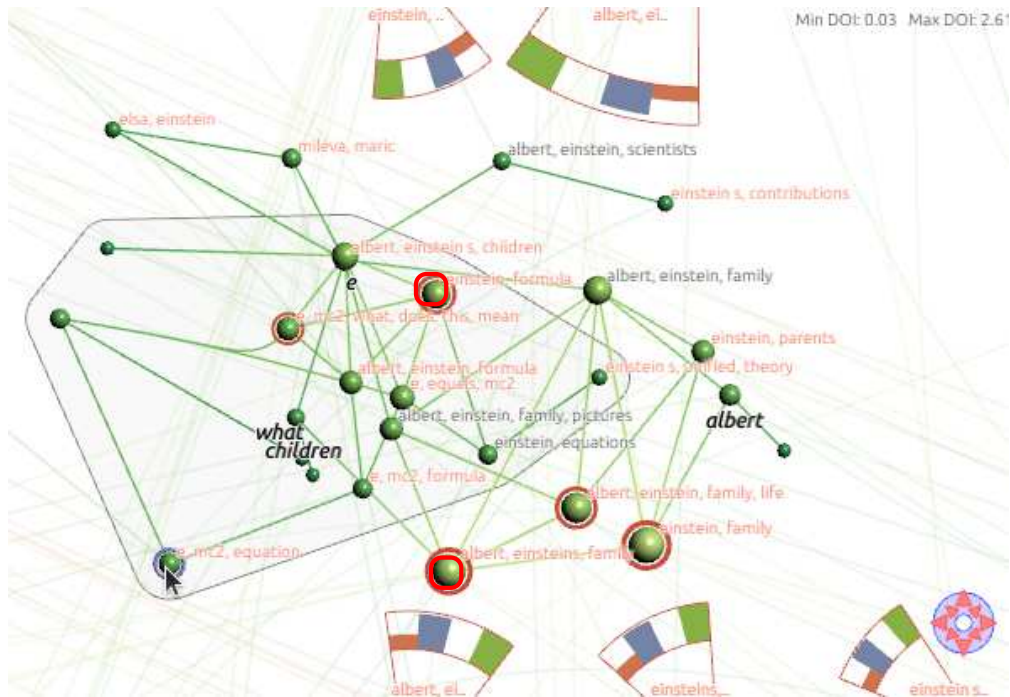


Figure 1.12: Gladisch et al. [GST13] provide *orienting* guidance to support the exploration of large graphs. In particular, the system guides the user towards interesting areas of the graph. (Figure taken from [CGM19a], Courtesy of Christian Tominsky, ©2019 Wiley. Used with Permission)

the graph, thereby informing the user about the possible next steps to take. Similarly, Crnovrsanin et al. [Crn+11], upon selection of a node in a graph, can recommend actions to continue the exploration.

Verification and Knowledge Generation. Whereas the previous works focused on findings, the following ones deal with arranging those findings into valuable insights and new knowledge.

Yang et al. [Yan+07] designed an approach for managing discoveries in visual analysis. The system supports the organization of facts and findings by suggesting clustering of a given discovery based on semantic similarity. Shrinivasan et al. [SGL09] presented an approach for helping the user in the activity of connecting the dots. Based on the current line of inquiry, the system suggests findings, notes, and concepts and how to arrange them together. Chen and Scott [CBY10] developed an approach for semi-automated annotation to support insights into externalization activities and the reconstruction of the process that leads to this insight. The work by Hossain et al. [Hos+11] guides the user through the process of arranging facts from a collection of documents with the aim of creating a story. The system provides the user with a set of paths (stories) connecting

an initial and a final document. The user can navigate and explore the suggestions as well as adjust the provided stories.

1.6 Steering the Guidance Process

Outline *In the following section, we complement our initial answer to S1 by discussing how a user can steer the analysis providing guidance to the analysis system. In doing this, we also discuss what constitutes guidance in the first place. Contentwise, this section covers parts published in [CGM19a] and [Cen+18a].*

The definition we introduced earlier stresses that guidance is applied in interactive contexts in which a user must take decisions to progress. However, this can be said of any VA approach. This leads to another question: Does this mean that any VA approach can actually be considered a guidance approach?

If we apply our definition to existing approaches, it is indeed difficult to discern which system is classified as guidance and which is not. After all, VA approaches are all designed to "support" users solving some kind of tasks. As our definition provides room for uncertainties, we worked to refine it and identify the ground rules for determining what constitutes guidance in the first place. In terms of research goals, we extended and complemented the initial question S1 with the following question:

"How can we clearly separate guidance from other approaches that do not provide guidance?"

In order to clearly separate guidance from other approaches, a simple question has to be asked: Is it the system or the user the one who makes decisions, takes actions, and is responsible for promoting the analysis? There are only three valid answers to this question. These answers define the boundaries of guidance approaches.

Although the majority of the existing VA, InfoVis, and HCI approaches are designed to facilitate specific tasks, we can argue that most of them do not readily qualify as guidance approaches. In some cases, they are only **user-initiative** (see Figure 1.13) methods, in which the user alone advances the analysis process and generates ideas about how to reach the analysis target [AS03; Shn10]. Consequently, we cannot consider such methods as guidance approaches. No matter how or how well the system may facilitate the analysis, in the end, it will always be up to the user to derive the necessary information to proceed toward the analysis target.

On the other extreme of the scale are **system-initiative** methods (see Figure 1.13). These methods encompass automated analysis processes that autonomously direct users to the completion of a task [PF91; FPS96], taking over actions and making choices automatically [Lie+95; Mae94; HH98]. System-initiative approaches do not comply with the definition of guidance either.

Instead, guidance is a **mixed-initiative** process [Hor99b] (see Figure 1.13). Guidance is like a dialogue between the human and the machine in which users communicate,

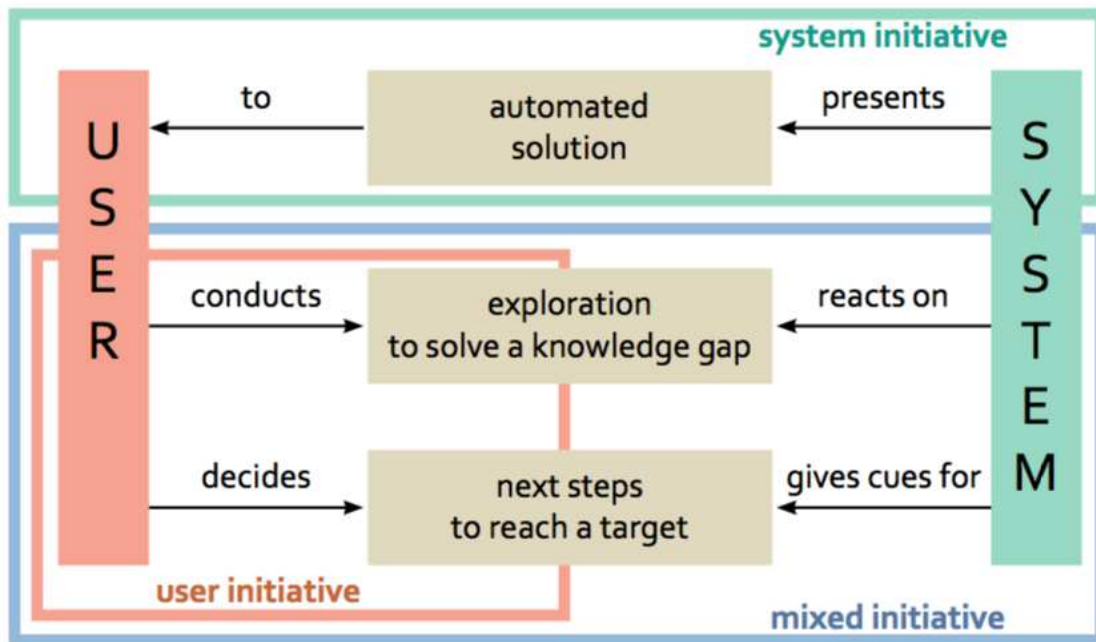


Figure 1.13: Guidance is a mixed-initiative process. On the one hand, the user explicitly or implicitly expresses his/her analysis target and a possible knowledge gap that hinders progression, by interacting with the system. On the other hand, the system reacts to the user's actions and gives cues that help to decide which steps to take to reach the target. (Figure taken from [Cen+18a], ©2018 ACM. Used with Permission)

implicitly or explicitly, their own needs as input and the system provides possible answers to alleviate problematic situations. As we have already seen, answers can be given at different levels of sophistication (or guidance degrees). A detailed discussion of what constitutes guidance in the first place can be found in [Cen+18a; CGM19a]. An excerpt of these publications is reported in the following.

As guidance is a mixed-initiative process, in the remainder of this section, we review the ways guidance is provided by the user to the system, closing the guidance loop and the mixed-initiative analysis process we started describing in the previous section. User guidance may serve to close the guidance loop after system guidance has been provided or to initiate guidance in the first place. Theoretically, as a consequence of user guidance, a system should provide further additional visual cues to the user in order to acknowledge a change in the analysis course (due to the received input) and initiate (once again) the guidance process. Guidance provided by the user can have different purposes. In the following section, we start by describing what we refer to as the *direction* of user guidance and how such guidance can be derived from the user's actions.

1.6.1 Guidance Direction: Feedback and Feedforward

The aim of user guidance is to foster changes and steer the analysis. We discern approaches that allow the user to push the analysis forward by evaluating the results and the guidance the system has produced *in the past* as well as those that allow the user to input directly what results or what kind of guidance suggestions he/she expects to see *in the future*. Following the terminology used in the cognitive science, we use the term feedback for the former and feedforward actions for the latter. We can further specify the quality of such actions, which can be positive or negative. Thus, we have four combinations: **positive or negative feedback** and **positive or negative feedforward**.

Positive or Negative Feedback. Most approaches enable the user to provide feedback (either positive or negative) in response to system actions. Hence, in this subsection we present a few examples to show how these approaches implement the feedback loop. For an overview, refer to Table 1.1.

To guide view selection in volume visualization, Zheng et al. [ZAM11] proposed an approach that suggests optimal viewpoints. The user can provide positive feedback on these suggestions by selecting the most promising ones. In response, the system updates the suggestions, pointing to new and promising but thus far non-viewed directions. Fujishiro et al. [Fuj+97] designed a system supporting the design of appropriate visualizations. As the user interacts with the tool, the system proposes and suggests additions to the actual design. The user, by selecting the most appropriate additions, guides and provides feedback to the guidance mechanism, influencing future suggestions. A similar approach was proposed by Müller et al. [Mül+08]. The Morpheus system supports the interactive exploration of subspace clustering by presenting suitable results. Based on the discovered knowledge, the user can give feedback to the system for improving the suggestions. In this case, the feedback loop enables the user to set parameters and, thus, discover meaningful subspace clusters. Andrienko et al. [And+09] provided guidance to support the visual clustering of trajectories. In this context, users can modify the cluster result computed by the system by excluding one or several subclusters from the cluster itself. Other aspects of this approach relate to the feedforward concept; thus, we will discuss it further in the next paragraph. Stein et al. [Ste+15; Sac+14a] proposed a visual analytics approach supporting the analysis of soccer matches. By extracting the most interesting features from the data, the system is able to propose to the user interesting events that characterize the match. The user can steer the exploration process by confirming or rejecting previously unlabeled events and using them as additional training data for the classifier. Finally, Click2Annotate [CBY10] supports semi-automated insights annotation. If the user is not satisfied with the annotations generated by the system, he/she can modify and affect the outlook of future annotations by dragging and dropping statistical measures into the annotation itself.

Positive or Negative Feedforward. Another group of approaches enable the user to directly input what he/she wants to obtain from the analysis. Therefore, the following

techniques differentiate themselves from the previous in that the user is focused on the future of the analysis instead of the past. This is usually achieved by sketching examples of what the user expects to obtain. Sketch-based information retrieval is a very popular field. However, in the context of this review, we only consider the literature related to guidance techniques.

Chegini et al. [Che+18] developed a system supporting the visual exploration of patterns in large scatter-plot matrices. Usually, the analysis space is huge, so to reduce users' effort, the system recommends suitable patterns for close-up investigations. On the other hand, the user can actively input what he/she is currently looking for in the data; the user can directly draw or select patterns representing the searched output. Andrienko et al.'s [And+09] work supports the clustering of trajectories, allowing users to provide feedforward actions. Users can, in fact, split, combine, and create their own clusters, thereby suggesting directly to the system how the clustering algorithm should categorize the data. The work by Janetzko et al. [Sac+14a] comprises elements that can be related to feedforward actions. In fact, it allows the user to steer the exploration of semantically meaningful soccer events by integrating the possibility for the user to sketch and describe dangerous situations that should be taken into consideration. Migut et al. [MGW11] guide the classification of psychiatric patients. The user can steer the model-building process by indicating to the system prototypes of patients in which they are interested. The iCluster system [DFB11] helps the user cluster large document collections by providing recommendations. The system learns and subsequently adjusts the suggestions as the user interacts with the tool showing the system how he/she would like to organize the documents. The approach conceived by O'Donovan et al. [OAH15] aims to guide the design of visual layouts. Users can specify their own intents in the form of constraints and by sketching partial layouts, thereby steering the guidance process.

1.6.2 Guidance Inference

Whereas the previous paragraphs analyzed the direction of user guidance, in this section we describe how users can convey such guidance to the system. The most common manner for inferring user guidance is by analyzing the user's direct actions with the interface widgets, such as by using drop-down lists, buttons, check-boxes, etc., as sources of guidance. Among others, we might consider, for instance, the user providing weights for mining algorithms as input or users selecting visual parameters for a visual mapping as well as data annotations for insight generation.

Different taxonomies are available to discern how a user can provide input [AES05; Yi+07; Shn96; And+11]. Existing taxonomies on this topic are mainly focused on assigning a meaning to certain actions, with the aim of understanding users' intents. Instead, in the context of this survey, we analyze how this interaction affects the guidance process.

The other way to derive guidance is by interpreting users' *indirect actions*. We chose this focus in line with the notion of "*user is the loop*" developed by Endert et al. [End+14]. They shifted the focus from approaches exploiting direct actions toward the creation of

new approaches in which the user does not simply take part in the analysis process, but is part of it. This new concept fosters the creation of more immersive tools that directly learn from user interactions instead of waiting for direct user input.

Direct Actions. Performing direct actions is the most common way to provide input and guidance to the system. We analyze tools and techniques that offer direct user input through interface widgets. We can also take into consideration the temporal aspects of such interaction, such as when exploiting a history of actions to derive the user's intent.

Direct Manipulation of Parameters. To support data transformation, May et al. [May+11] proposed a method that highlights interesting data features. On the user side, the direct selection of those interesting features causes the recalculation and updates those measures and metrics, closing the guidance loop. Bernard et al. [Ber+12b] supported the design of preprocessing pipelines for time-series data. Although the system points to the most promising parameters for each processing step, the user can steer the process by selecting appropriate weights directly.

Bouali et al. [BGV16] provided guidance for the generation of visualizations. The system proposes a set of suggestions from which the user can choose. The selection of the most appropriate visual mapping provides input for the creation of the next visualization generation. Similarly, in DesignScape [OAH15], previews of design suggestions are shown to the user, who can select the most promising one. Many other approaches exploit direct user feedback to generate visual mappings [Won+16; Koo+08; Got+10].

Ankerst et al. [AEK00] supported the generation of decision trees. The direct selection of the next split points steers the tree construction process. Choo et al. [Cho+10] proposed a system to guide the data classification process. Users can directly select an area in the data view containing uncategorized items and subsequently re-run clustering algorithms to optimize the process. Similarly, Garg et al. [GRM10] helped the user associate elements to proper clusters. The user is directly involved in the manipulation of the clustering parameters to split and join clusters. Other approaches follow a similar strategy to model building [MGW11; And+09; MW10; Jeo+09; DFB11; Pir+96].

Similar concepts exist for data exploration and annotation. Domino and StratomeX [Gra+14; Lex+12] supported this task by suggesting how different parts of data are related. Thus, the user can interact, connect, and arrange data chunks based on the suggestions. Streit et al. [Str+12] allowed exploration steering by presenting the next analysis steps. Based on the given task, the user can choose the most appropriate analysis direction. Ip et al. [IV11] guided the exploration of large images. By means of direct actions performed on the interface, the user can modify the selection of interesting views and image areas. Shrinivasan et al. [SGL09] helped the construction of data stories by structuring the data according to a given start and end document. The user can directly choose among the suggested structures and affect the composition of the story as well as the story-line. Many other approaches are present in the literature [Yan+07; HB05; Joh+05; Ste+09; CBY10; Hos+11; Fia12].

History. This category comprises those approaches that go beyond capturing single actions. A temporal component is also taken into consideration. Hence, complex action patterns and sequences of interactions are compared within each other to understand the user’s analytical intents. The system can exploit such findings to finetune the guidance, provide better suggestions, and steer the analysis process accordingly.

Gotz et al. [GW09] supported visualization creation by considering complex interaction patterns. The captured patterns are compared with a knowledge base to understand the visual task the user is performing. This, in turn, influences the suggestions of visualizations that best fit the inferred task. Horvitz et al. [Hor+98] supported data exploration by modeling users’ time-varying needs via Bayesian networks. The suggestions proposed by the system to pursue an analytical goal are influenced by the users’ actions. Temporal series of actions are interpreted and based on the inferred goal to propose a next step. Yang et al. [Yan+07] provided guidance to extract valuable information nuggets hidden in the data based on the user’s preferences. In this case, such interests are also directly inferred from the user’s action history and are the base for the retrieval of new information nuggets.

Indirect Actions. Indirect actions involve providing feedback by acting on the data, rather than explicitly stating intentions through the interface widgets.

Spatial Actions. Strictly connected with implicit feedback is the concept of spatialization. For instance, a user provides guidance to the system by acting directly on the data. The system thus learns and infers weights, parameters, and preferences from the user’s actions. In particular, spatialization is derived from how the user modifies the spatial properties of the data (e.g., moving and grouping data).

Endert et al. [EFN12] designed ForceSpire to guide the visual exploration of text documents. To achieve that, this tool modifies the spatialization of data items on a canvas in such a way that the rendered layout reflects the user’s notion of similarity among documents. A decisive aspect of this approach is that distance metrics are implicitly suggested by the user through what the authors called semantic interactions. The movement and juxtaposition of data points (documents) are interpreted, and the weights determining the similarity metrics are implicitly changed. To the best of our knowledge, ForceFire is the only approach using semantic interaction and the implicit feedback paradigm through spatialization. However, although not directly connected with Endert’s ForceSpire [EFN12], other approaches (marked with an empty black circle in Table 1.1) utilize similar expedients to allow users to signal their intents and implicitly provide feedback to the system. Garg et al. [GRM10] supported model-building activities allowing the user to resolve inconsistently categorized data points by moving them to different clusters. Finally, Jeong et al. [Jeo+09] allowed the user to move data items to different clusters to influence the dimension reduction algorithm.

1. GUIDANCE IN VISUAL ANALYTICS

Papers	System Guidance										User Guidance			
	Guidance Objective							Guidance Deg.			Inference		Direction	
	Transf	Map	Par	ModV	ModB	Expl	Know	Or	Dir	Pre	Dir	Ind	Bck	Fwd
Total: 53	5	7	5	2	12	28	4	35	18	3	51	5	48	5
[May+11]	•							◦	•		•		•	
[Ber+12b]	•								•		•		•	
[Kan+11b]	•								•		•		•	
[Kan+12]	•								•		•		•	
[HHK15]	•								•		•		•	
[BGV16]		•						•			•		•	
[Won+16]		•							•		•		•	
[Fuj+97]		•							•		•		•	
[Got+10]		•							•		•		•	
[GW09]		•							•		•		•	
[Koo+08]		•						•			•		•	
[OAH15]		•							•		•			•
[Mül+08]			•			•		•			•		•	
[GRM10]			•		•			•			•	◦	•	
[MW10]			•		•			•			•		•	
[And+09]			•		•			•			•			•
[DLB13]			•			•			•		•		•	
[AEK00]				◦	•			•			•		•	
[ZAM11]				•		•			•		•		•	
[Cho+10]					•			•			•		•	
[MGW11]					•	•		•			•			•
[Jeo+09]					•			•			•	◦	•	
[EFN12]					•	•			•		•	•	•	
[DFB11]					•				•		•		•	
[KPN16]					•			•			•		•	
[Ber+17]					•			•			•		•	
[EFN12]					◦	•		•			•	•	•	
[Adl+10]						•			•		•			•
[HB05]						•		•			•	◦	•	
[Joh+05]						•		•			•		•	
[Yan+07]						◦	•		•	◦	•		•	
[IV11]						•		•		•	•		•	
[Ste+09]						•		•			•		•	
[Als+11]						•		•			•		•	
[Lex+12]						•		•			•		•	
[Gra+14]						•		•			•		•	
[Hor+98]						•				•	•		•	
[Str+12]						•		•	•		•		•	
[PS08]						•		•			•		•	
[Ber+14]						•		•			•		•	
[May+12b]						•		•			•		•	
[Crn+11]						•		•			•		•	
[GST13]						•		•			•		•	
[Tol+12]						•		•			•		•	
[ST15]						•		•			•		•	
[JN15]						•		•			•		•	
[Lub+12]						•		•			•		•	
[KB06]						•		•			•		•	
[WM13]						•		•			•		•	
[CBY10]							•	•			•		•	
[SGL09]							•		•		•		•	
[Hos+11]							•	•			•			•

Table 1.1: Table summarizing the classification of guidance papers. Columns represent the different aspects we took into consideration, while papers are listed as rows. The rows are sorted according to the guidance objective they support. We considered approaches providing guidance for different **objectives**: approaches supporting data **Transformation**, visual **Mapping**, **Parameter** setting, **Model Visualization**, **Model Building**, **Exploration** and **Knowledge** generation. We considered three **Guidance degrees**: **O**rienting, **D**irecting, and **P**rescribing guidance. We also describe the user side of guidance, in particular the guidance **Inference**: **D**irect and **I**ndirect actions. Finally, we provide details of the guidance **Direction**: whether it is **F**eedback or **F**eedforward. Empty black circles (◦) show approaches that offer this functionality but to a minor extent. (Table taken from [CGM19a], ©2019 Wiley. Used with Permission)

1.7 The Effects of Guidance

Outline *We now shift our focus to S2 – "What are the benefits of using guidance during visual analytics?". To answer this question, we performed two user studies. In the first, we analyzed how users with different knowledge and expertise reacted to multiple types of guidance. We not only analyzed performance, but also focused on their personal experience with guidance and investigated if too much or too little guidance can be counterproductive. In the second study, we examined how guidance influences the way users solve tasks. The section is based on [CGM19b] and [Cen+18b].*

Having defined guidance based on either its "user" or "system" aspects, we proceeded to investigate if guidance can really benefit the analysis or if it is counterproductive in some instances.

With this scope in mind, we set up two user studies. With the first study, which is described in its entirety in 3, we wanted to understand how users' previous knowledge and the type of tasks influence the effectiveness of different types of guidance. In doing so, we aimed to evaluate not only the effects of guidance on the performance of the users involved, but also their influence on users' feelings and state of mind, investigating, for instance, if too much guidance can frustrate experienced users or, conversely, can help reduce frustration.

In the second study, we investigated the effects of a specific guidance solution on the way users performed different exploration tasks [Cen+18b]. We observed how their resolution strategies varied when guidance was provided compared to the scenario where no support was given.

1.7.1 User Study 1: Effects on Performance and Mental State

The first user study is based on a simple assumption: Three dimensions play an important role in the design of guidance in VA — namely, (1) the task type, (2) the knowledge of the user, and (3) the degree of guidance provided. Following a user-centered approach, we set up an analysis environment in which we alternatively varied one of these factors and analyzed if and how users were able to complete the given tasks.

The results we obtained shed light on the appropriateness of certain degrees of guidance in relation to different tasks as well as the overall influence of guidance on the analysis outcome, in terms of performance, and changes in the user's mental state. Previous studies (see Chapter 3) have shown that many negative emotions, among which frustration and sadness are the most typical, grow in users experiencing issues during the analysis. We examined whether guidance can put a remedy to such problematic situations, contributing to instill a positive analysis experience counteracting the rise of such "negative" feelings.

However, first we make a small digression to explain the three factors we considered in this user study and then provide details of the setup and study outcome.

Knowledge and Task Types. The first factor we consider is the knowledge required to complete a task. Two types of knowledge are usually required to complete visual data analysis: operational and domain knowledge [Che+09].

Our aim is to test whether the type of knowledge involved influences the effectiveness of the provided degree of guidance — that is, if a certain type of guidance can compensate for either a lack of domain or operational knowledge. According to the distinction between operational and domain knowledge, it is possible to delineate two types of tasks:

- *Exploratory tasks* are related to operational knowledge and therefore to the ability to interact with the tool. Exploratory tasks rely on basic interaction abilities, like the ability to choose among different interaction means (e.g., filter, selection) and the user’s ability of using them effectively. One exploratory task is portrayed in Figure 1.14.
- *Domain tasks* require domain-specific knowledge to be successfully completed. These tasks are related to the ability to reason and connect a given domain concept to the task and data under analysis.

User Knowledge. A second dimension we consider is the degree of the user’s competence. Usually, when there is a knowledge gap, the user might have a hard time completing tasks and the analysis may stall. To counteract this situation, more or less guidance should be administered.

Our hypothesis is that different degrees of guidance would have different effects on users with different levels of knowledge. For this reason, we distinguish between:

- *Knowledgeable users*, possessing the knowledge required to complete the task (i.e., operational or domain knowledge) and
- *Novice users*, who may not possess the knowledge required to complete the task, with an exception made for previous expertise.

Guidance Degrees. Finally, we distinguish among three different degrees of guidance. Our assumption is that users with different knowledge may require different degrees of assistance and that receiving more or less support while solving the task would reverberate on the overall performance and mental state. According to the guidance degrees described earlier, we discuss how users reacted to:

- *No guidance:* When no guidance is provided, users have to solve the tasks on their own. This is translated into the provision of simple visualizations, without any further support.
- *Directing guidance:* This kind of guidance aims at providing different analysis options. Therefore, on top of the basic visualizations, we indicate possible analysis paths. For instance, interesting data subsets are recommended to the user, but the system may also recommend actions to take.

- *Prescribing guidance*: This is the highest degree of guidance. It aims at providing step-by-step instructions to reach the result. Among the different analysis paths and recommendations (see directing guidance), the system picks one and provides it to the user, who must follow the indications (the different steps) to reach the final result.

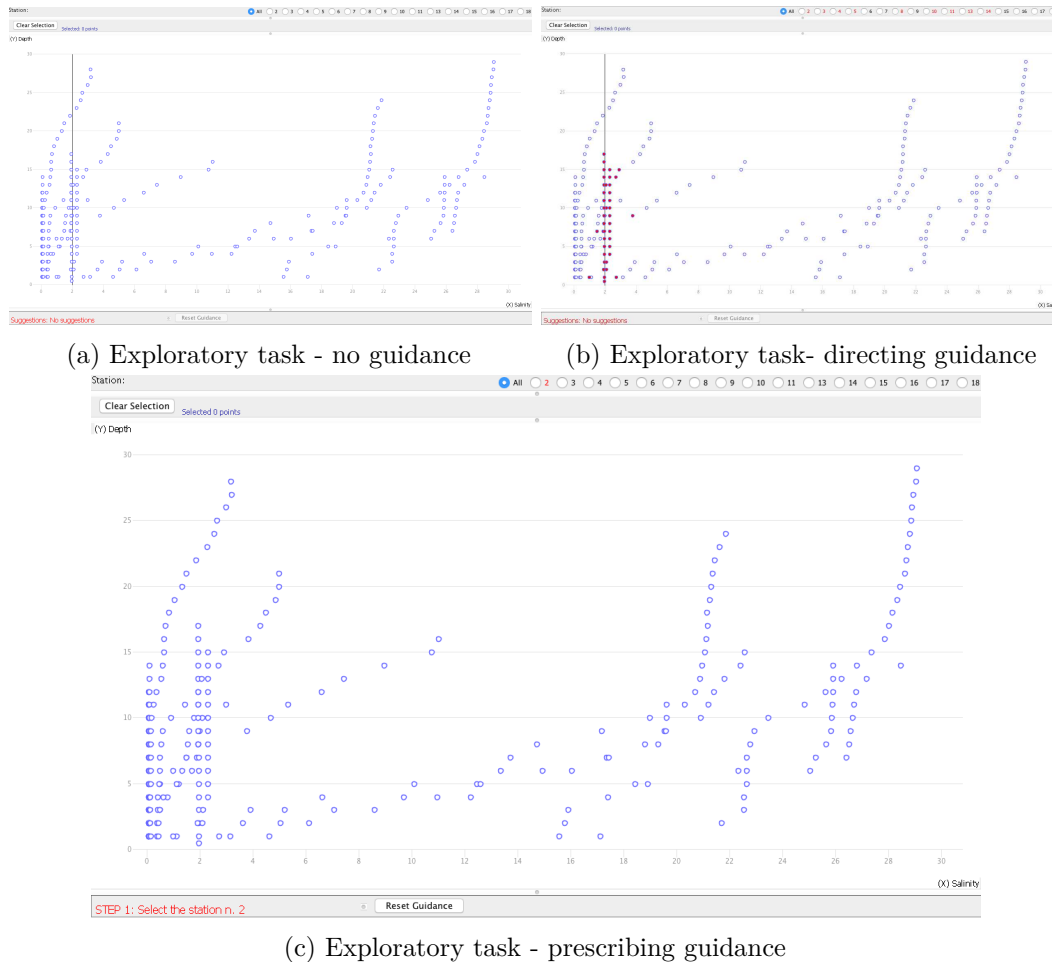


Figure 1.14: An example exploratory task users had to solve during the user study we set up. The task can be supported with different guidance. (a) no guidance: a scatter plot shows values of salinity (x-axis) in relation to the change of water depth (y-axis). User have to select and filter data points by themselves. (b) directing guidance: The system highlights interesting data points and filtering options; (c) prescribing guidance: the system provides step-by-step instructions to solve the task (bottom left of the image). (Figure taken from [CGM19b], ©2019 Elsevier. Used with Permission)

Study Setup and Aims

We compared the results obtained by the users in situations in which we alternatively varied one of the aforementioned aspects. Specifically, we asked some users to solve different tasks in a simulated analysis environment. The comparison was done by analyzing how the variation of the aforementioned aspects contributed to the change in (H1) users' performance and (H2) their mental state (i.e., their emotions).

As for (H1), we analyzed the effects of guidance on novice users and evaluated if the positive effect of guidance is mitigated in knowledgeable users. Our assumption is that knowledgeable users may still benefit from guidance in terms of reduced task completion time.

As for (H2), we evaluated how the mental state of users was altered when receiving different degrees of guidance. Specifically, if guidance increased the user's confidence or if in some situations guidance can frustrate the user.

To verify our assumptions, we designed a user study comprising six specific tasks (three exploratory and three domain tasks) that we asked 65 participants with different expertise and supported with different guidance degrees to solve. As an example, in Figure 1.14 we can see how one of the exploratory tasks can be supported with different degrees of guidance, by highlighting interesting data cases (b), guiding the interaction with tool (c), or without guidance at all (a). By varying the aforementioned factors, we analysed how users with different users with different expertise were able to solve the tasks we provided them.

The evaluation environment and the evaluation procedure were designed to avoid learning effects on study participants so that gained expertise in one kind of task that did not affect the performance on the others. The answers were collected in the form of questionnaires that we administered to the participants at the end of each task. Subjective feelings were measured on a 5-point scale while performance measures, like timings and errors, were automatically recorded by the analysis system. We subsequently analyzed the results with statistical tests to verify differences among the results obtained by the various study groups, according to performance and mental state variations due to guidance. Additional details of the study can be found in Chapter 3; here we report only an excerpt.

Outcome

The statistical test we ran indicated that guidance has an overall positive effect on users' performance. What's more interesting, though, is that the positive effects also relate to the user's mental state.

As summarized in Table 1.2, guidance was particularly successful for novice users solving exploratory tasks. In other words, it is easy to compensate for a lack of operational knowledge with simple guidance (e.g., suggesting what button to push). In addition, the tests highlight that, for domain tasks, at least a minimum of knowledge should be possessed by the users. We think that this is due to the fact that a certain degree of

Key Findings

1. While it is no surprise that a high degree of guidance had positive effects on the performance of novice users, it is remarkable that guidance, especially the prescribing degree, had significant positive effects on performance and mental state also of knowledgeable users for almost all combinations of task types (H1).

2. Guidance was particularly effective to account for the lack of operational knowledge. For domain tasks, the users should possess at least a minimum of knowledge to interpret correctly the suggestions. This indicates that missing operational knowledge is easier to compensate by guidance than missing domain knowledge (H1).

3. Knowledgeable users were not frustrated by high degrees of guidance while there was a positive effect on confidence and the subjective assessment of the difficulty of the task (H2).

4. Participants' subjective assessment of appropriateness of guidance degree was reflected in better performance, and more positive mental state, which reflects the importance of providing an appropriate degree of guidance for the given user (H2).

5. Knowledge plays an important role for positive performance and mental state especially when solving domain tasks. However, prescribing guidance may compensate for the lack of knowledge in many aspects (additional finding).

6. Knowledge may also compensate for a lack of guidance. Knowledgeable users with no guidance obtained similar performances to novice users provided with directing guidance, for both exploratory and domain tasks (additional finding).

7. Domain tasks evoked more frustration than exploratory tasks in novice users, since trial and error can compensate for a lack of operational knowledge while not for a lack of domain knowledge (additional finding).

Table 1.2: Key findings of the first user study. In this table, *additional finding* refers to results that were not taken directly from the hypotheses, but inferred from them. Please refer to Section 3.7.4 for the details.

knowledge is required not only to understand the question but also to interpret the suggestions received correctly.

The study revealed that guidance may have a bad impact on the analysis if the guidance degree does not match the knowledge gap and users' expectations. From the results we obtained, we noticed an increased number of errors in novice users receiving directing guidance. Novice users tended to trust excessively the guidance suggestions, even if sometimes they could be misleading. In addition, some knowledgeable users reported increased frustration when supported with prescribing guidance.

The tests showed that directing guidance was beneficial for knowledgeable users who were able to interpret and judge correctly the provided suggestions. When assisted with this kind of guidance, participants typically obtained performances similar to the one obtained by other users supported with prescribing guidance. On the other hand, directing guidance produced no improvements (same results as no guidance) when provided to

novice users. Thus, for novice users the prescribing degree of guidance seems better suited for solving domain-related tasks.

As for the mental state, the results showed that knowledgeable users were usually not frustrated by high degrees of guidance. We could also see how the provision of guidance had a positive effect on the reported confidence and the subjective assessment of the difficulty of the task (i.e., with guidance the users thought the tasks were easier to solve). Moreover, domain tasks evoked more frustration than exploratory tasks in novice users. Our interpretation is that trial and error can compensate for a lack of operational knowledge but cannot completely compensate for a lack of domain knowledge.

For a complete discussion of the results please refer to Chapter 3.

1.7.2 User Study 2: Changes of Strategy

In a second study, we changed perspective and focused on testing whether guidance can affect how users solve tasks. We considered users' strategies with and without guidance. This section refers to content published in [Cen+18b].

With this aim in mind, we designed and implemented a guidance-enriched tool (see Figure 1.15) to support the exploration of cyclical patterns in univariate time-series data [Cen+18b]. In the following section, we describe the design process of the analytical approach and the guidance. Afterwards, we describe how we evaluated the produced guidance with six visualization experts.

Scenario and Guidance Design

The exploration of cycles in a time series is a challenging task. The literature describes two main strategies to find cycles in data. First, it is possible to employ algorithms to detect them. However, this strategy has the disadvantage that it is often difficult to assess what patterns make sense. In addition, multiple algorithms can deliver different results, so users experience issues when they need to select or interpret the output of the automatic algorithms. Second, it is possible to employ a visual interactive solution. However, what typically happens is that the user is not sure to have analyzed all the cycles (i.e., when the analysis is complete). This leads to the user performing time-consuming iterations of trial-and-error analysis to reach an adequate confidence about the obtained results.

Our idea, in this study, was to combine the two alternatives following the VA paradigm, but giving it a boost and enhancing the whole process with guidance. In summary, we mixed a standard spiral plot visualization, which supports the visual detection of cycles, with the output of two algorithms. Thanks to their output and to a careful encoding we provided the user with guidance suggesting what cycles to investigate.

To design the guidance aspects, we chose the idea of scented widgets [WHA07]. As in Willet et al.'s approach, we embedded visual suggestions about what cycles to explore directly into interface widgets. For a smooth interaction, we also encoded additional

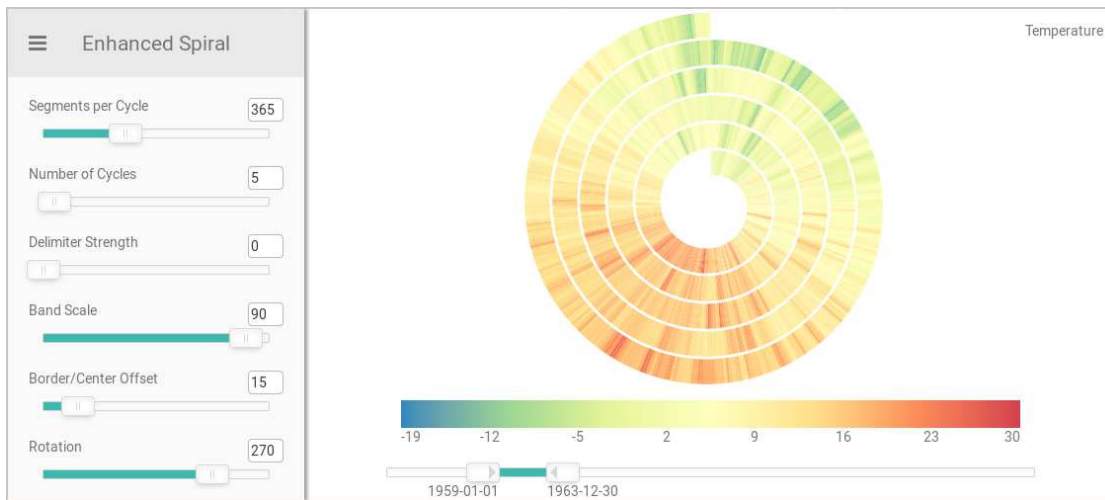


Figure 1.15: Spiral visualization (center) and user interface with sliders (left). The spiral visualizes a subset of 1,825 temperature values from a dataset with more than 25,000 days worth of data. 365 segments per spiral cycle are visualized to emphasize the yearly temperature fluctuation. (Figure taken from [Cen+18b], ©2018 IEEE. Used with Permission)

visual hints directly in the main spiral plot. Thanks to this solution, users were free to explore the suggestions provided by the system, reducing analysis timings and avoiding long trial-and-error sessions.

In line with the previously provided definition of guidance, we characterize the guidance solution we designed as orienting. The guidance is generated for addressing an unknown target. The knowledge gap we aim to tackle can also be defined as a data problem. In fact, the unknown target refers to the number of segments that should be visualized in each spiral cycle so that patterns appear. As we can see in 1.17 when the user selects the correct cycle length value, the colored segments of the spiral plot align to signal a cycle—in this specific case, an annual cycle. The guidance is generated using the output of two algorithms. In other words, algorithms constitute the input to the guidance process. The output is instead suggestions of cycle length values, which we chose to encode (the means) in the interface widgets that controls the appearance of the spiral.

Algorithms. The suggestions to guide the user were produced thanks to the combination of two algorithms: the Discrete Fourier Transform (DFT) [WW89] and the Chi-squared Periodogram (CSP) [SB78]. Both are typically used to find cyclical patterns in time-series data. The DFT is a classic algorithm that approximates the time series with a linear combination of basic functions. The CSP on the other hand, works by constructing a Chi-square periodogram structure, which is an efficient method that is usually exploited for spotting periodical patterns and circadian rhythms.

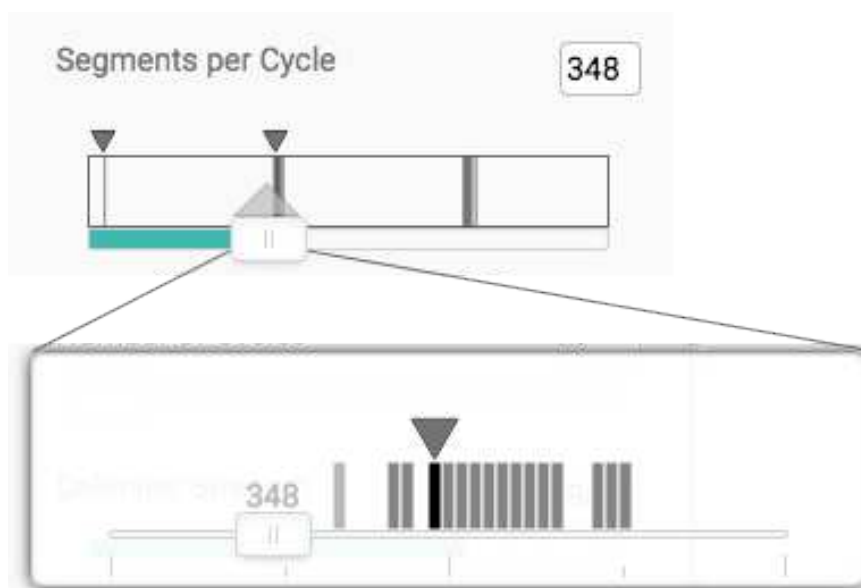


Figure 1.16: Guidance suggestions represented in a slider. The statistical output of CSP algorithm is used to delineate areas of interest while the different shades of gray/black encode probabilities. The output of the DFT algorithm is encoded with small triangles pointing to precise points within these areas. A further triangular shape handle, is used to show precisely the currently selected value. In the bottom part of the image, a zoomed view of the selection allows for a fine-grained value selection. (Figure taken from [Cen+18b], ©2018 IEEE. Used with Permission)

The selected algorithms have the same goal but work in a complementary way. The DFT output is a discrete set of cycles. For instance, if we imagine applying the algorithm to a dataset collecting the temperatures of a city, it will probably find a cycle with a length of 365, which corresponds to a yearly cycle (i.e., temperatures repeat each year with a similar pattern). Other values could be computed, such as 30, which represents a monthly cycle and can be related to the change of seasons.

Meanwhile, the CSP works by producing a set of candidates and associating them with a (continuous) probability value, ranging between 0 and 1. If we think of the variation of temperatures in a city in multiple years, the CSP algorithm will certainly assign high probabilities (close to 1) to cycle lengths in the vicinity of 365 (e.g., [360–370]) or to 30 in a way that resembles the results of the DFT. As years and seasons do not have a precise length, the CSP output will most certainly be complementary to the results of the DFT, as the results produced by the CSP algorithm show the small oscillations of the cycle length. The CSP was also chosen because it can work with recurring patterns. In our example scenario, cycle lengths around 730, corresponding to a biannual cycle, would have also been given a high probability.

Enhancing Sliders. We encoded the guidance suggestions directly on top of the sliders used to modify the appearance of the spiral plot. The visual encoding we chose is presented in Figure 1.16. We encoded the results of DFT and CSP with two different visual cues. The output of the DFT algorithm is encoded with downward pointing triangles. If the DFT algorithm computes, for instance, yearly and monthly cycles, triangles at the top of the slider would appear in correspondence to values close to 365 and 30 days.

The additional output provided by the CSP was instead encoded with bars of different brightness (see Figure 1.16). As the probabilities span a continuous spectrum ranging between 0 and 1, to keep the visualization as simple as possible, we grouped them into three main sets: high, medium, and low probability. We then visualized them with three different shades of gray, from dark to light, with darker shades representing higher probabilities.




Figure 1.17: Our guided solution for the exploration of cyclical patterns in time-series. A classical spiral plot is enhanced with data-driven guidance mechanisms to support the identification of patterns. We statistically determine cycle lengths that reveal strong patterns and visually indicate these interesting cycle lengths while the user interacts with the slider. (Figure taken from [Cen+18b], ©2018 IEEE. Used with Permission)

Enhancing the Spiral. Although the guidance-enriched slider provides enough support for the analysis, it is still necessary for the user to look back and forth between the spiral and the slider to visually verify the selection of specific cycle length values. As this continuous movement can be a source of distraction for the user, we complemented the described guidance with a further visual cue so that the user could spot the appearance of cycles just by looking at the spiral. In the specific case, we added a glow surrounding the spiral to indicate how far the current configuration is off from the closest suggested cycle length.

For instance, if the current cycle length value is far from the one suggested by the algorithms, the glow will have a larger radius. When approaching a suggested cycle length, the radius will decrease to the point where the glow is very thin and sharp to indicate that a suitable configuration has been reached. Although minimal, the glow is a source of additional guidance (see Figure 1.18).

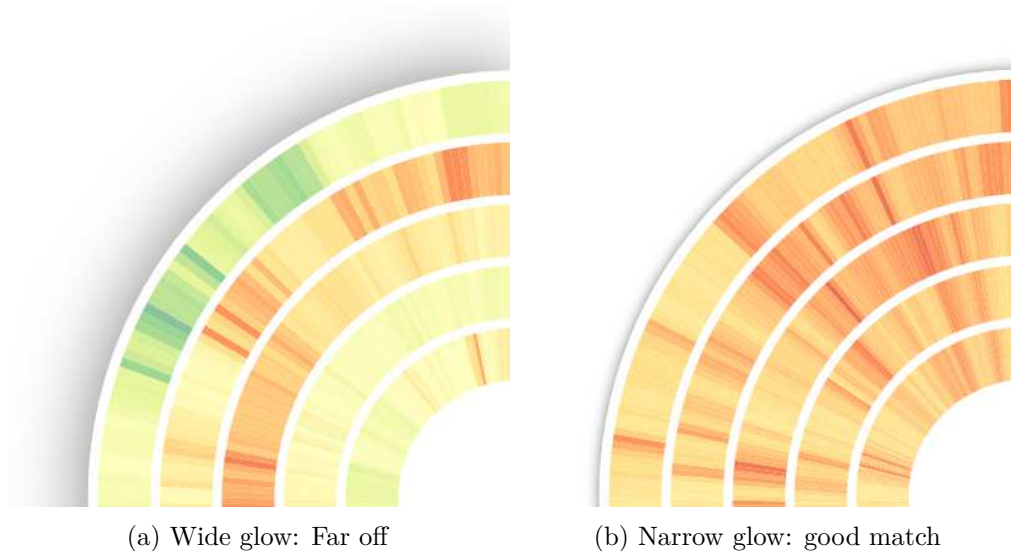


Figure 1.18: The distance to a suggested cycle length candidate is encoded as a glowing ring surrounding the spiral. a) A wide glow indicates that the current configuration is far off from a suggested cycle length. b) A narrow glow suggests that number of segments per cycle matches the cycle length in the data. (Figure taken from [Cen+18b], ©2018 IEEE. Used with Permission)

Evaluation

Keeping in mind the goal of evaluating the effects of guidance and answer for S2, we conducted a semi-structured qualitative evaluation using the described guidance-enriched prototype.

Similar to User Study 1, we asked a number of participants to perform exploratory tasks, at first without any assistance and then supported by the guidance suggestions we introduced earlier. We collected qualitative feedback, in written form, after each task to understand how the users performed under the two different conditions.

We had in mind two hypotheses when designing the study:

- H1** The guidance mechanism brings benefits to the analysis in terms of an increased trust and confidence of the user towards the obtained results.

H2 The implementation of guidance relieves the users from mental workload, allowing them to focus their attention on reasoning and on confirming work hypotheses.

The value of our evaluation derives not only from the evaluation of changes in users' mental states (as in User Study 1), but also from the evaluation of how much the integration of guidance affected users' resolution strategies. We wanted to evaluate not only if the designed guidance had positive effects on users' performance but also if the guidance was in line with the user's way of reasoning and if it affected the analysis strategies, if the guidance distracted the user, or if it facilitated insight discovery. Finally, we wanted to evaluate users' trust and confidence toward the guidance suggestions.

Study Participants. Unlike in User Study 1, in this study we performed a qualitative study focusing on a smaller number of subjects. We recruited six visualization experts directly from our research group. At the time of the study, the study participants were all pursuing their doctorates in the field of VA and were familiar with the visual and interactive means we used in our setup. For obvious reasons, they were not involved in the design phase, so they were totally unaware of what they were going to experience.

Datasets. As we prepared two tasks, we also chose to use two datasets for the study: one containing real data and one containing artificial data created by us. The first dataset contained information about the weather in the city of Rostock, Germany. It included about 25,000 entries and different measurements, like humidity, air pressure, and quantity of precipitation. For the purposes of our study, we chose to focus solely on the temperature measurements. The interface and the dataset are shown in Figure 1.15 and a yearly cycle is displayed.

We created the second dataset artificially by modeling a sine wave with a cycle length of 13 days. We created this artificial dataset to simulate the case of cycles not following the common calendar subdivision and provide an additional challenge to the participants. In this artificial dataset, we also introduced a variable amount of white noise to avoid the cycles being too easily recognizable.

A final reason why we chose these datasets, especially the temperature measurements dataset, was that we had already used them in the past and knew in advance the cycles they contained. We used such knowledge as a foundation of truth to compare with the participants' findings.

Procedure. We randomly subdivided the participants into two groups and asked them to solve an exploratory task without and with guidance. After a short introduction, the participants were free to interact with the analysis system and the sliders. After a few more minutes we started the study and asked the participants to solve the following task:

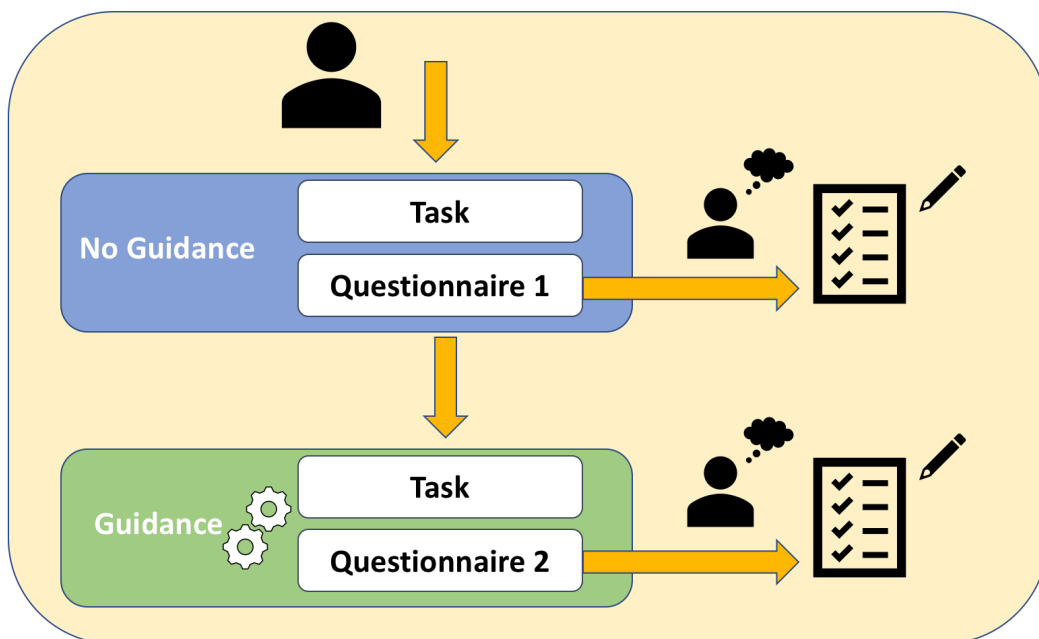


Figure 1.19: The procedure of the user study # 2. We asked the participants to perform the same task under different conditions (guidance/no guidance). After each execution we asked questions regarding their trust and confidence, to understand if guidance improved the analysis. (Figure taken from [Cen+18b], ©2018 IEEE. Used with Permission)

Find all the cyclical patterns present in the dataset and report, for each of them, the cycle length. You are also encouraged to think aloud about the task while you solve it.

They had to find all the cyclical patterns in the dataset and possibly reason about their relevance (e.g., the relevance of a cycle that is a multiple of another cycle may be considered less relevant). We asked them to think aloud and explain their thoughts while solving the task. An external person sat next to the participants to take notes of the analysis.

The task was initially carried out without guidance support. This means the study participants performed the analysis on their own and did not receive any suggestions about cycles to explore. Then we asked them three questions to get a first idea about how they performed the analysis and their impressions. We asked them (1) if they followed a specific strategy, (2) if they felt confident about the results they found, and (3) if they believed they could detect all relevant cycles.

In the second phase of the study (see Figure 1.19), the participants performed the same task but on the other dataset. This time the guidance suggestions were enabled. At the end of the analysis, we asked them six questions targeted at discovering if the guidance

affected the way they performed the task. Specifically, we asked them (1) if they followed a different analysis strategy (compared to the first phase) and if the guidance influenced their choice; (2) whether they felt more confident about the results they found; (3) if they felt that they may have missed some important results; (4) what trust they had in the analytical algorithms; (5) if they felt that the suggestions were leading them to unwanted analysis paths toward wrong results; and (6) if they found the suggestions to be useful in solving the task and followed their way of reasoning. Finally, we collected comments about the guidance solution and design choices.

Results

From the questionnaires and by observing the participants, we could derive how they performed under the two conditions and how guidance influenced the analysis. In particular, the questionnaires allowed us to evaluate how guidance changed the analysis strategies. Summarizing it in a few words, the study shows that our combination of visual and algorithmic means was effective and had a positive impact on the data analysis. Some details are provided next.

Detection of Cyclical Patterns. Although we did not precisely measure the correctness of task nor completion times, we noticed that, when supported with guidance, the users found more cycles and, more importantly, reasoned more about the results. We noticed they were more inclined to talk and formulate hypotheses (e.g., what phenomenon was reported in the data) and to reason aloud (H2). When guidance was enabled, the participants were able to rate the relevance of the cycles they found, order them, and reason about recurrences and multiple cycles. Without guidance they had more approximated answers and, in most cases, did not find all the cycles.

Analysis Strategies. Without guidance, the participants followed quite closely a trial-and-error workflow. When guidance was provided, they changed their strategy.

Without any support, the main way of solving the task was characterized by first getting an overview of the dataset followed by an exploration phase during which the users tested different values for the cycle length. After this exploration, users proceeded with a deeper inspection of a selected group of cycle length values. Some of them explored the values from the highest to the lowest; others performed it the other way around. Without guidance the analysis was instead characterized by a thorough exploration of all the possible cycle lengths, followed by a confirmation phase in which the most promising cycle lengths were inspected in detail.

The introduction of guidance had a substantial effect on the analysis strategy. Users spent most of their time evaluating the cycles suggested by the algorithms instead of exploring all the different cycle length values. We noticed that, thanks to guidance, the participants developed a deeper understanding of the data. They reasoned more about the possible meaning of the phenomena represented in the data and also more research

hypotheses than without guidance (H2). This change of strategy was reported by all the participants.

Confidence and Trust. As in User Study 1, we also analyzed how guidance affected participants' feelings. Although the participants we recruited can be considered experts in VA, the majority of them still reported an increased confidence in the results when guidance was enabled (H1). They also reported having the impression that they completed tasks faster and more easily, although we did not record timings precisely and could not confirm such statements. This claim concurs with the findings of User Study 1, in which the addition of guidance gave the impression that the tasks were easier to solve.

The participants stated they trusted the provided guidance suggestions (H1). However, we noticed that the participants did not consider the analysis as finished when all the suggestions were explored, but always spent some extra time to continue the analysis. We investigated this behavior further and discovered that it was mainly due to a problem of trust. When questioned about the behavior, the participants stated that their main problem with trusting the provided suggestions was their lack of knowledge about the algorithms that computed those recommendations. Some stated they felt like the guidance could have potentially missed some cycles and that not all the results were displayed. The participants also reported that their confidence in the guidance mechanism increased after a certain time. This feedback confirms the positive effects of guidance (H1). However, it also shows that shedding light on the functioning mechanisms behind the guidance process is important to raise trust and the overall effectiveness of guidance.

1.8 Designing Effective Guidance

Outline *In order to answer S3 – "How is it possible to design effective guidance?" – in this section, we describe the research we pursued to define the requirements and the process of designing effective guidance in VA. This section is based on [Cen+20], which is reported in its entirety in Chapter 4.*

As we have seen in the previous sections, no single approach in the literature covers the whole spectrum of guidance characteristics nor offers assistance to each and every VA step. We think that this limitation hinders the real effectiveness of guidance approaches in practice.

In an effort to take a first step toward an effective and comprehensive guidance in VA, in the remainder of this section, we describe our work to define the characteristics of *effective guidance* and illustrate a procedural framework to design it. We conclude the section illustrating a design walk-through that should illustrate how the framework can be used in practice.

1.8.1 Requirements for Effective Guidance

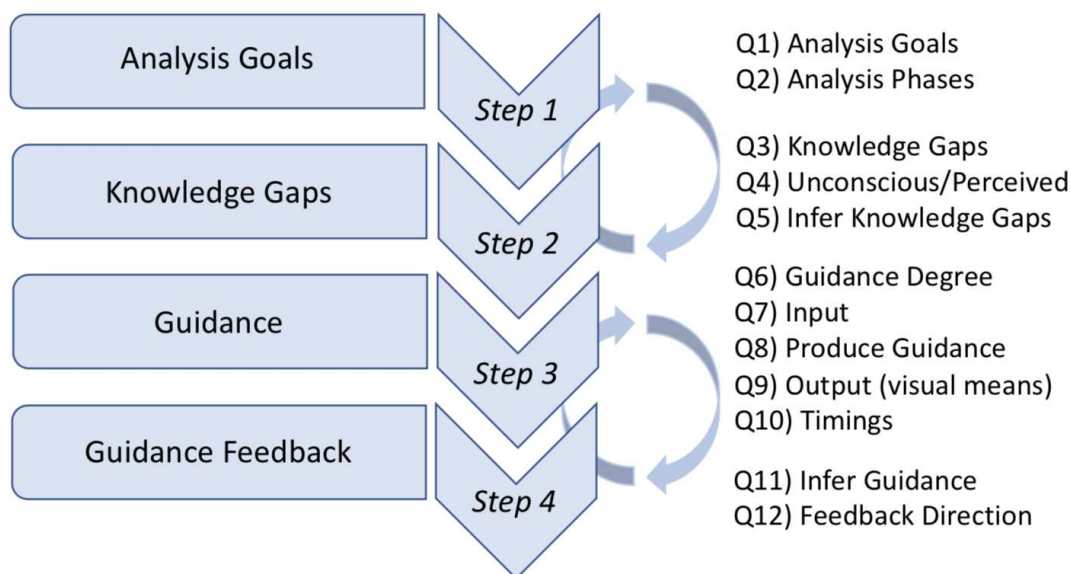
R0 - *Effective* support is what we see as the end goal of guidance-enriched VA.

Effective guidance refers to mechanisms that should help analysts complete a given task while overcoming possible issues that could arise during the process. More formally, effective guidance is help that is able to solve a knowledge gap. If we think of the guidance process, many things could go wrong as many factors are involved in the way guidance is administered to the users. These factors could all play a role in the resulting effectiveness. For instance, one factor is allowing the user to control the degree of guidance; if the analysts cannot control the degree of guidance, they could become frustrated and stop using it, rendering the guidance totally ineffective. Even more, if the guidance, or the way it is produced, is not *trustworthy*, the suggestions provided will not be used by the user.

From an initial discussion in [CGM19a] that was later expanded in [Cen+20], we list a set of qualities that influence the effectiveness of guidance in practice and that should be taken into account when designing guidance-enriched VA approaches. In summary, in order to be effective, guidance has to be:

- R1 **Available:** *Guidance is there for you.* Users should be aware that guidance is available and that support can be provided or requested at any time. Designers should make available interactive means to request guidance and appropriate visual means to convey guidance.
- R2 **Trustworthy:** *Guidance will help you.* Any generic data analysis task includes a certain degree of variability. Guidance should be regarded as a support to overcome the uncertainty involved and not be a source of further confusion. Designers should take care of specific ways to encode and provide guidance to make it trustworthy and accepted by users without being an additional source of misinformation. Trust, once lost, is hard to restore.
- R3 **Adaptive:** *Guidance will adapt to the situation.* Usually, as the analysis evolves, so do the problems users encounter. The guidance system must know what the actual state of the analysis is in order to deal with dynamically changing knowledge gaps. Designers should implement mechanisms to capture the analysis phase, provide interfaces for inferring the knowledge gap, and provide guidance accordingly.
- R4 **Controllable:** *Guidance can be tuned if necessary and the user needs to be in control of the analysis.* Guidance is a mixed-initiative process [Hor99b; Cen+18a]. Therefore, the designed solution should enable users to steer the analysis, choose between alternative recommendations, turn off the guidance if not needed, and provide means to ask for assistance in the first place.

Guidance Design Framework



Effective Guidance




Figure 1.20: The guidance design framework. The framework aims to support the design of effective guidance. We list a set of steps (Step 1–4) as well as quality criteria (R1–R5) that should guide designers during the design process. .

R5 **Non-disruptive**: *Guidance will not annoy or mislead you.* A final quality that we expect to be supported by the guidance process is that it should not disrupt the analysis flow and the analysts’ mental map. The guidance should be provided without requiring users to exit their state of flow.

1.8.2 A Framework for Guidance Designers

Having described the meaning of effective guidance and the complementary qualitative criteria and requirements that concur with its realization, we studied how guidance could be designed in practice using a framework upholding the given quality criteria.

Our work led us to detail a framework composed of four nested steps that aim to identify and delineate a set of issues (i.e., knowledge gaps) that the user might encounter during the analysis. The framework is portrayed in Figure 1.20. Given these gaps, the framework pushes the designer to design appropriate countermeasures (e.g., guidance suggestions) to solve them together with all the details to implement the guidance in a specific analysis environment. The quality criteria mentioned earlier should serve as guidelines for choosing among different design options.

Step 1: Analysis Goals

When approaching the task of designing effective guidance, the first thing to do is identify the analysis goals. These should be fully described not only in general terms, but also (when possible) in full details. Being able to characterize subtasks and subgoals could improve the outcome of the process greatly. The first questions that designers should answer are:

Q1 - What are the analysis goals?

Q2 - In which analysis phases issues might occur?

The VA process can be decomposed in different phases, as detailed in Section 1.3. For instance, in the first phase of the analysis the data is typically preprocessed and prepared for the subsequent exploration. Problems arising in all such phases are those that we want to address with guidance. Breaking the analysis down into its steps will allow designers to think of guidance for easier-to-solve problems that, when combined, will concur to solve more complex issues.

To reason about the analysis and its fundamental phases and determine the analysis goals, any model of VA can be used [And+18; Sac+14b; CMS99].

Step 1 Risks, Threats, and Countermeasures. While completing this first design step, we should be careful to consider the following risks:

Threats and Risks to Step 1: Underestimation and overestimation of the analysis goals are major threats to this step. The designer may identify too many or too few activities/tasks/goals requiring guidance. This may lead to the design of insufficient (a lack of proper support for critical tasks) or excessive guidance (users could be frustrated).

Possible Solutions : Collaborate with domain experts who could provide crucial information about how to structure the tasks and support the identification of analysis objectives. The implementation of means for the user to control and finetune the guidance, as required by R4, could also be a viable solution to counteract such threats.

Step 2: Knowledge Gap

After identifying the analysis goal and structuring it into its fundamental parts, guidance designers need to focus on possible knowledge gaps that could affect the user on the path to the analysis goals.

Q3 - What knowledge gaps might hinder the analyst from proceeding with the analysis?

Types of Knowledge Gaps. One way to complete this step is by reasoning — namely, for each goal and subtask, what are the types of knowledge gaps that could hinder the completion of the analysis? As previously discussed, designers should consider a number of knowledge gaps:

- 1) Data:** This knowledge gap is related to the lack of knowledge about the data, which typically affects the data preprocessing phases. However, a knowledge gap in the data domain may also affect subsequent phases, such as data exploration, if the user does not know what data to explore.
- 2) Tasks:** This knowledge gap is related to the lack of knowledge (i.e., what are the steps) to complete a sub-task. This knowledge gap could also refer to the definition of tasks and goal themselves.
- 3) VA Methods/Algorithms:** This knowledge gap is related to the choice of appropriate visual and analytical methods to apply to the data.
- 4) Knowledge and Insights Management:** This knowledge gap is related to the lack of expertise in interpreting patterns or issues managing the knowledge itself.

While the main aim of Step 2 is defining the knowledge gap, a number of accessory questions help realize a good design.

Perceived and Unconscious Knowledge Gaps. Related to defining the knowledge gap is defining methods to understand if the analyst is aware of the detected knowledge gap(s).

Q4 - Are analysts aware or unaware of their knowledge gaps?

Simply speaking, a perceived knowledge gap is one of which analysts are aware. The opposite is considered an unconscious knowledge gap. For instance, an unconscious knowledge gap could be related to analysts being unaware of missing values, noise, or quality problems in the data. In general, unconscious knowledge gaps could have a negative impact on the analysis and, hence, should be carefully addressed.

Identification of the Knowledge Gap. Although, with Q3, we reason generically about all possible knowledge gaps, from a practical point of view, designers should think of methods to identify the knowledge gap as the analysts pursue their goals.

Q5 - How can potential knowledge gaps be identified during the analysis?

In summary, to answer this question, designers have two alternatives:

- **Knowledge Gap Interface:** Enables the analysts to communicate the knowledge gap to the system.
- **Knowledge Gap Inference:** Enables the system to derive the knowledge gap from analysts' behavior.

The easiest solution is to let the analysts inform the system about the issues they are directly experiencing. In this first case, designers need to implement a knowledge gap interface through which the user can communicate with the system about possible problems. Obviously, this could only work if the analysts are aware of the issues in the first place, as detailed in Q4. The other alternative is that the knowledge gap is indirectly inferred from the actions of the analysts using an inference mechanism.

Step 2 Risks, Threats, and Countermeasures. Similar to Step 1, major threats to the proper conclusion of Step 2 are as follows:

Threats and Risks to Step 2. Underestimation and overestimation of the analysis goals are also major threats to this step. Generally speaking, these problems are related to the completeness of the designed guidance solution.

Possible Solutions. In an attempt to minimize this risk, we propose the following countermeasures:

1. **Design for the top-N knowledge gaps:** To foster design completeness, as an initial step, designers could think of ways to design guidance for the most problematic (top-N) knowledge gaps.
2. **Design adaptive guidance:** In a more advanced scenario, designers could aim for adaptive guidance mechanisms. This allows designers to define only the boundaries within which the guidance can be provided.
3. **Let the analysts guide themselves:** The generation of dynamic guidance cannot always be pursued. Hence, designers should aim for mechanisms to help analysts guide themselves. This corresponds to providing the user with the minimum amount of assistance, putting them in the position to make the best informed choice. This could be helpful in situations like exploratory analysis, when analysis goals cannot be precisely defined, but users might still need guidance.

Step 3: Guidance Generation

Having characterized the knowledge gaps, we have to define guidance to solve them. This will be the goal of the third design step.

Guidance Characteristics We start by reasoning about the guidance we want to obtain, based on the characteristics of the knowledge gaps identified earlier. **Guidance Degree.** Initially, we should reason about how much guidance is necessary to solve the knowledge gap.

Q6 - What degree of guidance is needed? What mechanisms can be employed to switch among different degrees?

Guidance Input. In a subsequent step, we should see what inputs are available to produce the necessary guidance type.

Q7 - What input is available?

In addition to the input sources described in Section 1.4, designers should consider user preferences and possible subjective biases.

Algorithms and Procedures to Calculate Guidance. Knowing about the possible input available and the characteristics of the guidance we want to obtain, we can finally identify suitable algorithms to produce it. This corresponds to answering the following question:

Q8 - What algorithms and procedures are needed to generate guidance?

Algorithms for producing guidance refer to how guidance is generated. As we have seen in Section 1.5, the choice of a specific algorithm depends on the specific scenario and what guidance type designers want to implement. Unfortunately, a complete list of algorithms to generate guidance is beyond the scope of this thesis.

Guidance Output. Once produced, the guidance output must be provided to the analyst. The use of visual means is typical:

Q9 - What are appropriate means to communicate the guidance output?

Appropriate means need to be selected while also considering the existing visual encoding of the data. As we have detailed in Section 1.5, the guidance designer may choose to provide the computed guidance in the form of simple text. Other frequently used expedients to convey guidance are highlighting and changing the color of interesting data items [HB05]. Motion and animation could also be used to communicate guidance [Joh+05]. Finally, glyphs and visual artifacts are the most common way for encoding guidance suggestions [CGM19a].

Identification of the Moment to Provide Guidance. The final challenge in defining appropriate guidance is to identify the moment to provide it:

Q10 - When should the guidance be provided?

Although the opposite might seem true, it is not possible to provide guidance at any time during the analysis. A moment, or a period, does exist when providing guidance is considered desirable. In all other moments, providing guidance could have adverse effects, like frustrating the user or interfering with the current line of analysis [PC18]. This challenge is strictly related to the research on "interruptibility" which has been studied in the field of HCI [DCG09; Lit+11; TO17; QWJ04; BLC20; Züg+18; LCD15; Ras18; SVV99; CCH00].

Step 3 Risks, Threats, and Countermeasures. A non-satisfactory execution of this design step could derive from the following threats. In general, a non-satisfactory completion of this step could hinder the implementation of trustworthy (R2), adaptive (R3), and non-disruptive (R5) guidance.

Threat 1: Introduction of biases.

Possible Solutions. Biases represent a systematic deviation from what is generally recognized as rational judgment. Unfortunately, biases are subjective by nature, and a generalized solution cannot be devised. As general advice, it is recommendable to conduct the design in collaboration with the end-users. Considering iterative cycles, alternating design and evaluation phases could also help mitigate such problems.

Threat 2: Choice of a wrong guidance degree.

Possible Solutions. Implement strategies and mechanisms to enable the user to change the guidance degree, when possible.

Threats 3: Choice of a wrong timing.

Possible Solutions. Although guidance might theoretically be required at any time, it is worth mentioning that, in practical analysis scenarios, it is indeed likely that guidance is needed only at distinct points in time. These moments, in fact, correspond to those situations in which the user is required to make a decision or judgment [Sil91]. Therefore, to select the appropriate timing for guidance, the role of designers is to identify these decisional points in the analysis, as well as what constitutes a critical decision, in the first place. Providing guidance and limiting the alternatives available at such points can make a difference in a successful analysis.

Step 4: Guidance Feedback Loop

As we mentioned, guidance is a mixed-initiative approach [Hor99b]. Step 4 aims to close the guidance loop, enabling users to steer and provide feedback to the guidance system in such a way that the whole analysis is controllable and the user can address it according to his/her goals. Following our discussion in Section 1.6, designers should think of two main aspects: (1) mechanisms to derive guidance for the system from analysts' actions (usually in the form of feedback), which we referred to as the guidance inference mechanism; and (2) the direction of such guidance as it can be directed toward the past (i.e., feedback) or the future (i.e., feedforward).

Inferring Guidance for the System

Q11 - How can the system derive guidance from the analyst's actions?

Interaction sequences can be used by the analyst to provide feedback about the received guidance. Two kinds of feedback can be identified:

- **Direct feedback:** the analyst moves sliders or uses other controls to change the guidance parameters directly.
- **Indirect feedback:** the analyst acts on the data. Analysts move the data, group the data, and label the data, which affects the guidance algorithms indirectly.

Through direct interaction with the user interface and widgets, analysts can finetune the guidance process. For instance, if analysts are not satisfied with the clustering suggested by the system, they may use interaction to adjust the results. The designer should implement mechanisms to integrate this additional input in future guidance suggestions and, sticking with the clustering example, adapt the suggested grouping and improve future guidance output.

The second feedback method is *indirect* [End+14] feedback. This happens when the users' feedback is derived from their interaction with the data (and not with the widgets or with the user interface).

Direction of Feedback. In a last design step, designers should reason about the direction of users' feedback:

Q12 - What is the direction of the analysts' feedback?

As mentioned, the guidance directions can be past and future (i.e., feedback and feedforward, respectively, as discussed in Section 1.6). Similar to relevance feedback, positive (negative) feedback is meant to provide a positive (negative) evaluation of the guidance the system has provided in the previous analysis loop. Feedforward actions, either positive or negative, should instead enable analysts to provide hints for how the guidance should occur in future analysis cycles.

Step 4 Risks, Threats, and Countermeasures. We conclude this section with an analysis of the threats to this design step.

Threats and Risks to Step 4. A major threat to the design is the unsatisfactory realization of R4: controllable guidance. In other words, the system does not allow the user to steer the course of the analysis.

Possible Solutions. In some situations, limiting the alternatives available to the analysts can be desirable, such as when providing prescribing guidance. However, this cannot be assumed as a general design pattern. Designers should find a suitable balance between restricting and guiding the analyst.

1.8.3 Applying the Framework: A case-study

Having described the design framework, we show how it can be applied with a design walk-through. In the following discussion, we illustrate how a designer carried out a real design of guidance following the four design steps introduced herein (see Table 1.3 for a summary). This section constitutes a summary of a more articulated discussion about the applicability and goodness of fit of our guidance design framework. The reader can find additional means to evaluate our framework and further design examples in Chapter 4.

	Questions
Q1	What are the analysis goals?
Q2	In which analysis phases issues might occur?
Q3	What knowledge gaps might hinder the analyst from proceeding with the analysis?
Q4	Are analysts aware or unaware of their knowledge gaps?
Q5	How can potential knowledge gaps be identified during the analysis?
Q6	What degree of guidance is needed?
Q7	What input is available?
Q8	What algorithms and procedures are needed to generate guidance?
Q9	What are appropriate means to communicate the guidance output?
Q10	When should the guidance be provided?
Q11	How can the system derive guidance from the analyst's actions?
Q12	What is the direction of the analysts' feedback?

Table 1.3: Summary of the questions designers should ask when designing guidance.

Problem Description. We grounded the design walk-through in the field of statistics. In particular, we applied our framework to design guidance in a parallel project we ran in our research group. We describe how a designer designed guidance to support a blind-source separation task (BSS). This task can be generally formulated as separating a generic signal into its components without any assumption about the characteristics of the original signal. This is of practical interest in various disciplines. We can imagine, for instance, the need to separate and discern mixed signals composed by the heartbeats of a mother and her child during pregnancy.

Step 1: Analysis Goal. The guidance designer initiated the design, gaining confidence with the topic. He set up a series of interviews with the statisticians to answer Q1 and Q2. Through the interviews, he learned that typically, without guidance, BSS tasks are completed using only R [Tea14]. In particular, in a typical BSS workflow, the data is processed by appositely created algorithms that decompose the original measurements into their components. Later, the output is visualized using static visual representations and undergoes additional inspections. If the result (i.e., the original signals) is sufficiently precise, the task can be considered concluded. However, it is typical for the statisticians (i.e., the users) to iterate over and over the aforementioned phases, comparing the results with those obtained in previous iterations, until a satisfactory approximation is achieved. The goal of the statisticians is not to calculate the optimal solution, which would be unfeasible, but one that represents a very good approximation of the original signal. As it we have seen, the use of visualizations is, yes, already part of the workflow. However, only static images are used. Hence, the designer decided to give additional emphasis to this aspect by introducing a VA methodology and enhancing the workflow with guidance.

Step 2: Knowledge Gap. While dealing with Step 1 to understand the problem, the designer also approached Step 2. Thanks to the interviews, the designer was also able to draft an initial list of knowledge gaps, as required by Q3. In total, the statisticians involved in the project listed three main knowledge gaps (defined as KG-1, KG-2, and KG-3 below).

KG-1: The parameter space is huge, which hinders the possibility of an exhaustive search of the optimal solution.

KG-2: Exploration of the results is error-prone and time consuming since newly obtained data has to be compared with each previously obtained result.

KG-3: Interpretation of the results produced by the BSS algorithms is challenging, as it usually involves specific domain knowledge.

Ultimately, it is clear that the statisticians were aware of the knowledge gaps (Q4). The designer discovered during the interviews that some statisticians relied on a sort of rule-of-thumb methodology to solve the BSS tasks. Therefore, he decided to shed more light on such established but implicit practices and determine if it were possible to formalize and exploit them to solve the knowledge gaps in a more systematic way during the analysis applying a guidance approach.

Step 3: Guidance Generation. Having detailed the knowledge gaps, the designer moved to design solutions to address them. Step 3 was approached by first analyzing the input sources available (Q7). These were identified as the data, the implicit knowledge of the statisticians and, possibly, their domain knowledge, which depends on the specific application scenario.

To address KG-1, the designer designed a guidance approach that can be framed as orienting (Q6) by exploiting the data input (Q7). As the user loads the data, the tool

automatically calculates statistics about the signals and, thanks to them, it can infer what parameters would make sense to use. These "suggestions" about parameter settings can be integrated directly in the sliders that control the selection of those very parameters (see Figure 4.7a). One challenge faced by the statisticians that is strictly connected with KG-1 was selecting parameters during the first phases of the analysis, when almost nothing is known about the data. To support this scenario, the designer exploited the domain knowledge to indicate to the statisticians which parameters could be chosen first, with that specific data and in that specific domain.

Once the parameters have been chosen, the system launches the BSS algorithms, and the statisticians can proceed with the second part of the workflow, which requires comparing, exploring, and interpreting the obtained results. According to KG-2 and KG-3, analysts must discern whether the results obtained make sense, comparing them with those obtained previously. To support these tasks, the designer designed a mechanism consisting of a machine learning algorithm (Q8) that classifies and groups together the signal components produced by the BSS algorithms. This solution is visualized in Figure 4.7b (Q9). Thanks to this guidance, the analyst/statistician is immediately presented with a reasonable classification of the signals and can immediately start reasoning about the obtained results. In addition, the designer provided interaction facilities and visualizations to allow the statisticians to easily compare the obtained signals with those of previous runs. In particular, this was obtained thanks to the superimposition of old and newly obtained results and to the storage of the history of interaction and output results. This additional guidance can be characterized as directing guidance (Q6).

Step 4: Guidance Feedback Loop. Finally, the designer approached Step 4 to define feedback mechanisms and let the statisticians steer the guidance process. As the guidance mechanisms described earlier make strong assumptions about the knowledge that the analyst might need to conclude the analysis, the designer also designed feedback mechanisms to steer the course of the analysis. In particular, whenever a suggestion was provided, it was always possible for the analyst to select something else and change, for instance, the output of the classification algorithms by moving the signals. In addition, the designer also added a mechanism that allows the system to learn. In particular, the system stored all the produced results for reusing them in future analyses. For instance, if the analysis reached a positive conclusion, the results of the classification algorithms were added to a knowledge base to improve future calculations of the classification algorithms. The same happened when a correct data domain was inferred and proper parameters suggested. All these small details add to the support of R4 (adaptive guidance).

In conclusion, our design walk-through shows how the framework carefully considers all major aspects involved in the design of guidance and how it can be considered a rather complete tool to reason about guidance in VA. As it has been written, the design of guidance poses many challenges and requires designers to foresee issues arising during the analysis. Our framework helps in this respect, as it encourages designers to consider all these aspects thoroughly in a step-wise process.

Additional details can be found in Sections 4.6 and 4.6.4.

1.9 Scientific Contributions and Future Perspectives

1.9.1 Contributions

Thus far, we have described our work and the research we pursued to characterize guidance-enriched VA approaches. In Section 1.4, we began describing our work to answer S1:

S1 Is it possible to devise a general framework and a common guidance definition embodying the current state-of-the-art approaches and literature?

We formalized guidance as a computer-assisted process aimed at solving a knowledge gap encountered by the users of an interactive VA system [Cen+17b]. Our work did not stop at the definition of guidance; we also presented a set of characteristics of the guidance process (i.e., the knowledge gap, the input, the output, and the guidance degree), which we combined in a conceptual model integrating VA and guidance, as shown in Figure 1.6.

In a further effort to exhaustively answer S1, we adjusted our definition of guidance [Cen+17a] by describing its mixed-initiative nature [CGM19a; Cen+18a] and expanding the initial characterization to account for the process in which the user guides the system. In this regard, we described the interactive mechanisms the user can exploit to steer the guidance process—namely, direct and indirect as well as positive and negative feedback and feedforward actions [CGM19a].

We continued with S2:

S2 What are the benefits (if any) and, in general, what are the effects of using guidance during visual analytics?

We designed and conducted two user studies (see Section 1.7). In the first user study, we studied the effects of different degrees of guidance (i.e., no guidance, directing, and prescribing) on users with various degrees of previous knowledge (i.e., expert and novice users) to solve exploratory and domain tasks [CGM19b]. The results show that guidance has positive effects on the analysts, who are less frustrated and more confident about the results when receiving guidance. The results also show how guidance can easily compensate for a lack of operational knowledge, which is required for solving exploratory tasks, but it cannot compensate completely for a lack of domain knowledge. We hypothesized that this is due to the fact that users must possess some domain competences to at least interpret and understand the guidance suggestions. In the second study, we confirmed the positive effects of guidance in terms of the frustration and the general positive feelings guidance can induce if provided during a problematic

phase of the analysis, but also explored how guidance affects the resolution strategies adopted by the users [Cen+18b]. In particular, we observed a switch from a mainly trial-and-error approach, when users were not provided with guidance, to a confirmatory analysis followed by additional shorter exploration sessions, when users were supported with guidance.

Finally, we moved to S3:

S3 How is it possible to design effective guidance to support users throughout the visual analytics process?

We described the characteristics of effective guidance by listing a set of qualitative criteria and requirements that a guidance approach should fulfill in order to be effective [Cen+20]. We also introduced a step-wise framework for designing guidance that incorporates such qualitative criteria. The design framework we conceived is an iterative process composed of four steps that consider all the most important questions a designer should answer to complete the design. In this regard, the qualitative criteria we listed serve as a guideline to support designers in choosing among multiple design alternatives. Finally, we described how we used our framework to design guidance in a real-life scenario, providing a solid basis for the evaluation of the design framework itself.

Hence, we can return to our main research question, introduced at the beginning of this thesis:

- *How can we devise guidance methods for supporting users performing visual analytics tasks?*

In this thesis, we have described the characteristics, effects, and design process of guidance, with a clear focus on VA (guidance-enriched VA). We have proposed guidance as a promising way to enable a better collaboration between the human and the computer, supporting the mixed-initiative process that has long been advocated in the literature as a way to achieve successful data analysis while gaining insights.

In other words, we have introduced guidance as a solution to the issues encountered by users performing VA and have paved the way for the adoption of effective guidance-enriched VA approaches.

We conclude by listing possible future directions for our research.

1.9.2 Future Directions

Detecting the Knowledge Gap

We have discussed in detail the characteristics of the guidance process and, in particular, the knowledge gap. How the knowledge gap can be determined and captured during

the analysis is an open challenge in contemporary literature. As the knowledge gap is the focus of any guidance approach, it is clear why its identification is crucial for designing a successful guidance solution. In the past, extensive work has been done to infer knowledge gaps a priori. For instance, Wang et al. hypothesise that ambiguous areas of a graph might need to be explored carefully, hence guiding the exploration and the user towards these precises ambiguous spots [Wan+16]. In other instances, ontologies and taxonomies have been exploited to match the user interaction with known tasks and behaviors [GW09]. However, how the knowledge gap can be properly inferred during the analysis is a challenge that has not yet been completely addressed. The research has also led to the use of a posteriori strategies that employ pattern-matching algorithms to compare the interaction history of the user with known patterns of tasks and known behaviors to understand the user's intent [GW09]. However, these strategies have their limitation in that exploratory tasks, which are typical in VA analyses, do not possess a clear structuring. Indeed, detected patterns of behavior may have different meanings according to when those actions were taken and the level of granularity used for grouping those users' actions.

Closely related to the detection of the knowledge gap is how it can be conveyed to the system. The research on mechanisms for the analyst to communicate knowledge gaps is still far from providing a definite answer to this question. This involves finding ways to encode knowledge and communicate it effectively to a computational system.

Providing Timely Guidance

Guidance is effective only if provided at the right moment. A further challenge is how the precise moment to provide guidance to the user can be determined in practice during the analysis. The literature describes approaches that assumed that guidance could be provided immediately after a problem (i.e., a knowledge gap) was detected [Sil91]. In other scenarios, literature approaches looked at implicit signals to administer guidance. For instance, the inactivity of the user is quite often interpreted as a signal that the user needed support [Hor+98]. However, these assumptions are not representative of real analysis scenarios. How can we be sure that the inactivity does not simply represent a moment of reasoning or that the user only wants a pause from the analysis? What's the threshold that separates a normal analysis and a stalled analysis? A promising way to solve this problem could come directly from the user. As we have seen, emotions are deeply related to a need for guidance. Frustration and sadness grow, for instance, when users cannot solve given tasks [PC18]. Hence, a viable solution could be using such emotions to determine the moment at which guidance is needed. The challenge at that point would be determining the thresholds (e.g., of frustration), above which guidance should be triggered.

Practical Guidelines for Guidance Designers

With our theoretical and conceptual foundation of guidance, we explored and described the possibilities offered by guidance approaches in VA. Although guidance aims to support

the user when using VA tools, we have not yet considered guidelines to apply in the development phase of VA. We already completed a first step in this direction with our work describing a design framework for guidance [Cen+20]. However, how this framework can be bridged with practice and applied to different scenarios to solve task-specific issues is still an open challenge in our field. In particular, when designing guidance for different data (e.g., time-oriented data) and different tasks, it can be difficult for designers to choose an appropriate guidance technique for a specific problem from among many alternatives. Therefore, it is important to provide both guidance for users and practical guidelines for developers. Such guidelines could, for example, suggest how certain scenarios are best supported with a certain degree of guidance.

Evaluating Effectiveness

Evaluating visualization techniques is a notoriously challenging task. Evaluating VA is even more challenging as it requires us to look at not only the appropriateness of the chosen visual encoding, but also how the visualizations are related to the computational models used and how they all together concur with the completion of the analysis. It is easy to imagine that evaluating guidance-enriched VA would make the evaluation task even more complex and challenging. In our work, we defined the requirements for effective guidance [Cen+20; CGM19a]. Effective guidance can be intuitively seen as a way to overcome users' issues. However, how the effectiveness of the designed guidance could be evaluated in practice is still an open challenge in the field. Metrics typically used to establish the goodness of fit of a visual approach, like performance improvements or better completion timings, might be insufficient for evaluating if guidance was successful. One available option for this research direction could be evaluating if guidance enables users to make unexpected discoveries. Another viable solution could be to abandon the classic performance metric, reestablish users' central role in the analysis, and evaluate how guidance affects their reasoning process, going beyond current evaluative methods. Similarly to what we did in the works described in this thesis, the effectiveness of guidance could be evaluated by looking at how it is able to affect, and especially contribute to improving, the user's state of mind to establish a beneficial user experience.

Comprehensive Guidance in Visual Analytics

All the previous challenges would concur with reaching the overarching goal of comprehensive guidance in VA. As we have stated many times, no single approach in the literature covers the whole spectrum of possibilities offered by guidance. Table 1.1 clearly shows this point. Guidance is never provided throughout the analysis, but approaches tackle only single problems arising at precise points of the analysis. Typical techniques suggest how the user should clean the data or which data to inspect, but no approach extends guidance to more than one knowledge gap. For this reason, we need what we may call comprehensive guidance. Our end goal is to push the research toward a deeper integration of guidance and VA, starting from the initial phases of its design, which is crucial for effective support in complex real-life scenarios. With our design framework,

we completed a first step in this direction [Cen+20]. Our initial work describes the foundations upon which comprehensive guidance can be designed and developed. However, it does not exhaust the overall problem in that it does not specify how the guidance and the VA process can be integrated. A promising idea comes directly from our work, which intuitively relates, through the second design step, the design of guidance for the VA process. The idea is to use the VA process, such as described by Sacha et al. [Sac+14b], as a horizontal design axis of an integrated design process, while we could use our design framework to provide details about how the different phases of the analysis could be enhanced with guidance.

Generating Guidance

What algorithms and methods exist to generate guidance? In Section 1.5, we investigated how guidance is produced and generated by literature approaches. However, the exact algorithms and methods are scenario- and task-specific. Developing ontologies and task taxonomies could be helpful for determining the kind of guidance necessary for specific situations.

Managing the Guidance Degree

As we have seen in Section 1.7.1, there is a natural correlation among the user's knowledge, the task, and the degree of guidance needed to complete it. Therefore, we need guidelines and methodologies to determine, given a knowledge gap, a suitable degree of guidance that neither frustrates nor limits the user's interactions. A further challenge is allowing a smooth transition between guidance degrees. Guidance is defined as a dynamic process that typically adapts to the situation at hand. Thus, as the analysis proceeds, so do the needs of the user and so should guidance. Hence, in addition to an appropriate degree of guidance, we also need to determine how to switch among the degrees so as not to distract or abruptly interrupt the analysis.

1.10 Publications

The results of our work have been published mainly in international scientific journals, such as Computer Graphics Forum (CGF), Transactions on Visualization and Computer Graphics (TVCG), and Visual Informatics. We also presented our findings at specialized conferences, like IEEE VIS, EuroVis, and EuroVA. The publications themselves were always realized in collaboration with great fellow scientists. This also helped us build a noteworthy international scientific network and contributed to the dissemination of the results. The following is a full list of publications published in the last five years, most recently updated in July 2020.

1.10.1 Cumulative Journal Publications

The following three journal articles are reported in their entirety in Chapters 2, 3 and 4 as part of this cumulative dissertation:

- [Cen+17b] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-j. Schulz, M. Streit, and C. Tominski. „Characterizing Guidance in Visual Analytics“. In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (Jan. 2017), pp. 111–120. ISSN: 1077-2626. DOI: 10.1109/TVCG.2016.2598468

My Contributions: This paper describes the characterization of guidance in VA. I had the scientific and organizational lead of this paper. Being this my first paper, the lead was done in cooperation with senior co-authors. Initially, the paper was considered to be a state-of-the-art report, so it had a different structure and a different approach to the problem and presented a slightly different perspective on guidance approaches. The initial work was merged with the co-authors contributions into a paper’s new structure and submitted to IEEE VAST. In particular, I contributed to defining the guidance input, shaping and refining the guidance degrees, and, in part, initially characterizing the knowledge gaps. I also contributed an initial definition of guidance, taking inspiration from the literature and work in other communities. I drafted the initial version of the paper and contributed ideas for the discussion and future work section. I was responsible for the consistency of the whole paper, for the submission process as well as for merging the different inputs of the co-authors. Finally, I would like to thank all the co-authors for their critical reflections and contributions.

- [CGM19b] D. Ceneda, T. Gschwandtner, and S. Miksch. „You get by with a little help: The effects of variable guidance degrees on performance and mental state“. In: *Visual Informatics* 3.4 (2019), pp. 177–191. ISSN: 2468-502X

My Contributions: I was the lead author of this paper, which describes a user study to investigate the effects of guidance on VA users in different analysis scenarios. I designed, structured, implemented, and conducted the study and was responsible for the statistical analysis of the results. I also contributed the idea of focusing the study not only on performance metrics, but also showing the effects of guidance on the user’s mental state, which is crucial for instilling a positive analysis environment. I drafted the paper and was responsible for the entire submission process. Finally, I would like to thank all the co-authors for their critical reflections and contributions.

- [Cen+20] D. Ceneda, N. Andrienko, G. Andrienko, T. Gschwandtner, S. Miksch, N. Piccolotto, T. Schreck, M. Streit, J. Suschnigg, and C. Tominski. „Guide Me in Analysis: A Framework for Guidance Designers“. In: *Computer Graphics Forum* 39.6 (2020), pp. 269–288. DOI: 10.1111/cgf.14017

My Contributions: I was the lead author of this paper, which describes a framework for designing guidance and a set of qualitative criteria that have to be sustained during the design phase to obtain effective guidance. I contributed the initial idea of a design framework and defined the initial list of design steps. Working with the design phases, I also identified a set of risks and issues that might arise during their implementation, as well as possible measures to counteract them. I also contributed an initial list of quality criteria that was partly drafted in [CGM19a]. I was responsible for conducting and describing the design walk-through that showed the applicability of the framework in real analysis scenarios. Finally, I would like to thank all the co-authors for their critical reflections and contributions.

1.10.2 Additional Publications

Additional articles that are included in this thesis:

- [Cen+18a] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, M. Streit, and C. Tominski. „Guidance or No Guidance? A Decision Tree Can Help“. In: *Proc. of the International Workshop on Visual Analytics (EuroVA)*. Eurographics Digital Library, 2018, 19–23. DOI: 10.2312/eurova.20181107
- [CGM19a] D. Ceneda, T. Gschwandtner, and S. Miksch. „A review of guidance approaches in visual data analysis: A multifocal perspective“. In: *Computer Graphics Forum* 38.3 (2019), pp. 861–879
- [Cen+18b] D. Ceneda, T. Gschwandtner, S. Miksch, and C. Tominski. *Guided Visual Exploration of Cyclical Patterns in Time-series*. Visualization in Data Science (VDS at IEEE VIS 2018). 2018

Bibliography

- [Adl+10] S. Adler et al. „Overlay of patient-specific anatomical data for advanced navigation in surgery simulation“. In: *Proc. of the First International Workshop on Digital Engineering*. ACM. 2010, pp. 52–58.
- [AEK00] M. Ankerst, M. Ester, and H.-P. Kriegel. „Towards an effective cooperation of the user and the computer for classification“. In: *Proc. of the Workshop on Visual Analytics and Knowledge Discovery: Integrating Automated Analysis with Interactive Exploration*. ACM. 2000, pp. 179–188.
- [AES05] R. Amar, J. Eagan, and J. Stasko. „Low-level components of analytic activity in information visualization“. In: *Proc. of IEEE Symposium on Information Visualization*. IEEE. 2005, pp. 111–117.
- [Als+11] J. Alsakran et al. „STREAMIT: Dynamic visualization and interactive exploration of text streams“. In: *Proc. of the IEEE Pacific Symposium on Visualization*. IEEE. 2011, pp. 131–138.
- [And+09] G. Andrienko et al. „Interactive visual clustering of large collections of trajectories“. In: *Proc. of IEEE Symposium on Visual Analytics Science and Technology*. IEEE. 2009, pp. 3–10.
- [And+11] G. Andrienko et al. „A conceptual framework and taxonomy of techniques for analyzing movement“. In: *Journal of Visual Languages & Computing* 22.3 (2011), pp. 213–232.
- [And+18] N. Andrienko et al. „Viewing Visual Analytics as Model Building“. In: *Computer Graphics Forum* 37.6 (Jan. 2018), pp. 275–299.
- [AP13] D. Archambault and H. C. Purchase. „Mental map preservation helps user orientation in dynamic graphs“. In: *Graph Drawing*. Springer, 2013, pp. 475–486. DOI: 10.1007/978-3-642-36763-2_42.
- [AS03] C. Ahlberg and B. Shneiderman. „Visual information seeking: Tight coupling of dynamic query filters with starfield displays“. In: *The Craft of Information Visualization*. Elsevier, 2003, pp. 7–13.
- [Ber+12a] J. Bernard et al. „Visual-Interactive Preprocessing of Time Series Data“. In: *Proc. of SIGRAD*. Vol. 81. Linköping University Electronic Press, 2012, pp. 39–48. URL: http://www.ep.liu.se/ecp_article/index.en.aspx?issue=081;article=006.

- [Ber+12b] J. Bernard et al. „Visual-interactive preprocessing of time series data“. In: *Proc. of SIGRAD*. Vol. 081. Linköping University Electronic Press. 2012, pp. 39–48.
- [Ber+14] J. Bernard et al. „Visual-interactive Exploration of Interesting Multivariate Relations in Mixed Research Data Sets“. In: *Computer Graphics Forum*. Vol. 33. Wiley Online Library. 2014, pp. 291–300.
- [Ber+17] J. Bernard et al. „Visual-interactive semi-supervised labeling of human motion capture data“. In: *Electronic Imaging 2017.1* (2017), pp. 34–45.
- [BGV16] F. Bouali, A. Guettala, and G. Venturini. „VizAssist: an interactive user assistant for visual data mining“. In: *The Visual Computer* 32.11 (2016), pp. 1447–1463.
- [BH19] L. Battle and J. Heer. „Characterizing exploratory visual analysis: A literature review and evaluation of analytic provenance in tableau“. In: *Computer Graphics Forum*. Vol. 38. 3. Wiley Online Library. 2019, pp. 145–159.
- [BL09] E. Bertini and D. Lalanne. „Surveying the complementary role of automatic data analysis and visualization in knowledge discovery“. In: *Proc. of the Workshop on Visual Analytics and Knowledge Discovery: Integrating Automated Analysis with Interactive Exploration*. ACM. 2009, pp. 12–20.
- [BLC20] O. Barral, S. Lallé, and C. Conati. „Understanding the effectiveness of adaptive guidance for narrative visualization: a gaze-based analysis“. In: *Proceedings of the 25th International Conference on Intelligent User Interfaces*. 2020, pp. 1–9.
- [BPH05] A. Bernstein, F. Provost, and S. Hill. „Toward intelligent assistance for a data mining process: An ontology-based approach for cost-sensitive classification“. In: *IEEE Transactions on Knowledge and Data Engineering* 17.4 (2005), pp. 503–518. DOI: 10.1109/TKDE.2005.67.
- [BS06] E. Bertini and G. Santucci. „Give chance a chance: modeling density to enhance scatter plot quality through random data sampling“. In: *Information Visualization* 5.2 (2006), pp. 95–110.
- [CBY10] Y. Chen, S. Barlowe, and J. Yang. „Click2annotate: Automated insight externalization with rich semantics“. In: *Proc. of IEEE Symposium on Visual Analytics Science and Technology*. IEEE. 2010, pp. 155–162.
- [CCH00] M. Czerwinski, E. Cutrell, and E. Horvitz. „Instant messaging: Effects of relevance and timing“. In: *People and computers XIV: Proceedings of HCI*. Vol. 2. 2000, pp. 71–76.
- [Cen+17a] D. Ceneda et al. „Amending the Characterization of Guidance in Visual Analytics“. In: *CoRR* abs/1710.06615 (2017). arXiv: 1710.06615. URL: <http://arxiv.org/abs/1710.06615>.

- [Cen+17b] D. Ceneda et al. „Characterizing Guidance in Visual Analytics“. In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (Jan. 2017), pp. 111–120. ISSN: 1077-2626. DOI: 10.1109/TVCG.2016.2598468.
- [Cen+18a] D. Ceneda et al. „Guidance or No Guidance? A Decision Tree Can Help“. In: *Proc. of the International Workshop on Visual Analytics (EuroVA)*. Eurographics Digital Library, 2018, 19–23. DOI: 10.2312/eurova.20181107.
- [Cen+18b] D. Ceneda et al. *Guided Visual Exploration of Cyclical Patterns in Time-series*. Visualization in Data Science (VDS at IEEE VIS 2018). 2018.
- [Cen+20] D. Ceneda et al. „Guide Me in Analysis: A Framework for Guidance Designers“. In: *Computer Graphics Forum* 39.6 (2020), pp. 269–288. DOI: 10.1111/cgf.14017.
- [CGM19a] D. Ceneda, T. Gschwandtner, and S. Miksch. „A review of guidance approaches in visual data analysis: A multifocal perspective“. In: *Computer Graphics Forum* 38.3 (2019), pp. 861–879.
- [CGM19b] D. Ceneda, T. Gschwandtner, and S. Miksch. „You get by with a little help: The effects of variable guidance degrees on performance and mental state“. In: *Visual Informatics* 3.4 (2019), pp. 177–191. ISSN: 2468-502X.
- [Che+09] M. Chen et al. „Data, information, and knowledge in visualization“. In: *Proc. of IEEE Computer Graphics and Applications* 29.1 (2009), pp. 12–19. DOI: 10.1109/MCG.2009.6.
- [Che+18] M. Chegini et al. „Interactive Visual Exploration of Local Patterns in Large Scatterplot Spaces“. In: *Computer Graphics Forum*. Vol. 37. Wiley Online Library. 2018, pp. 99–109.
- [Cho+10] J. Choo et al. „iVisClassifier: An interactive visual analytics system for classification based on supervised dimension reduction“. In: *Proc. of IEEE Symposium on Visual Analytics Science and Technology*. IEEE. 2010, pp. 27–34.
- [CMS99] S. K. Card, J. D. Mackinlay, and B. Shneiderman. *Readings in information visualization - using vision to think*. Academic Press, 1999. ISBN: 978-1-55860-533-6.
- [Crn+11] T. Crnovrsanin et al. „Visual recommendations for network navigation“. In: *Computer Graphics Forum*. Vol. 30. Wiley Online Library. 2011, pp. 1081–1090.
- [DCG09] S. K. D’Mello, S. D. Craig, and A. C. Graesser. „Multimethod assessment of affective experience and expression during deep learning“. In: *International Journal of Learning Technology* 4.3-4 (2009), pp. 165–187.
- [DFB11] S. M. Drucker, D. Fisher, and S. Basu. „Helping users sort faster with adaptive machine learning recommendations“. In: *International Conference on Human-Computer Interaction*. Springer. 2011, pp. 187–203.

- [Dix+04] A. Dix et al. *Human-Computer Interaction*. 3rd. Pearson Education, 2004.
- [DLB13] M. Dörk, H. Lam, and O. Benjelloun. „Accentuating visualization parameters to guide exploration“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM. 2013, pp. 1755–1760.
- [DR00] M. Derthick and S. Roth. „Data exploration across temporal contexts“. In: *Proc. of Conference on Intelligent User Interfaces*. ACM, 2000, pp. 60–67. DOI: 10.1145/325737.325782.
- [DR01] M. Derthick and S. F. Roth. „Enhancing data exploration with a branching history of user operations“. In: *Knowledge-Based Systems* 14.1-2 (2001), pp. 65–74.
- [EFN12] A. Endert, P. Fiaux, and C. North. „Semantic interaction for visual text analytics“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM. 2012, pp. 473–482.
- [End+14] A. Endert et al. „The human is the loop: new directions for visual analytics“. In: *Journal of Intelligent Information Systems* 43.3 (2014), pp. 411–435.
- [Eng96] R. Engels. „Planning Tasks for Knowledge Discovery in Databases Performing Task-Oriented User-Guidance“. In: *Proc. of Conference on Knowledge Discovery and Data Mining (KDD)*. Portland, Oregon: AAAI Press, 1996, pp. 170–175.
- [Fia12] P. O. Fiaux. „Solving intelligence analysis problems using biclusters“. PhD thesis. Virginia Tech, 2012.
- [FPS+96] U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, et al. „Knowledge Discovery and Data Mining: Towards a Unifying Framework.“ In: *Proc. of Knowledge Discovery in Databases*. Vol. 96. 1996, pp. 82–88.
- [FPS96] U. Fayyad, G. Piatetsky-Shapiro, and P. Smyth. „From data mining to knowledge discovery in databases“. In: *AI magazine* 17.3 (1996), p. 37.
- [Fuj+97] I. Fujishiro et al. „GADGET: Goal-oriented application design guidance for modular visualization environments“. In: *Proc. of IEEE Conference on Visualization*. IEEE, 1997, pp. 245–252. DOI: 10.1109/VISUAL.1997.663889.
- [Got+10] D. Gotz et al. „HARVEST: An intelligent visual analytic tool for the masses“. In: *Proc. of Intelligent Visual Interfaces for Text Analysis*. ACM, 2010, pp. 1–4. DOI: 10.1145/2002353.2002355.
- [Gra+14] S. Gratzl et al. „Domino: Extracting, comparing, and manipulating subsets across multiple tabular datasets“. In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2023–2032. DOI: 10.1109/TVCG.2014.2346260.
- [GRM10] S. Garg, I. Ramakrishnan, and K. Mueller. „A visual analytics approach to model learning“. In: *Proc. of IEEE Symposium on Visual Analytics Science and Technology*. IEEE. 2010, pp. 67–74.

- [GS86] B. R. Gaines and M. L. Shaw. „Foundations of dialog engineering: the development of human-computer interaction. Part II“. In: *International Journal of Man-Machine Studies* 24.2 (1986), pp. 101–123.
- [GST13] S. Gladisch, H. Schumann, and C. Tominski. „Navigation Recommendations for Exploring Hierarchical Graphs“. In: *International Symposium on Visual Computing*. Rethymnon, Crete: Springer, 2013, pp. 36–47. DOI: 10.1007/978-3-642-41939-3_4.
- [GW09] D. Gotz and Z. Wen. „Behavior-driven visualization recommendation“. In: *Proc. of Conference on Intelligent User Interfaces*. ACM, 2009, pp. 315–324. DOI: 10.1145/1502650.1502695.
- [HB05] J. Heer and D. Boyd. „Vizster: Visualizing online social networks“. In: *Proc. of IEEE Symposium on Information Visualization* (2005).
- [Heg04] M. Hegarty. „Diagrams in the mind and in the world: Relations between internal and external visualizations“. In: *International Conference on Theory and Application of Diagrams*. Springer, 2004, pp. 1–13.
- [HH98] D. Heckerman and E. Horvitz. „Inferring informational goals from free-text queries: A Bayesian approach“. In: *Proc. of Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers Inc. 1998, pp. 230–237.
- [HHK15] J. Heer, J. M. Hellerstein, and S. Kandel. „Predictive Interaction for Data Transformation.“ In: *CIDR*. 2015.
- [Hor+98] E. Horvitz et al. „The Lumiere project: Bayesian user modeling for inferring the goals and needs of software users“. In: *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc. 1998, pp. 256–265.
- [Hor99a] E. Horvitz. „Principles of Mixed-Initiative User Interfaces“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1999, pp. 159–166. DOI: 10.1145/302979.303030.
- [Hor99b] E. Horvitz. „Principles of mixed-initiative user interfaces“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM. 1999, pp. 159–166.
- [Hos+11] M. S. Hossain et al. „Helping Intelligence Analysts Make Connections.“ In: *Scalable Integration of Analytics and Visualization* 11 (2011), p. 17.
- [IV11] C. Y. Ip and A. Varshney. „Saliency-assisted navigation of very large landscape images“. In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 1737–1746.
- [Jeo+09] D. H. Jeong et al. „iPCA: An Interactive System for PCA-based Visual Analytics“. In: *Computer Graphics Forum*. Vol. 28. Wiley Online Library, 2009, pp. 767–774.

- [JN15] L. Jiang and A. Nandi. „Snaptoquery: providing interactive feedback during exploratory query specification“. In: *Proc. of the VLDB Endowment* 8.11 (2015), pp. 1250–1261.
- [Joh+05] J. Johansson et al. „Revealing structure within clustered parallel coordinates displays“. In: *Proc. of IEEE Symposium on Information Visualization*. IEEE, 2005, pp. 125–132.
- [Kan+11a] S. Kandel et al. „Research directions in data wrangling: Visualizations and transformations for usable and credible data“. In: *Information Visualization* 10.4 (2011), pp. 271–288.
- [Kan+11b] S. Kandel et al. „Wrangler: Interactive visual specification of data transformation scripts“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2011, pp. 3363–3372.
- [Kan+12] S. Kandel et al. „Profiler: Integrated statistical analysis and visualization for data quality assessment“. In: *Proc. of the International Working Conference on Advanced Visual Interfaces*. ACM, 2012, pp. 547–554.
- [KB06] G. Krishnamoorthy and P. Brusilovsky. „Personalized guidance for example selection in an explanatory visualization system“. In: *E-Learn: World Conference on E-Learning in Corporate, Government, Healthcare, and Higher Education*. Association for the Advancement of Computing in Education (AACE), 2006, pp. 2122–2127.
- [Kei+08a] D. Keim et al. „Visual Analytics: Scope and Challenges“. In: *Visual Data Mining*. Ed. by S. J. Simoff, M. H. Böhlen, and A. Mazeika. Springer, 2008, pp. 76–90. DOI: 10.1007/978-3-540-71080-6_6.
- [Kei+08b] D. A. Keim et al. „Visual analytics: Scope and challenges“. In: *Visual Data Mining*. Springer, 2008, pp. 76–90.
- [Kei+10] D. Keim et al. „Mastering the information age: solving problems with visual analytics“. In: (2010).
- [KNS04] M. Kreuzler, T. Nocke, and H. Schumann. „A history mechanism for visual data mining“. In: *Proc. of IEEE Symposium on Information Visualization*. IEEE, 2004, pp. 49–56. DOI: 10.1109/INFVIS.2004.2.
- [Koo+08] D. Koop et al. „VisComplete: Automating suggestions for visualization pipelines“. In: *IEEE Transactions on Visualization and Computer Graphics* 14.6 (2008), pp. 1691–1698. DOI: 10.1109/TVCG.2008.174.
- [KPN16] J. Krause, A. Perer, and K. Ng. „Interacting with predictions: Visual inspection of black-box machine learning models“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2016, pp. 5686–5697.
- [Lam+11] H. Lam et al. „Empirical studies in information visualization: Seven scenarios“. In: *IEEE Transactions on Visualization and Computer Graphics* 18.9 (2011), pp. 1520–1536.

- [LCD15] H. Lopez-Tovar, A. Charalambous, and J. Dowell. „Managing smartphone interruptions through adaptive modes and modulation of notifications“. In: *Proceedings of the International Conference on Intelligent User Interfaces*. 2015, pp. 296–299.
- [Lex+12] A. Lex et al. „StratomeX: Visual Analysis of Large-Scale Heterogeneous Genomics Data for Cancer Subtype Characterization“. In: *Computer Graphics Forum* 31.3pt3 (2012), pp. 1175–1184. DOI: 10.1111/j.1467-8659.2012.03110.x.
- [Lie+95] H. Lieberman et al. „Letizia: An agent that assists web browsing“. In: *IJCAI* 1995 (1995), pp. 924–929.
- [Lit+11] G. C. Littlewort et al. „Automated measurement of children’s facial expressions during problem solving tasks“. In: *Face and Gesture 2011*. IEEE. 2011, pp. 30–35.
- [Liu+18] S. Liu et al. „Steering data quality with visual analytics: The complexity challenge“. In: *Visual Informatics* 2.4 (2018), pp. 191–197. ISSN: 2468-502X. DOI: <https://doi.org/10.1016/j.visinf.2018.12.001>.
- [LJ12] A. Labrinidis and H. V. Jagadish. „Challenges and opportunities with big data“. In: *Proceedings of the VLDB Endowment* 5.12 (2012), pp. 2032–2033.
- [Lub+12] M. Luboschik et al. „Heterogeneity-based Guidance for Exploring Multiscale Data in Systems Biology“. In: *Proc. of IEEE Symposium on Biological Data Visualization*. IEEE, 2012, pp. 33–40. DOI: 10.1109/BioVis.2012.6378590.
- [MA14] S. Miksch and W. Aigner. „A matter of time: Applying a data–users–tasks design triangle to visual analytics of time-oriented data“. In: *Computers & Graphics* 38 (2014), pp. 286–290. DOI: 10.1016/j.cag.2013.11.002.
- [Mac86] J. Mackinlay. „Automating the design of graphical presentations of relational information“. In: *ACM Transactions on Graphics* 5.2 (1986), pp. 110–141. DOI: 10.1145/22949.22950.
- [Mae94] P. Maes. „Agents that Reduce Work and Information Overload“. In: *Commun. ACM* 37 (1994), pp. 30–40.
- [May+11] T. May et al. „Guiding feature subset selection with an interactive visualization“. In: *Proc. of IEEE Symposium on Visual Analytics Science and Technology*. IEEE. 2011, pp. 111–120.
- [May+12a] T. May et al. „Using Signposts for Navigation in Large Graphs“. In: *Computer Graphics Forum* 31.3pt2 (2012), pp. 985–994. DOI: 10.1111/j.1467-8659.2012.03091.x.
- [May+12b] T. May et al. „Using signposts for navigation in large graphs“. In: *Computer Graphics Forum*. Vol. 31. Wiley Online Library. 2012, pp. 985–994.
- [Maz09] R. Mazza. *Introduction to information visualization*. Springer Science & Business Media, 2009.

- [McC+18] J. McCormack et al. „Multisensory immersive analytics“. In: *Immersive Analytics*. Springer, 2018, pp. 57–94.
- [Men+98] C. Meng et al. „Visualizing histories for selective undo and redo“. In: *International Conference on Human-Computer Interaction*. IEEE. 1998, pp. 459–464.
- [MGW11] M. Migut, J. C. van Gemert, and M. Worring. „Interactive decision making using dissimilarity to visually represented prototypes“. In: *Proc. of IEEE Symposium on Visual Analytics Science and Technology*. IEEE. 2011, pp. 141–149.
- [Mil56] G. A. Miller. „The magical number seven, plus or minus two: Some limits on our capacity for processing information.“ In: *Psychological review* 63.2 (1956), p. 81.
- [Mül+08] E. Müller et al. „Morpheus: interactive exploration of subspace clustering“. In: *Proc. of the Workshop on Visual Analytics and Knowledge Discovery: Integrating Automated Analysis with Interactive Exploration*. ACM. 2008, pp. 1089–1092.
- [Mun09] T. Munzner. „A nested model for visualization design and validation“. In: *IEEE Transactions on Visualization and Computer Graphics* 15.6 (2009), pp. 921–928.
- [MW10] M. Migut and M. Worring. „Visual exploration of classification models for risk assessment“. In: *Proc. of IEEE Symposium on Visual Analytics Science and Technology*. IEEE. 2010, pp. 11–18.
- [OAH15] P. O’Donovan, A. Agarwala, and A. Hertzmann. „Designscape: Design with interactive layout suggestions“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM. 2015, pp. 1221–1224.
- [PC05] P. Pirolli and S. Card. „The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis“. In: *Proc. of the International Conference on Intelligence Analysis*. Vol. 5. McLean, VA, USA. 2005, pp. 2–4.
- [PC18] P. Panwar and C. M. Collins. „Detecting negative emotion for mixed initiative visual analytics“. In: *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. 2018, pp. 1–6.
- [PF91] G. Piatetski and W. Frawley. *Knowledge discovery in databases*. MIT Press, 1991.
- [PHG07] H. C. Purchase, E. Hoggan, and C. Görg. „How important is the “mental map”? – An empirical investigation of a dynamic graph layout algorithm“. In: *Graph Drawing*. Springer, 2007, pp. 184–195. DOI: 10.1007/978-3-540-70904-6_19.

- [Pir+96] P. Pirolli et al. „Scatter/gather browsing communicates the topic structure of a very large text collection“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1996, pp. 213–220.
- [PS08] A. Perer and B. Shneiderman. „Systematic yet flexible discovery: guiding domain experts through exploratory data analysis“. In: *Proc. of Conference on Intelligent User Interfaces*. ACM, 2008, pp. 109–118. DOI: 10.1145/1378773.1378788.
- [QWJ04] L. Qu, N. Wang, and W. L. Johnson. „Choosing when to interact with learners“. In: *Proceedings of the International conference on Intelligent User Interfaces*. 2004, pp. 307–309.
- [Ras18] S. Raschka. „Model evaluation, model selection, and algorithm selection in machine learning“. In: *arXiv preprint arXiv:1811.12808* (2018).
- [Sac+14a] D. Sacha et al. „Feature-driven visual analytics of soccer data“. In: *Proc. of IEEE Symposium on Visual Analytics Science and Technology*. IEEE, 2014, pp. 13–22.
- [Sac+14b] D. Sacha et al. „Knowledge generation model for visual analytics“. In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 1604–1613.
- [SB78] P. G. Sokolove and W. N. Bushell. „The chi square periodogram: its utility for analysis of circadian rhythms“. In: *Journal of theoretical biology* 72.1 (1978), pp. 131–160.
- [Sch+13] H.-J. Schulz et al. *Towards a Characterization of Guidance in Visualization*. Poster at IEEE Conference on Information Visualization (InfoVis). Atlanta, USA, 2013.
- [Sch+17] M. Schwärzler et al. „Sketch-based guided modeling of 3D buildings from oriented photos“. In: *Proceedings of the 21st ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. ACM, 2017, p. 9.
- [SGL09] Y. B. Shrinivasan, D. Gotzy, and J. Lu. „Connecting the dots in visual analysis“. In: *Proc. of IEEE Symposium on Visual Analytics Science and Technology*. IEEE, 2009, pp. 123–130.
- [Shn10] B. Shneiderman. *Designing the user interface: strategies for effective human-computer interaction*. Pearson Education India, 2010.
- [Shn96] B. Shneiderman. „The eyes have it: A task by data type taxonomy for information visualizations“. In: *Proc. of IEEE Symposium on Visual Languages*. IEEE, 1996, pp. 336–343. DOI: 10.1109/VL.1996.545307.
- [Sil91] M. S. Silver. „Decisional guidance for computer-based decision support“. In: *MIS quarterly* (1991), pp. 105–122.
- [SM86] S. Smith and J. Mosier. *Guidelines for Designing User Interface Software*. Tech. rep. ESD-TR-86-278. Mitre Corporation, Bedford MA, 1986.

- [ST15] A. Sarvghad and M. Tory. „Exploiting analysis history to support collaborative data analysis“. In: *Proceedings of the 41st Graphics Interface Conference*. Canadian Information Processing Society. 2015, pp. 123–130.
- [Ste+09] C. A. Steed et al. „Guided analysis of hurricane trends using statistical processes integrated with interactive parallel coordinates“. In: *Proc. of IEEE Symposium on Visual Analytics Science and Technology*. IEEE. 2009, pp. 19–26.
- [Ste+15] M. Stein et al. „Visual soccer analytics: Understanding the characteristics of collective team movement based on feature-driven analysis and abstraction“. In: *ISPRS International Journal of Geo-Information* 4.4 (2015), pp. 2159–2184.
- [Str+12] M. Streit et al. „Model-driven design for the visual analysis of heterogeneous data“. In: *IEEE Transactions on Visualization and Computer Graphics* 18.6 (2012), pp. 998–1010. DOI: 10.1109/TVCG.2011.108.
- [SVV99] C. Speier, J. S. Valacich, and I. Vessey. „The influence of task interruption on individual decision making: An information overload perspective“. In: *Decision Sciences* 30.2 (1999), pp. 337–360.
- [TC05] J. J. Thomas and K. A. Cook. *Illuminating the path: The research and development agenda for visual analytics*. Tech. rep. Pacific Northwest National Lab.(PNNL), Richland, WA (United States), 2005.
- [Tea14] R. C. Team. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. 2013. 2014.
- [Ter95] L. G. Terveen. „Overview of human-computer collaboration“. In: *Knowledge Based Systems* 8.2 (1995), pp. 67–81.
- [TO17] K. Tsubouchi and T. Okoshi. „People’s Interruptibility in-the-Wild: Analysis of Breakpoint Detection Model in a Large-Scale Study“. In: *Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers*. UbiComp ’17. Association for Computing Machinery, 2017, pp. 922–927. ISBN: 9781450351904. DOI: 10.1145/3123024.3124556. URL: <https://doi.org/10.1145/3123024.3124556>.
- [Tol+12] A. Toledo et al. „Personalized recommendation in interactive visual analysis of stacked graphs“. In: *ISRN Artificial Intelligence* (2012).
- [Tre85] A. Treisman. „Preattentive processing in vision“. In: *Computer vision, graphics, and image processing* 31.2 (1985), pp. 156–177.
- [Tri] Trifacta. trifacta.com. Accessed: July 31st, 2020.
- [Tuk77] J. W. Tukey. *Exploratory data analysis*. Vol. 2. Reading, Mass., 1977.
- [Van06] J. J. Van Wijk. „Views on visualization“. In: *IEEE Transactions on Visualization and Computer Graphics* 12.4 (2006), pp. 421–432.

- [Wal+18] E. Wall et al. „A heuristic approach to value-driven evaluation of visualizations“. In: *IEEE Transactions on Visualization and Computer Graphics* 25.1 (2018), pp. 491–500.
- [Wan+16] Y. Wang et al. „AmbiguityVis: Visualization of Ambiguity in Graph Layouts“. In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), pp. 359–368. DOI: 10.1109/TVCG.2015.2467691.
- [WHA07] W. Willett, J. Heer, and M. Agrawala. „Scented widgets: Improving navigation cues with embedded visualizations“. In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (2007), pp. 1129–1136. DOI: 10.1109/TVCG.2007.70589.
- [Wij06] J. van Wijk. „Views on Visualization“. In: *IEEE Transactions on Visualization and Computer Graphics* 12.4 (2006), pp. 421–433. DOI: 10.1109/TVCG.2006.80.
- [WM13] E. Wu and S. Madden. „Scorpion: Explaining away outliers in aggregate queries“. In: *Proc. of the VLDB Endowment* 6.8 (2013), pp. 553–564.
- [Won+16] K. Wongsuphasawat et al. „Voyager: Exploratory Analysis via Faceted Browsing of Visualization Recommendations“. In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), pp. 649–658. DOI: 10.1109/TVCG.2015.2467191.
- [WW89] J. Wu and S. Wei. *Time series analysis*. Hunan Science and Technology Press, ChangSha, 1989.
- [Yan+07] D. Yang et al. „Managing discoveries in the visual analytics process“. In: *Proc. of the Workshop on Visual Analytics and Knowledge Discovery: Integrating Automated Analysis with Interactive Exploration* 9.2 (2007), pp. 22–29.
- [Yi+07] J. S. Yi et al. „Toward a Deeper Understanding of the Role of Interaction in Information Visualization“. In: *IEEE Transactions on Visualization and Computer Graphics* 13 (2007), pp. 1224–1231.
- [ZAM11] Z. Zheng, N. Ahmed, and K. Mueller. „iView: A feature clustering framework for suggesting informative views in volume visualization“. In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 1959–1968.
- [Züg+18] M. Züger et al. „Sensing interruptibility in the office: A field study on the use of biometric and computer interaction sensors“. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 2018, pp. 1–14.



Die approbierte gedruckte Originalversion dieser Dissertation ist an der TU Wien Bibliothek verfügbar.
The approved original version of this doctoral thesis is available in print at TU Wien Bibliothek.

Part II

Papers



Die approbierte gedruckte Originalversion dieser Dissertation ist an der TU Wien Bibliothek verfügbar.
The approved original version of this doctoral thesis is available in print at TU Wien Bibliothek.

Characterizing Guidance in Visual Analytics

This paper first appeared in: IEEE Transactions on Visualization and Computer Graphics 23(1); 111-120, 2017. ©2017 IEEE. Reused with Permission. (See [Cen+17])

Full List of Authors Besides Davide Ceneda, the following authors contributed and provided valuable input to this paper: Theresia Gschwandtner, Thorsten May, Silvia Miksch, Hans-Jörg Schulz, Marc Streit, and Christian Tominski.

2.1 Abstract

Visual analytics (VA) is typically applied in scenarios where complex data has to be analyzed. Unfortunately, there is a natural correlation between the complexity of the data and the complexity of the tools to study them. An adverse effect of complicated tools is that analytical goals are more difficult to reach. Therefore, it makes sense to consider methods that guide or assist users in the visual analysis process. Several such methods already exist in the literature, yet we are lacking a general model that facilitates in-depth reasoning about guidance. We establish such a model by extending van Wijk's model of visualization with the fundamental components of guidance. Guidance is defined as a process that gradually narrows the gap that hinders effective continuation of the data analysis. We describe diverse inputs based on which guidance can be generated and discuss different degrees of guidance and means to incorporate guidance into VA tools. We use existing guidance approaches from the literature to illustrate the various aspects of our model. As a conclusion, we identify research challenges and suggest directions for future studies. With our work we take a necessary step to pave the way to a systematic development of guidance techniques that effectively support users in the context of VA.

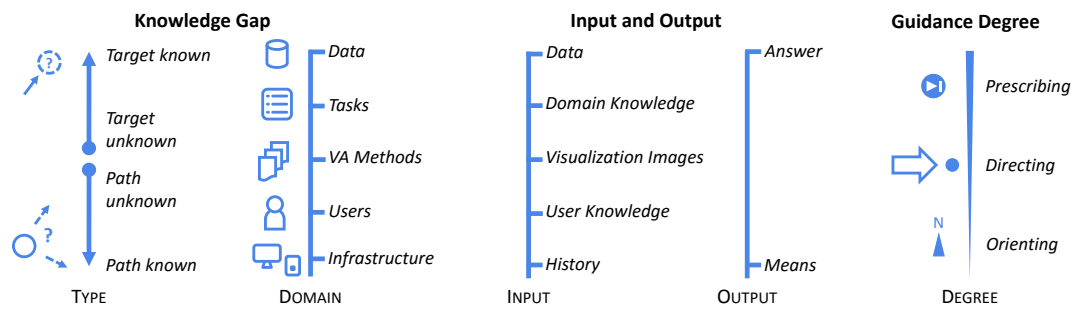


Figure 2.1: Guidance can be characterized in terms of the main aspects: knowledge gap, input and output, as well as guidance degree.

2.2 Introduction

Thomas and Cook [TC05] define visual analytics (VA) as a technology that supports discovery by combining automated analysis with interactive visual means. The key idea is to establish a synergy of computational power and human reasoning. In recent years, a large number of VA approaches have been developed for diverse data, analytical problems, and user requirements. They are particularly useful in situations where complex problems have to be solved. Consequently, these methods are often not as simple to use as one would wish they were. Analytical computations usually require the user to set parameters, while suitable values are not clear upfront. Visual representations of complex phenomena tend to be more demanding to interpret than plain information graphics. And also in terms of interaction there are many more things to control, in order to make proper progress in the data analysis process.

The problem is that users, which are typically experts in their domain, but novices when it comes to VA, could be easily overwhelmed. Which method to use, how to set parameters, or how to get from one part of the data to another? Particularly when visual analysis methods are not applied on a regular basis, but only occasionally, such questions are not easily answered, a fact that hinders the effective use of VA in practice. What is needed are solutions that guide the user during data analysis and exploration. We see appropriate guidance as a key factor for significant improvements of the overall quality of data-intensive analytical work. In this context, the study and the development of tools for and models of guidance in VA is an important research topic.

While there are already a few approaches that offer guidance to users, there is only limited knowledge about the general mechanisms and underlying structures of guidance. Therefore, the goal of this paper is to contribute to a conceptual characterization of guidance. In section 2.3, we study the design space of guidance and develop a general model of guidance in the context of VA. We build upon the initial characterization of guidance by Schulz et al. [Sch+13] and revise it with respect to the knowledge gap of users, the input and the output of a guidance generation process, as well as the degree to which guidance is provided (see Figure 2.1). Van Wijk's [Wij06] model of visualization

serves as the basis for the development of a first model of guided VA. Our new model includes the fundamental building blocks of guidance and attaches them properly to the classic components of VA.

section 2.4 bridges the gap between our conceptual considerations and guidance in practice. The individual dimensions and categories of the design space will be used to structure a review of existing approaches, which offer guidance in diverse ways. Selected examples from our own previous work will be described in more detail. In section 2.5, we focus on open research questions related to guidance. With this we hope to stimulate the development of effective guidance approaches and systems in the future.

In summary, the key research contributions of this work are (1) a characterization of guidance in VA, (2) a conceptual model of guided VA, (3) a review of guidance approaches, and (4) a compilation of open research challenges.

2.3 Guidance: Terminology and General Concepts

In this section, we characterize the main aspects of *guidance*. In order to make this concept clear, we will first take a look at an illustrating example that deliberately leaves out any VA-specific aspects.

2.3.1 An Illustrating Non-VA Example

We imagine a smart car, supporting its driver in the journey to a destination. If the driver is confident about how to get there, he or she will drive the car, while the car provides guidance by showing the names of the traversed streets, highlighting the position of stops or traffic lights, and streaming the weather conditions for the current day. If the driver does not know how to reach the destination, the car could provide a higher degree of guidance by displaying turn-by-turn navigation instructions. These could also include alternative paths fulfilling certain constraints (e.g., avoid traffic jams or refuel required). Finally, in an advanced scenario, it is the car that drives autonomously to the destination, taking on each decision, changing paths if needed, but leaving the driver the freedom of taking over the steering wheel to deviate from the route or act in unexpected situations.

With this car example we sketch three different scenarios in which a system offers support to a human operator. By exploiting information derived from different sources and sensors, the system provides the driver with different degrees of assistance in order to address different needs: driving autonomously, searching for routes, and displaying additional information.

The example already hints at some of the important questions related to guidance. What are the needs of the human? How much guidance is provided by the system, and how is it conveyed to the driver? Based on what information is the guidance generated? In the next paragraphs, we will look at these questions in detail and through the lens of VA.

2.3.2 Definition of Guidance

Guidance is a broad term with much room for interpretation. To arrive at a crisp definition of guidance in VA, it makes sense to first review how the term is used in general and in related areas. Naturally, definitions provided in dictionaries are generic. According to two dictionaries, guidance can be defined as “advice or information aimed at resolving a problem or difficulty” [Oxf] or “the act or process of guiding someone or something” [Mer]. These definitions are quite interesting, because they highlight guidance as a *process* aiming at solving a *problem*.

Another perspective of guidance is given in the field of human-computer interaction. Engels [Eng96] outlines the main dimensions of guidance: the ‘What’, clarifying the problem, composed by an initial state and a goal state, and the ‘How’, aimed at solving the discrepancies between the two states by decomposing the main problem in a sequence of sub-problems that are easier to solve. Instead of focusing on the process itself, Smith and Mosier [SM86] emphasize the importance of interactivity and the visual nature of guidance defined as a “pervasive and integral part of interface design that contributes significantly to effective system operation”. They also include guidance in their guidelines on visual interface design. The importance of guidance is also underlined by Dix et al. [Dix+04]. Since each analysis system might be used by different kinds of users, it is inevitable that not everyone will understand it. This is where guidance is essential, in the sense of *knowing where you are or what will happen*. Guidance has to be unobtrusive to the user, and adaptive to the particular context, as the type of assistance a user requires varies and depends on many factors.

In the visualization literature, one can find several notions that are similar or related to guidance, including *recommendations*, *incentives*, or *assistance*. Schulz et al. [Sch+13] group these different notions under the common term *guidance*. In their thinking, guidance refers to methods that have the goal of providing dynamic support to users, for example, when exploring data or when finding the best visual mapping for presenting analysis results. In addition to that, they also consider guidance in terms of suggesting a suitable domain expert and an appropriate computational infrastructure to carry out particular tasks.

From the diverse interpretations of guidance in various fields, we derive a definition of guidance in the context of VA:

Guidance is a computer-assisted *process* that aims to actively resolve a *knowledge gap* encountered by users during an *interactive* visual analytics session.

According to this definition, guidance is a dynamic process that aims to support users in a particular task. In general, any task can be decomposed into a series of actions or decisions that lead to a desired result. Guidance provides support for at least one of these actions in situations where a user is unable to identify, judge, or execute the action.

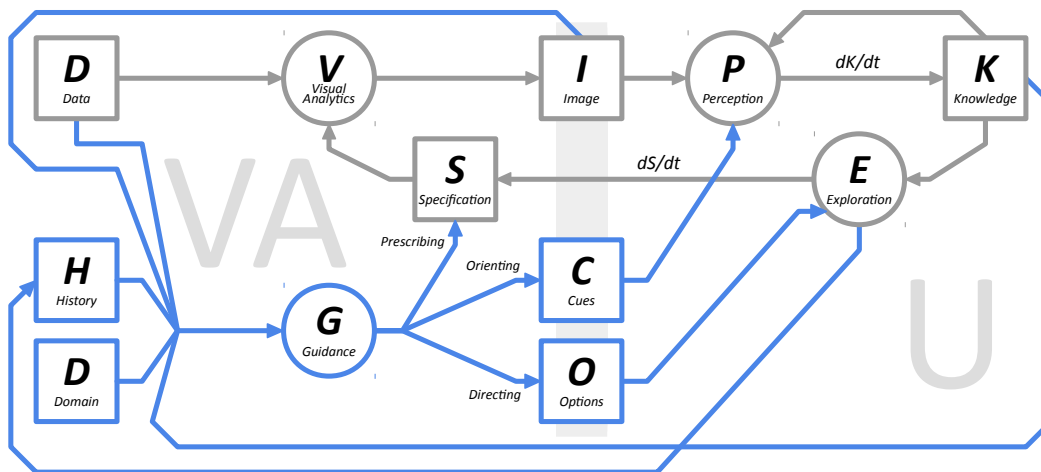


Figure 2.2: Components of guidance (in blue) attached to van Wijk's [Wij06] model (in gray). Aspects of guided VA are shown to the left, while user aspects (U) are on the right. Guidance considers the user's knowledge (or lack thereof) and may build upon various inputs, including data, interaction history, domain conventions, and visualization images. Different degrees of guidance are possible. Orienting uses visual cues to enhance perception. Directing supports exploration by providing alternative options. Prescribing directly operates on the specification. Guidance positively affects the user's knowledge in a dynamic process that eventually converges to zero knowledge gap.

Our definition also includes cases where the desired result is not known in advance, and thus, the actual task must be derived from previous actions. Yet, we do not consider guidance to take over the reasoning part. For example, guidance is not supposed to retrospectively *explain* what is shown in visual data representations and how or why it came about. Instead, guidance provides prospective assistance so that users can make sense of the data on their own.

It is important to note that our definition focuses on the human perspective of guidance in that the system is guiding the human user [Hor99]. There is also the notion of human users guiding algorithms to improve analysis results, but this is not what we are addressing here. This will become clearer in the next paragraphs, where we sketch a model of guided VA.

As a starting point for a first model of guidance in the context of VA, we use van Wijk's [Wij06] model of visualization. We make a slight modification though in that we replace the term *visualization* by *visual analytics*. This makes clear that our model covers both visual and analytical methods. The model is shown in gray in Figure 2.2. Boxes represent artifacts, such as data or images, while circles represent functions that process some input and generate some output. Visual and analytical means (V) transform data [D] into images [I] based on some specifications [S]. The images are then perceived (P) to generate some knowledge [K]. Based on their accumulated knowledge, users can

interactively explore (E) the data by adjusting the specifications (e.g., choose a different clustering algorithm or change the perspective on the data). As such, van Wijk's model effectively conveys the iterative and dynamic nature of knowledge generation mediated through VA. This makes it perfectly suited to be expanded to a model of guided VA.

We attach new guidance-related components to the model, shown in blue in Figure 2.2. A central position is taken by the guidance generation process (G). It is hooked up first and foremost with the user's knowledge [K]. The reason is that before we can take any measures of guidance, we need to know what the particular problem of the user is. Similar to the worldview gap [AS05], we coin the term *knowledge gap* to capture the actual deficit that hinders continuation of the data analysis. The guidance generation process (G) is further connected to sources of information based on which guidance can be generated. These sources include the original data [D], visualization images [I], interaction history or provenance [H], and domain conventions or models [D]. Taken together, these components represent the *input* to the guidance generation process.

On the *output* side, results of a guidance generation process can be delivered in various ways. Figure 2.2 illustrates three different scenarios. Orienting provides basic guidance through visual cues [C]. Directing offers useful options or alternatives [O] that the user may or may not choose to follow. Prescribing directly operates on the specification [S] in order to automatically generate suitable visual results.

The main goal of guidance is to create and maintain an environment in which users are able to make progress and perform their tasks effectively. This dynamic progressive procedure is well expressed by the knowledge change (dK/dt) occurring as a consequence of the guided visual analysis and the interactive adjustment (dS/dt) of the specification. A critical concern is that knowledge is acquired through perception and cognition (P). So the leverage point of guidance is to facilitate perception and cognition at different *degrees*, for example, by showing visual cues alongside the visualization, by offering options that, if chosen, lead to an improved visualization, or by taking over control and circumventing progress-hindering obstacles automatically.

In summary, we can identify three main characteristics of guidance: (1) the reasons why guidance is needed, i.e., the knowledge gap, (2) the inputs that are used to provide guidance as well as the output, and how the output is conveyed to the user, and (3) the expressed guidance degree. In the following, we will describe these aspects in detail.

Knowledge Gap

The knowledge gap pertains to the question: *What does the user need to know to make progress?* There are many different pieces of information that the user may need to know before progress can be made. It could be that a suitable color map has to be chosen before a certain data characteristic becomes visible. Or it may be necessary to visit different parts of the data before high-level relations can be discerned.

While a knowledge gap can come in myriad ways, there are two distinct *types* of knowledge gaps:

Target unknown means the user does not know the desired result. For example, the analyst has no idea about the clustering outcome to be generated.

Path unknown means that the user does not know how to reach the desired result. For example, given some ground truth, the analyst does not know which algorithm to choose and how to parametrize it, in order to extract the ground truth.

Figure 2.1 illustrates the axis of known and unknown target and path that characterize the knowledge gap. Another perspective on the knowledge gap is the *domain* to which it pertains. There are five domains that are particularly relevant in VA:

Data: The user needs guidance in terms of data subsets or features. Guidance could (semi-)automatically identify such subsets or features based on some kind of “interestingness” definition, such as degree-of-interest functions or recommender systems.

Tasks: The user needs help in structuring a goal into a series of tasks that solve the goal. This is a high-level gap that guidance could narrow by hinting at what to do next. It is independent of the actual choice of VA methods to be used.

VA Methods: The user needs help with the available visual, analytical, and interactive methods. Guidance in this space could suggest suitable visualization techniques or algorithm parametrizations. This also relates to enhancements by means of providing additional information about VA methods.

Users: It is unclear who should carry out a task. When analysts work collaboratively, guidance could provide advice as to who would be a suitable expert to work on a specific task. This avoids situations where users are assigned to tasks that do not match their expertise.

Infrastructure: The user is unsure which infrastructure to employ. Guidance in this case means recommending hardware (e.g., display wall or touch-enable surface) and software (e.g., analytical mining tools or interactive exploratory tools).

Users may or may not be aware of the gap. It can very well be that a user does not even know that a certain procedure has to be performed before useful analytical results can be generated. This makes capturing the knowledge gap difficult. If users are aware of it, they can actively make it known to the system. If not, the system has to infer the knowledge gap, for example, by detecting deviations from domain conventions or long dwell times during exploration.

Input and Output

The input is concerned with the question: *What is the basis for generating the guidance?* When we look at the output of the guidance generation process, we are facing two questions: *What is the answer to the user’s problem and how is the answer presented?*

The inputs are the foundations upon which guidance is generated. In the context of VA systems, we identified the following useful sources of information.

Data includes all kinds of information readily available or derivable from the data to be analyzed. Concrete examples are raw data, statistical properties of the data, data topology, or meta-data.

Domain Knowledge refers to information that originates from the application domain. This could be expert systems, domain models, workflows, or conventions.

Visualization Images include the visual data representations and information about mapping parameters. They can be useful for understanding what the user is actually seeing.

User Knowledge is about information that users input to the system, including annotations or degree of interest (DOI) functions, or information that the system can infer from the user.

History relates to keeping track of interactive changes. This includes logging interaction steps, employed algorithms, applied parameterizations, or visited parts of the data.

Concerning the output of the guidance generation process, there are two aspects to be considered: finding of a suitable *answer* and using appropriate *means* to convey the answer to the user.

Answer: Conceptually, finding the answer boils down to developing a function that takes the knowledge gap plus additional input and computes a suitable result.

$$\textit{guidance}(\textit{gap}, \textit{input}) \rightarrow \textit{answer}$$

This definition is abstract and broad enough to consider many different situations. Iterations of the function converge to the goal of zero knowledge gap, where each iteration conveys a variable amount of knowledge to the user, depending on the user's expertise and perceptual and cognitive abilities. In this sense, guidance is an active process and the user is included in the loop.

We distinguish *direct* from *indirect* answers. Usually, the knowledge gap should be answered directly. For example, if a user has a problem in finding a suitable value for a clustering parameter, the guidance generation process should provide promising candidates. On the other hand, guidance could provide indirect answers. Staying with the same example, the guidance could hint at interesting structures in the data, whose analysis (note the indirection) may help the user fine-tune the clustering parameter.

Means: Once computed, the answer has to be communicated to the user. This is a critical step. The goal is to induce an impulse in the user so as to enhance perception or to trigger exploratory actions. It is typical in VA settings that the answer is presented visually. This could mean adjusting the visualization mapping, providing visual enhancements, or including additional user interface elements. Yet, we do not consider the means to be limited exclusively to the visual channel. Depending on the context in which guidance is used, answers can be provided by exploiting non-visual channels as well, including sounds or tactile feedback.

Guidance Degree

The guidance degree is about the question: *How much guidance is provided?* For the car example mentioned earlier, we already saw that guidance can be provided at different levels. The same holds true for guidance in VA. The *guidance degree* specifies the extent to which guidance is required and actually provided. The guidance degree is not static, but varies over time as tasks, data, and procedures change through the course of a VA session. This enables guidance to be fine-tuned to the requirements at hand. For example, if a user gets lost during data exploration, the guidance degree should be increased. If the user feels too restricted by the system-prescribed course, the guidance degree should be decreased.

The two extremes of the guidance degree are *no guidance* (no support given to the user) and *fully automated* (no options for the user to intervene). These are, however, only of theoretical relevance. In practice, the guidance degree is in between these extremes, with three characteristic scenarios being particularly interesting to look at:

Orienting: Providing merely orientation is at the low end of the guidance degree. The main goal is to build or maintain the user's mental map. Orienting in VA typically involves adopting the map metaphor for an abstract domain. Such a map may contain potential targets and paths as well as relations among them. Providing visual cues hinting at these targets and paths are a common strategy for implementing orientation. Visual overview technique may provide some kind of orientation as well.

Directing: Directing represents a medium degree of guidance. In contrast to orienting, directing approaches emphasize a certain preference for a future course of action. The system presents the user with a set of alternative options to produce the desired result or a set of similar results. The suggestions may differ in terms of quality and costs for different paths leading to the same result or, in terms of interest for paths, leading to similar or new results. Directing can benefit from preview techniques that help users make informed decisions for one or the other option.

Prescribing: With prescribing we reach a higher degree of guidance. In contrast to directing, prescribing approaches make decisions on steps to be taken on their own. Prescribing implements a largely automated process, which proceeds towards a

specified target. Such a process may cover any (sub-)task of analysis regardless of its scope. In the context of VA, it is important to visually present the intermediate steps of the process and the decisions that lead from one step to the next. In a sense, this degree of guidance can be compared to an interactive presentation. A user may interrupt the presentation and ask for details, or rewind/reverse it to revisit a nugget of knowledge that has been found earlier. Depending on the degree of automation, the user can recover control for a while and nudge the presentation to another path or even another target.

With these three scenarios we have completed sketching the key characteristics behind guidance. In the next section, we will use the developed characterization to structure a broader review of existing guidance approaches in the context of VA.

2.4 A Review of Guidance in Visual Analytics

There is no single comprehensive guidance approach for VA that covers all aspects that we discussed in the previous section. Yet, instantiations of specific aspects can be found in existing work. In this section, we apply our characterization to a selection of examples to showcase the state of the art and to show possible connections between complementing approaches.

2.4.1 Knowledge Gap

Type

The following examples illustrate the difference between guidance approaches allowing the user to find and specify solutions, and guidance approaches that allow the user to pursue the path towards a solution.

Target Unknown The *target* refers to a solution to a specific problem, such as a useful visualization. Usually such a solution is not purely deterministic, but instead is defined in guided interaction with the user. For instance, Fujishiro et al. developed Gadget [Fuj+97], that exploits a knowledge-base containing task descriptors with the aim of suggesting a set of possible visualizations effective for the given tasks. A further approach is BOZ [Cas91]. In this approach, tasks are modeled as logic rules that are associated with suitable visual encodings, that are proposed to the user to improve the analysis. The aim is to improve the user's performance. Both approaches provide support in choosing the correct target, in these cases a visualization. The users of automated techniques face similar problems. Choosing appropriate techniques for an analytical task or selecting their parameters are cases of unknown targets. As one of many examples, Krause et al. [KPB14] developed a tool to rank data features for modeling, offering guidance in the feature selection process. In this case the target is the set of most useful features.

Path Unknown The next two approaches address the problem of finding sequences of actions to achieve a goal, be it the creation of a view or the application of filters to a dataset. Willet et al. [WHA07] developed scented widgets, a technique that offers guidance in the data domain, to help users in completing a series of data transformation steps. These widgets are interactive elements in a graphical user interface that incorporate information about other users' activity. The hints provided by scented widgets level possible knowledge gaps and lead inexperienced users to significant results. The visual pre-processing by Bernard et al. [Ber+12] offers guidance in composing a sequence of steps for time series transformation. The effects of each step are demonstrated by input-output comparison of time series samples suggested by the system.

Domain

The guidance domain captures the subject matters with respect to which a knowledge gap can manifest. Most of the existing literature is concerned with guiding towards data of interest and suitable VA methods. Yet, the following approaches will illustrate how versatile the guidance domain can be beyond data and VA methods.

Data Finding data that are worthwhile to investigate in a large dataset is a known challenge in VA research. One of the most prominent ways of assisting this task is by capturing what makes a data item interesting to the user in a so-called *degree-of-interest function* and recommending those data items with high interest values to the user [GST13]. Aspects that factor into such a quantitative notion of interestingness are, for example, special data characteristics (e.g., uniqueness, extreme properties), novelty (e.g., whether a data item has been looked at before), or visual saliency (e.g., whether a data point is visible or overplotted). To infer automatically what parts of the data might interest the user is subject of the area of *user profiling* and in particular *preference elicitation* [HD12].

Tasks Given some data of interest, it is not necessarily obvious what to do with it. Step-by-step methodologies or analytical workflows that have been found to be generally good approaches in a certain domain can help in such cases to suggest promising analytical tasks and how to complete them.

VA Methods Offering guidance on VA methods means to point the analyst to concrete methods to solve a specific task. For instance, *VizAssist* [BGV15] provides such guidance by matching the data to be analyzed to suitable visual methods.

Users This kind of guidance aims to help the research of appropriate users to assign to a given task. This actually relates closely to the field of *expert finding* [MDH01], for which already some visual and analytical tools exist [May07]. In the field of VA, these methods are not yet picked up on.

Infrastructure Guidance can be provided to inform possible users about the computational infrastructure needed to perform a certain analysis. Radloff et al. [RFS12] present a framework for smart view management, that takes views, available display spaces, and analytical tasks into account to suggest favorable mappings onto available displays. In essence, it computes for each possible view-display mapping a *view quality score* that is weighted by the importance of the view for the task at hand. Thus, the framework suggests view configurations that maximize the sum of these weighted scores.

2.4.2 Input and Output

Input

Inputs are the sources of information that are used to generate guidance. Most approaches require a combination of sources to offer a useful solution. Our examples are categorized according to their primary source.

Data Gratzl et al. created Domino [Gra+14], a general technique for tabular data that permits the user to create, explore and extract heterogeneous data subsets and show their relationships by visually connecting them. Visual cues indicate compatible views, with respect to data properties. Lex et al. designed StratomeX [Lex+12], which is meant to support the analysis of cancer data. The tool shows the analyst the relationships among cancer subtypes data. Either in Domino and StratomeX, visual cues like lines or ribbons are used to make clearly visible to the user the relationships between the data.

Domain Knowledge Guidance can also be generated based on domain related knowledge: task knowledge, workflows, and conventions. The work by Streit et al. [Str+12] presents a step-by-step process for the analysis of heterogeneous data. The process aims to satisfy both experienced and inexperienced users improving orientation and analysis completeness by using tasks knowledge and providing the user a clear sequence of steps to reach a result. In general, there are many approaches that use domain knowledge to generate guidance. Some of those we have already discussed in previous sections of this paper [PS08; Cas91; Fuj+97].

Visualization Images This category focuses on guidance systems that exploit information derived from views, mappings, and visual elements. One example of taking visual features as input to generate user guidance is given by Wang et al. [Wan+16]. They devised a guidance approach in the field of graph drawing. It provides guidance by calculating an index about the ambiguity of the graph drawing (e.g., edge crossings or insufficient distances among nodes) and highlighting problematic graph regions. This approach considers the visualization at hand to guide the user on which areas to investigate further in order to uncover cluttered parts of the represented network.

User Knowledge User feedback, be it explicit (the user evaluates his/her experience directly) or implicit (the information is deduced from the user's actions and performances),

is also a valuable input for generating user guidance. While implicitly derived feedback avoids cumbersome feedback collection and does not interfere with the user's workflow, it may be subject to errors caused by misinterpreting the user's activities [OK98]. Mouse events, like clicks or hovering over specific regions of the display, could represent a source of implicit information about the user's preferences and interests. This feedback could be used, for instance, to steer a document retrieval operation or the search for a specific product in an e-commerce website [Hij04]. Gotz and Wen [GW09] present a comprehensive example of user and task based guidance. The interaction log of the user is matched with a set of interaction patterns derived from previous user behaviors. These patterns are used to identify the implicit task, which in turn is used to adapt the visualization.

History Another possible input for generating guidance is information derived from the user's past actions. Kreuzeler et al. [KNS04] and Derthick and Roth [DR00] present two similar solutions to assist the analysis by exploiting the user's history of actions, which are shown to the user to provide details about what can still be possibly analysed. Shrinivasan et al. [SW08] present a tool composed of three views, of which one is intended to show the analytical process history, one represents the findings, and the last one shows the dataset. These views enable the user to build a context that can help justify or prove a result or finding.

Output

The output of the guidance generation process is composed by the answer to the user's knowledge gap and by its (visual) representation. It may happen that the output of the guidance generation process does not fully solve the user's knowledge gap. However multiple output sequences may in the end converge towards zero knowledge gap.

Answer Although the answer is a response to a problematic situation, which is a direct consequence of the knowledge gap, there exist situations in which it can be only addressed indirectly.

Direct: The answer is given on the same domain as the knowledge gap. An example of direct answer can be found in the approach by Perer et al. [PS08]. The knowledge gap is related to the question: *Which are the steps to reach the result?* The answer to this problem is directly addressed in that the system provides the user with a list of steps to complete the task. Another example is the approach by May et al. [May+12]. The authors describe a process to support the user in finding and reaching interesting graph regions. In other words, the problem is locating these regions, and the system shows the user how they can be directly reached.

Indirect: Domino and Stratomex are approaches that address the knowledge gap in an indirect way [Gra+14; Lex+12]. In these approaches, the knowledge gap can be synthesized with the following sentence: *The user does not know what is the best way to visualize and compare data subsets*, and *The user would like to consider different data*

sources. In these examples, the two systems do not guide the user directly to results, but instead take care of the visualization of subsets or show relationships among them. In other approaches that fall into this category [DR00; SW08; KNS04] the knowledge gap relates generically to gaining insights. However, the user is just supported indirectly by showing them the history of actions.

Means Once an answer is computed, it has to be communicated to the user. This typically happens through a visual medium. In Stack'n'Flip [Str+12], the authors propose a visualization in which the sequence of steps needed to perform a task (i.e., the answer to a user need) is visually shown and added to the view: the path to follow is added below the main view together with the needed datasets. Alternative paths are of different color, while possibly related paths are highlighted. Jankun-Kelly and Ma [JM00] present an approach to guide the selection of parameter combinations in huge parameter spaces. The key idea is to present the user with a stack of two-dimensional spreadsheets showing all possible combinations of dimensions. Dedicated interaction techniques support the navigation in the parameter space. The user can then easily explore suitable parameter combinations for the problem at hand. Similarly, Lehmann et al. [Leh+15] describe how properties of data distributions in multidimensional visualizations can be visually encoded to ease their analysis and interpretation. Some examples of interaction facilities include Kreuzeler et al.'s [KNS04] or Derthick et al.'s [DR00] history mechanisms which support undoing and redoing of actions. Scented widgets [WHA07] are interactive elements of a user interface enhanced with visual guidance.

2.4.3 Guidance Degree

Another important aspect of guidance methods is the degree of assistance provided, which should meet the user's needs. We consider the degree as a continuous spectrum that spreads from orienting to prescribing guidance.

Orienting

Support for orientation is closely related to the goal of building and preserving a user's *mental map*. A mental map is the internal representation of the analysis the user develops as the analysis proceeds. The relevance of supporting the correct development of the mental map has been recognized in various studies, often in the field of graph drawing [PHG07; AP13].

A mental map for VA typically spatializes abstract relations. We present two groups of examples that operate in two different domains of the knowledge gap. Approaches in the first group primarily offer orientation in the *data* domain. These approaches aim at mapping relations between data subsets, patterns, attributes or models. Gratzl et al. [Gra+14] and Lex et al. [Lex+12] help users understand these relations by showing the connections between different parts of the data. Some of the relations may be known beforehand, others may be introduced during analysis. With a similar goal, Yang et

al's. [Yan+07] approach offers orientation in the 'pattern space'. It generates a map of patterns found during an entire session. The patterns are arranged according to their similarity, regardless of how and when the patterns actually have been defined.

Approaches in the second group primarily offer orientation in the *task* domain. These approaches aim at spatializing the series of tasks in the analytical process. This may include methods or intermediate results as well. Kreuzeler et al. [KNS04] sustain user's orientation by making explicit the history of actions, thus, providing guidance in trial-and-error systems. Shrinivasan et al. [SW08] subdivide the analysis process by assigning different views to the history of actions, datasets, and findings, with the aim of supporting the exploration. Finally, approaches like the one proposed by Streit et al. [Str+12], provide orientation but as a part of a broader guidance support: in this case data properties, relationships between datasets and predefined domain-specific workflows are exploited to provide assistance.

Directing

Directing approaches offer a ranking or preselection of alternatives, which can be inspected and finally selected by the user. Koop et al. [Koo+08] propose an approach for the creation and completion of visualization pipelines. The knowledge source is a database of previously created visualizations. While the user creates a pipeline, the user is offered suggestions for the most frequent completions. VizAssist [BGV15] and Voyager [Won+16] are recent examples for guiding the choice of visualizations in the context of an analytical process. Both approaches focus on guiding the selection of data and the mapping, rather than on guiding through the visualization design-space. Both use expert knowledge, automatically generated rankings about the data, and user intentions as guidance input. Remarkably, in VizAssist, user intentions are defined explicitly from a catalogue. In Voyager, implicit user intentions are defined incrementally via variable selection.

The guided improvement of visualizations can be complemented by techniques for improving analytical results as generated by different algorithms under different parametrizations. *Directing* approaches in this category display multiple, selectable parameter settings in relation to the quality of results. Bernstein et al. [BPH05] propose an approach for the assessment of classification models and modelers. Infuse by Krause et al. [KPB14] combine the assessment of classifier and feature selection methods. In terms of our characterization, these examples aim at bridging the knowledge gap in the domain of VA methods.

Prescribing

While techniques that provide directions allow users to follow or ignore them, prescribing guidance approaches purposefully limit user influence to traversing a fixed path of analysis. The reasons to do so can be manifold, for example, to reduce the learning curve for casual users by providing them with a simplified analysis experience [AT14], to streamline the analysis process in potentially "distraction-rich" datasets [Als+15], or to have the

analyst stick to an agreed upon standard operating procedure or best practice for better comparability or reproducibility of the results [Str+12].

On a user interface level, this guidance strategy is epitomized by the *wizard interface*. It leads users through a complex task by breaking it into a sequence of smaller tasks that can be carried out step-by-step. Streit et al. [Str+12] show a modern incarnation of such a wizard for visual analysis that departs from the classic modal dialog featuring two buttons to navigate back and forth among the subtasks. Their Stack'n'Flip interface, collects data visualizations that were already explored on one side, those that still need to be explored on the other side, and the one that is currently being explored in the middle of the screen. A linked visualization of the workflow serves as a navigation aid to go back and forth through this stack of visualizations. While still allowing deviations from the workflow, this interface discourages them and shows analysts how to get back on track.


On the view level, the prescriptive guidance strategy is embodied by the concept of providing a “tour” through the data. This idea originated from Asimov’s work on the *grand tour* in high-dimensional data spaces [Asi85]. At its core, it is an animation of different 2-dimensional projections of a multivariate dataset in an attempt to show the data from all possible angles. This idea has since been applied to other types of data, as well. For example, Yu et al. [Yu+10] present a mechanism that automatically constructs such an animated tour from events in time-varying data, whereas Wohlfart and Hauser [WH07] developed an approach that creates a guided and interactive visual story for volume data. While the story is completely defined by the system, the user is left the freedom of asking for details as well as interacting with the story playback. More abstractly, Dennis and Healey [BH02] provide a framework for data spaces in general, called *assisted navigation*. It can be used to generate tours that span certain elements of interest in data space as well as areas of interest in view space.

2.4.4 A Detailed Look at Selected Examples


In the previous paragraphs, we provided an exemplification of each single characteristics of guidance in VA. Next, we will be looking at three approaches in detail. To the best of our knowledge, no approach covers the whole guidance spectrum. Yet, the following examples highlight the most relevant factors when characterizing guidance.

Example 1: Heterogeneity-Based Guidance

Luboschik et al. [Lub+12] facilitate the exploration of multiscale data. The approach points the analyst to scales and regions within the data (*unknown targets*) that exhibit behavior of interest without the need for an exhaustive search. The main idea is to take the most fine-grained *data* as a *guidance input* and to step-wise aggregate it into more coarse-grained data. Pairs of subsequent data scales can then be compared by various metrics, detecting data features that were observable in the more detailed scale, but can no longer be found in the less-detailed aggregated scale. In other words, subsequent scales



(a) A lineplot (top) enriched with multiscale heterogeneity bands (bottom).



(b) A zoomed view of one of the spikes.

Figure 2.3: Orientation by means of visual cues [Lub+12]. (a) The lineplot shows clear spikes among millions of data points. The heterogeneity bands below the plot suggest that there is more to these spikes hidden at higher levels of granularity. (b) Zooming in on one of the spikes in confirms this assumption.

exhibit heterogeneous behavior. This information is then communicated to the user by *means of visual cues*, in this case colored heterogeneity bands that provide *orientation* towards regions that are worthwhile to zoom into. This way, the analyst is given a *direct answer* to the question where deviating behavior from the currently shown will emerge, while at the same time not having to bother with investigating other parts of the data where no such deviation occurs.




Figure 2.4: Orientation via signposts [May+12]. Signposts connect a small, yet detailed focus region of a graph to the invisible “context”. They label outbound edges that connect invisible regions along their shortest path.

Figure 2.3a shows an example of this approach, where a lineplot of millions of data points (top) is enriched with a display of multiscale heterogeneity bands (bottom) that measure how well slope changes are preserved between subsequent scales. The heterogeneity bands show three valleys and within them, very thin, suspicious peaks exactly at those points where the lineplot is at a maximum. Guided by this indicator of more nuanced behavior at these points, the analyst zooms into one of these instances in Figure 2.3b. One can immediately observe that the maximum is far from being as clearcut as the overview in Figure 2.3a suggested. Instead of a distinct tipping point, upward and downward movements are at a constant struggle against each other, until the latter gets the upper hand and reverses the strong upward trend. Without guidance, this interesting behavior of the data at a more detailed scale would have gone unnoticed or only be found by pure chance.

Example 2: Signposts for Navigation in Large Graphs

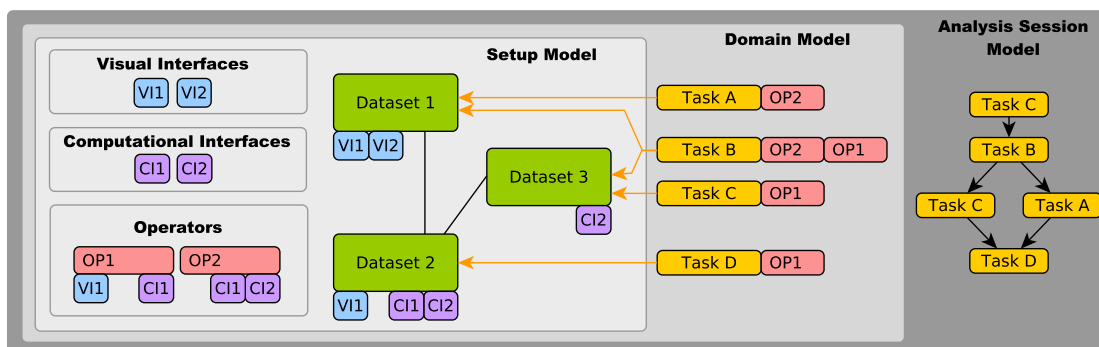
May et al. [May+12] support the orientation in large graphs by using glyphs representing signposts as shown in Figure 2.4. The sign posts are inspired by their real-world counterparts. Only a small subgraph is shown at any time. Orientation is supported by pointing to labeled regions of the graph outside the visible area. The signposts are attached to outbound edges connecting the focal area to the invisible regions along the shortest paths. The signs that are actually shown in the view are selected by the relative importance of regions in terms of distance, region size and overall graph coverage. Moving the visible subgraph triggers a recalculation of the relative importance of regions, and thus the selection of their signs.

In terms of the characterization of guidance, the signposts approach is a technique for orientation with the guidance domain being the graph data themselves. The primary knowledge gap addressed by the approach is literally an *unknown path*. A user will reach any region of interest by following breadcrumbs. Hence, the guidance output is a glyph, which indicates the beginning of the shortest path, and offers an affordance to short-cut movement directly to the target region. To associate a signpost to an intended target, a user requires meaningful names for any region given. The guidance input is based on interaction history and user knowledge. Firstly, the history of visited focal areas is maintained to assess region importance. Secondly, user-defined regions are stored as priority landmarks to ease revisiting.


Example 3: Model-Driven Guidance

In the work by Streit et al. [Str+12], analysts are guided through an analysis session based on a predefined comprehensive model as depicted in Figure 2.5a. The model, which is defined in an authoring process, consists of three stages: (1) a setup model, describing how heterogeneous datasets are connected and which visual and computational interfaces can operate on the datasets; (2) a domain model which defines domain specific tasks and their relation to the setup model; and (3) an analysis session model that defines a workflow as a sequence of tasks.

During an analysis, the setup model serves as a basis to orient the user during an analysis session in the domain of tasks and methods. Hence, the guidance degree is characterized by both orienting and directing. As the workflow is predefined, the path is known, while the target is unknown. The guidance input is covered by the three-stage model (data, visual and analytical interfaces, workflow, and domain specific tasks) together with the history of the analysis and further user input, such as user-defined thresholds. The guidance output is a tree-based meta-visualization that is used for both orienting and directing the analyst, as shown in the lower part of Figure 2.5b.



(a) Domain-specific three-stage model.



(b) Based on the model, stack'n'flip guides users through various analytical views.

Figure 2.5: Model-driven guidance [Str+12]. (a) A domain-specific model is defined in a three-stage process. (b) The model is then utilized to support users during the data analysis.

2.5 Discussion and Future Work

In the previous section, we have seen how existing guidance techniques can assist the user in various ways. The model, as introduced in this paper, is a first step to systematize the emerging field of guidance in VA. In this section, we identify open research questions and

derive suggestions for future work on guidance.

Refining the model Our model explains the embedding of guidance in VA scenarios. It comprises the fundamental components of guidance and their interplay. This helps us understand how guidance works in principle. A sensible next step for the future is to refine the model to develop a better understanding of the internals of guidance. For example, the core function of guidance, i.e., the guidance generation process, largely remains a black box. The illustrating examples implement it in one way or the other. Yet, it remains to be studied if one can extract a general procedure of *how* guidance is actually generated. Such a procedure could then be used as a blueprint for developing new guidance techniques. A sensible refinement to our framework may come by known models. Sacha et al. [Sac+14] expanded the original VA pipeline to highlight the strong synergy between human and machines while generating new knowledge. In the same way it is possible to look at the guidance model to spot where and how it is possible to provide assistance both to the human and to the machine loop.

Similarly, our understanding of the knowledge gap remains limited. Most existing approaches either implicitly infer knowledge gaps a-priori from overplotting and other ambiguities in the visualization (What parts of the data are not visible to the user?) or a-posteriori from interaction histories (What parts of the data the user has not explored yet?). It remains an open challenge to do the same during an ongoing analysis. Simple heuristics, such as long idle time, can be used to automatically detect stalled analysis sessions. Such methods provide but simple indicators of the fact that guidance is needed. For well-balanced and effective guidance, the knowledge gap needs to be specified in greater detail. A promising starting point is to consider established models from human-computer interaction. In Norman's action cycle [Nor13], the execution phase is associated with three layers of competence, *knowing why*, *knowing what*, and *knowing how*. All are needed for making progress in a human-in-the-loop analysis process. Distinguishing these layers will allow us to better attune guidance to the user's personal level of competence. To this end, a fundamental approach to identifying the knowledge gap during the analysis is needed. However, the back and forth between diverging processes (exploration) and converging processes (confirmation), which is typical for VA sessions, makes this a formidable research challenge.

Novel guidance approaches In the literature, there are a number of approaches that deal with guiding in selected aspects of VA. However, we did not find any guidance approach that covers the entire VA process. Here we see potential for future work on novel guidance techniques. New techniques could specifically address the lack of comprehensive guidance for the human-in-the-loop process and offer intertwined guidance on all phases of VA (e.g., how to transform data, modify calculations, and how to read and interact with the resulting visual representations). Just as we see a specialization of VA for specific data classes (e.g., multivariate data, graphs, text), we believe that it also makes sense to consider tailored guidance approaches. An example are guidance techniques for time-oriented data. The dimension of time has a rich structure and it is not always

clear to the analyst which facet of time to focus on (e.g., linear time vs. cyclic time). Navigation in time is another aspect where guidance could assist the user in visiting those parts of the data that potentially lead to interesting findings.

When we look at existing techniques, the majority of them generates guidance based on the data (e.g., [Gra+14; Lex+12]), past analytical actions (e.g., [KNS04; DR00; SW08]), or planned future analytical actions (e.g., [Str+12; PS08; Cas91]), such as workflows, analysis protocols, or standard procedures from the application domain. Only a few techniques (e.g., [Wan+16]) consider the visual representations as input to generate guidance. What other inputs can be useful, emotions [CEK13] for example? Another limitation is that current approaches typically consider only a single type of input. Particularly in the light of the different layers of competence as indicated before, there is a need to consider multiple sources of information. However, it is still an open question how various inputs can be combined in general.

On the output side of guidance, we have a similar situation: Most techniques provide only one degree of guidance: orienting, directing, or prescribing. Novel guidance approaches should support adaptive switching between guidance degrees in order to generate a richer experience. For example, if the user deviates often from the proposed route, orienting may be more suitable than directing or even prescribing. More research is needed to investigate mechanisms for triggering switches between degrees. What would be appropriate indicators (e.g., user input, situation monitoring) and suitable thresholds for automatic switching? Moreover, the guidance interface needs to be designed so as to make switches in the degree transparent to the user.

Regarding the human, existing approaches typically assume a single individual. Yet, VA is increasingly a collaborative effort of several analysts. So far, there are only very few approaches that offer guidance in collaborative scenarios. This is a largely open research question.

Evaluation of guidance Evaluating visualization techniques is notoriously challenging. VA with its mix of analytical, visual, and interactive methods is even harder to evaluate. On top of that, guidance adds considerably to the evaluation challenge. The tight coupling among the involved methods makes it difficult to set up controlled experiments. Already when investigating the visual embedding of guidance (what we refer to as *means*), a number of evaluation questions come to mind. For example, which means are appropriate for what tasks or which means are best suited for which degree of guidance?

Moreover, faster completion time and fewer errors alone might be insufficient to draw conclusions about the usefulness or utility of guidance approaches. An interesting alternative question is if guidance sends the user along worn-out paths or if it is able to suggest side tracks to allow for unexpected discoveries. One way to evaluate this is to simulate the use of guidance. To this end, one can pseudo-randomly select from the suggestions generated by guidance and mark the corresponding spot in the data or parameter space as visited. Useful guidance would lead to the relevant parts of the data or parameter space being gradually filled with marked spots.

Another suggestion to tackle the challenge of evaluation, is to consider self-reporting methods. Ideally, guidance would monitor the situations in which the user resorts to it and keep track of its use. This would allow for deriving conclusions about the utility of guidance depending on the different situations during visual data analysis. Moreover, the collected information can be used not only for evaluating guidance, but they could also serve to implement self-adapting or learning guidance. Certainly, this would require combining guidance with concepts known from artificial intelligence.

Guidance and guidelines With our work, we structure the space of guidance solutions. While guidance is to support the user in using VA tools, we have not considered *guidelines* that apply in the development phase of VA. Particularly with guidance for different data and different tasks, and maybe even for different users employing diverse infrastructures, it can become difficult to develop or choose an appropriate guidance technique for a given problem. Therefore, it is important to provide both guidance for users and guidelines for developers. By guidelines we mean established best practices that a developer can refer to when implementing VA approaches. Such guidelines could, for example, suggest how certain analytical situations are best supported with a certain degree of guidance. We see much potential for future research on guidelines enabling us to make the most of guided VA.

From guidance to mixed initiative visual analytics In this paper, we focused on guidance generated by the computer and provided to the user. Yet this thinking is limited in that it considers only one direction of guidance. Much of the potential of VA lies in the close cooperation of human and computer. To fully exploit this potential, it is necessary to include users assisting the computer in the guidance equation. The benefit of user interaction for complex problem solving has long been known [Weg97]. Yet it remains challenging to integrate human and computer on equal footing to obtain VA solutions that are truly mixed initiative. To tackle this challenges, we first need to better understand the back and forth between computers guiding humans and humans assisting the computer.

2.6 Conclusion

In summary, our work contributes to a better understanding of guidance in VA. We defined guidance as a dynamic, iterative, and forward-oriented process that aims to help users in carrying out analytical work using VA methods. Guidance was further characterized along the knowledge gap of the user, the input and output of the guidance generation process, and the degree of guidance that is actually provided to users. We developed a first conceptual model of guided VA based on van Wijk's model of visualization. A structured review of existing approaches illustrates diverse ways of how guidance can be applied in the context of VA. Finally, we identified open research questions to be addressed by future work on guidance.

2. CHARACTERIZING GUIDANCE IN VISUAL ANALYTICS

In conclusion, we established a basis for the comprehension and the development of assistive approaches that improve the insight generation process and ease the visual exploration and analysis of data.

Bibliography

- [Als+15] S. Alspaugh et al. *Rethinking the First Look at Data by Framing It*. Tech. rep. UCB/EECS-2015-208. University of California at Berkeley, 2015.
- [AP13] D. Archambault and H. C. Purchase. „Mental map preservation helps user orientation in dynamic graphs“. In: *Graph Drawing*. Springer, 2013, pp. 475–486. DOI: 10.1007/978-3-642-36763-2_42.
- [AS05] R. Amar and J. Stasko. „Knowledge Precepts for Design and Evaluation of Information Visualizations“. In: *IEEE Transactions on Visualization and Computer Graphics* 11.4 (2005), pp. 432–442. DOI: 10.1109/TVCG.2005.63.
- [Asi85] D. Asimov. „The grand tour: A tool for viewing multidimensional data“. In: *SIAM Journal on Scientific and Statistical Computing* 6.1 (1985), pp. 128–143. DOI: 10.1137/0906011.
- [AT14] G. Atemezing and R. Troncy. „Towards a linked-data based visualization wizard“. In: *Proc. of Conference on Consuming Linked Data (COLD)*. ACM, 2014, pp. 1–12.
- [Ber+12] J. Bernard et al. „Visual-Interactive Preprocessing of Time Series Data“. In: *Proc. of SIGRAD*. Vol. 81. Linköping University Electronic Press, 2012, pp. 39–48. URL: http://www.ep.liu.se/ecp_article/index.en.aspx?issue=081;article=006.
- [BGV15] F. Bouali, A. Guettala, and G. Venturini. „VizAssist: An Interactive User Assistant for Visual Data Mining“. In: *The Visual Computer* (2015), pp. 1–17. DOI: 10.1007/s00371-015-1132-9.
- [BH02] D. Brent and C. Healey. „Assisted navigation for large information spaces“. In: *Proc. of IEEE Conference on Visualization*. IEEE, 2002, pp. 419–426. DOI: 10.1109/VISUAL.2002.1183803.
- [BPH05] A. Bernstein, F. Provost, and S. Hill. „Toward intelligent assistance for a data mining process: An ontology-based approach for cost-sensitive classification“. In: *IEEE Transactions on Knowledge and Data Engineering* 17.4 (2005), pp. 503–518. DOI: 10.1109/TKDE.2005.67.

- [Cas91] S. Casner. „Task-analytic approach to the automated design of graphic presentations“. In: *ACM Transactions on Graphics* 10.2 (1991), pp. 111–151. DOI: 10.1145/108360.108361.
- [CEK13] D. Cernea, A. Ebert, and A. Kerren. „A Study of Emotion-triggered Adaptation Methods for Interactive Visualization“. In: *Proc. Workshop on User-Adaptive Visualization (WUAV)*. CEUR-WS.org, 2013, pp. 9–16. URL: http://ceur-ws.org/Vol-997/wuav2013_paper_03.pdf.
- [Cen+17] D. Ceneda et al. „Characterizing Guidance in Visual Analytics“. In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (Jan. 2017), pp. 111–120. ISSN: 1077-2626. DOI: 10.1109/TVCG.2016.2598468.
- [Dix+04] A. Dix et al. *Human-Computer Interaction*. 3rd. Pearson Education, 2004.
- [DR00] M. Derthick and S. Roth. „Data exploration across temporal contexts“. In: *Proc. of Conference on Intelligent User Interfaces*. ACM, 2000, pp. 60–67. DOI: 10.1145/325737.325782.
- [Eng96] R. Engels. „Planning Tasks for Knowledge Discovery in Databases Performing Task-Oriented User-Guidance“. In: *Proc. of Conference on Knowledge Discovery and Data Mining (KDD)*. Portland, Oregon: AAAI Press, 1996, pp. 170–175.
- [Fuj+97] I. Fujishiro et al. „GADGET: Goal-oriented application design guidance for modular visualization environments“. In: *Proc. of IEEE Conference on Visualization*. IEEE, 1997, pp. 245–252. DOI: 10.1109/VISUAL.1997.663889.
- [Gra+14] S. Gratzl et al. „Domino: Extracting, comparing, and manipulating subsets across multiple tabular datasets“. In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 2023–2032. DOI: 10.1109/TVCG.2014.2346260.
- [GST13] S. Gladisch, H. Schumann, and C. Tominski. „Navigation Recommendations for Exploring Hierarchical Graphs“. In: *International Symposium on Visual Computing*. Rethymnon, Crete: Springer, 2013, pp. 36–47. DOI: 10.1007/978-3-642-41939-3_4.
- [GW09] D. Gotz and Z. Wen. „Behavior-driven Visualization Recommendation“. In: *Proc. of Conference on Intelligent User Interfaces*. ACM, 2009, pp. 315–324. DOI: 10.1145/1502650.1502695.
- [HD12] C. Healey and B. Dennis. „Interest Driven Navigation in Visualization“. In: *IEEE Transactions on Visualization and Computer Graphics* 18.10 (2012), pp. 1744–1756. DOI: 10.1109/TVCG.2012.23.
- [Hij04] Y. Hijikata. „Implicit user profiling for on demand relevance feedback“. In: *Proc. of Conference on Intelligent User Interfaces*. ACM, 2004, pp. 198–205. DOI: 10.1145/964442.964480.

- [Hor99] E. Horvitz. „Principles of Mixed-Initiative User Interfaces“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1999, pp. 159–166. DOI: 10.1145/302979.303030.
- [JM00] T. Jankun-Kelly and K.-L. Ma. „A spreadsheet interface for visualization exploration“. In: *Proc. of IEEE Conference on Visualization*. IEEE, 2000, pp. 69–76. DOI: 10.1109/VISUAL.2000.885678.
- [KNS04] M. Kreuseler, T. Nocke, and H. Schumann. „A history mechanism for visual data mining“. In: *Proc. of IEEE Symposium on Information Visualization*. IEEE, 2004, pp. 49–56. DOI: 10.1109/INFVIS.2004.2.
- [Koo+08] D. Koop et al. „VisComplete: Automating suggestions for visualization pipelines“. In: *IEEE Transactions on Visualization and Computer Graphics* 14.6 (2008), pp. 1691–1698. DOI: 10.1109/TVCG.2008.174.
- [KPB14] J. Krause, A. Perer, and E. Bertini. „INFUSE: interactive feature selection for predictive modeling of high dimensional data“. In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 1614–1623. DOI: 10.1109/TVCG.2014.2346482.
- [Leh+15] D. J. Lehmann et al. „Visualnostics: Visual Guidance Pictograms for Analyzing Projections of High-dimensional Data“. In: *Computer Graphics Forum* 34.3 (2015), pp. 291–300. DOI: 10.1111/cgf.12641.
- [Lex+12] A. Lex et al. „StratomeX: Visual Analysis of Large-Scale Heterogeneous Genomics Data for Cancer Subtype Characterization“. In: *Computer Graphics Forum* 31.3pt3 (2012), pp. 1175–1184. DOI: 10.1111/j.1467-8659.2012.03110.x.
- [Lub+12] M. Luboschik et al. „Heterogeneity-based Guidance for Exploring Multiscale Data in Systems Biology“. In: *Proc. of IEEE Symposium on Biological Data Visualization*. IEEE, 2012, pp. 33–40. DOI: 10.1109/BioVis.2012.6378590.
- [May+12] T. May et al. „Using Signposts for Navigation in Large Graphs“. In: *Computer Graphics Forum* 31.3pt2 (2012), pp. 985–994. DOI: 10.1111/j.1467-8659.2012.03091.x.
- [May07] M. Maybury. „Discovering Distributed Expertise“. In: *Proc. of AAAI Fall Symposium on Regarding the Intelligence in Distributed Intelligent Systems*. AAAI, 2007, pp. 100–103.
- [MDH01] M. Maybury, R. D’Amore, and D. House. „Expert Finding for Collaborative Virtual Environments“. In: *Communications of the ACM* 44.12 (2001), pp. 55–56. DOI: 10.1145/501317.501343.
- [Mer] Merriam-Webster Online Dictionary. <http://www.merriam-webster.com/dictionary/guidance>. Accessed: July 31st, 2020.
- [Nor13] D. Norman. *The Design of Everyday Things*. Basic Books, 2013.

- [OK98] D. Oard and J. Kim. „Implicit feedback for recommender systems“. In: *Proc. of AAAI workshop on Recommender Systems*. AAAI, 1998, pp. 81–83.
- [Oxf] Oxford Online Dictionary. <http://www.oxforddictionaries.com/definition/english/guidance>. Accessed: July 31st, 2020.
- [PHG07] H. C. Purchase, E. Hoggan, and C. Görg. „How important is the “mental map”? – An empirical investigation of a dynamic graph layout algorithm“. In: *Graph Drawing*. Springer, 2007, pp. 184–195. DOI: 10.1007/978-3-540-70904-6_19.
- [PS08] A. Perer and B. Shneiderman. „Systematic yet flexible discovery: guiding domain experts through exploratory data analysis“. In: *Proc. of Conference on Intelligent User Interfaces*. ACM, 2008, pp. 109–118. DOI: 10.1145/1378773.1378788.
- [RFS12] A. Radloff, G. Fuchs, and H. Schumann. „Supporting Visual Analysis in Smart Meeting Rooms“. In: *Proc. of the International Workshop on Visual Analytics (EuroVA)*. Eurographics, 2012, pp. 1–5. DOI: <http://dx.doi.org/10.2312/PE/EuroVAST/EuroVA12/001-005>.
- [Sac+14] D. Sacha et al. „Knowledge Generation Model for Visual Analytics“. In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (Dec. 2014), pp. 1604–1613. DOI: 10.1109/TVCG.2014.2346481.
- [Sch+13] H.-J. Schulz et al. *Towards a Characterization of Guidance in Visualization*. Poster at IEEE Conference on Information Visualization (InfoVis). Atlanta, USA, 2013.
- [SM86] S. Smith and J. Mosier. *Guidelines for Designing User Interface Software*. Tech. rep. ESD-TR-86-278. Mitre Corporation, Bedford MA, 1986.
- [Str+12] M. Streit et al. „Model-driven design for the visual analysis of heterogeneous data“. In: *IEEE Transactions on Visualization and Computer Graphics* 18.6 (2012), pp. 998–1010. DOI: 10.1109/TVCG.2011.108.
- [SW08] Y. Shrinivasan and J. van Wijk. „Supporting the analytical reasoning process in information visualization“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2008, pp. 1237–1246. DOI: 10.1145/1357054.1357247.
- [TC05] J. Thomas and K. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Computer Society, 2005.
- [Wan+16] Y. Wang et al. „AmbiguityVis: Visualization of Ambiguity in Graph Layouts“. In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), pp. 359–368. DOI: 10.1109/TVCG.2015.2467691.
- [Weg97] P. Wegner. „Why Interaction Is More Powerful Than Algorithms“. In: *Communications of the ACM* 40.5 (1997), pp. 80–91. DOI: 10.1145/253769.253801.

- [WH07] M. Wohlfart and H. Hauser. „Story Telling for Presentation in Volume Visualization“. In: *Proc. of Eurographics Conference on Visualization*. Eurographics, 2007, pp. 91–98. DOI: <http://dx.doi.org/10.2312/VisSym/EuroVis07/091-098>.
- [WHA07] W. Willett, J. Heer, and M. Agrawala. „Scented widgets: Improving navigation cues with embedded visualizations“. In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (2007), pp. 1129–1136. DOI: [10.1109/TVCG.2007.70589](https://doi.org/10.1109/TVCG.2007.70589).
- [Wij06] J. van Wijk. „Views on Visualization“. In: *IEEE Transactions on Visualization and Computer Graphics* 12.4 (2006), pp. 421–433. DOI: [10.1109/TVCG.2006.80](https://doi.org/10.1109/TVCG.2006.80).
- [Won+16] K. Wongsuphasawat et al. „Voyager: Exploratory Analysis via Faceted Browsing of Visualization Recommendations“. In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), pp. 649–658. DOI: [10.1109/TVCG.2015.2467191](https://doi.org/10.1109/TVCG.2015.2467191).
- [Yan+07] D. Yang et al. „Managing Discoveries in the Visual Analytics Process“. In: *SIGKDD Explor. Newsl.* 9.2 (2007), pp. 22–29. ISSN: 1931-0145. DOI: [10.1145/1345448.1345453](https://doi.org/10.1145/1345448.1345453).
- [Yu+10] L. Yu et al. „Automatic Animation for Time-Varying Data Visualization“. In: *Computer Graphics Forum* 29.7 (2010), pp. 2271–2280. DOI: [10.1111/j.1467-8659.2010.01816.x](https://doi.org/10.1111/j.1467-8659.2010.01816.x).



Die approbierte gedruckte Originalversion dieser Dissertation ist an der TU Wien Bibliothek verfügbar.
The approved original version of this doctoral thesis is available in print at TU Wien Bibliothek.

You Get by with a Little Help: The Effects of Variable Guidance Degrees on Performance and Mental State

This paper first appeared in: Visual Informatics Journal 3(4); 177-191, 2019. ©2019 Elsevier. Reused with Permission (See [CGM19b])

Full List of Authors Besides Davide Ceneda, the following authors contributed and provided valuable input to this paper: Theresia Gschwandtner and Silvia Miksch.

3.1 Abstract

Since it can be challenging for users to effectively utilize interactive visualizations, guidance is usually provided to assist users in solving tasks. Guidance is mentioned as an effective mean to overcome stall situations occurring during the analysis. However, the effectiveness of a peculiar guidance solution usually varies for different analysis scenarios. The same guidance may have different effects on users with (1) different levels of expertise. The choice of the appropriate (2) degree of guidance and the type of (3) task under consideration also affect the positive or negative outcome of providing guidance. Considering these three factors, we conducted a user study to investigate the effectiveness of variable degrees of guidance with respect to the user's previous knowledge in different analysis scenarios. Our results shed light on the appropriateness of certain degrees of guidance in relation to different tasks, and the overall influence of guidance on the analysis outcome in terms of user's mental state and analysis performance.

3.2 Introduction

Mixed-initiative visual data analysis [Hor99] is an effective and powerful way to make sense of large data collections and support the completion of complex tasks. In this kind of analysis, the strengths of users and computational systems are joint to reach a common analytical goal. On the one hand, users are enabled to make sense of the data through external cognition. On the other, the computational system offers the means to execute complex calculations, elaborate statistics, or discover patterns [Gib77].

Although visual solutions have been proved to be effective in their scope [BS03; Kei+08; TC05], the research is still far from achieving an effective mixed-initiative integration in which the affordances of the user and the analysis system are balanced [Gib77; BL09; CGM19a]. Therefore, sometimes it can be challenging to effectively use and interact with sophisticated analytical solutions. As a consequence, the analysis may stall.

In the past, many approaches have been developed in the attempt to reduce the burden on users and help them to make sense of the data and the visual interfaces. Ceneda et al. [Cen+17] categorize these methods as *guidance*. Guidance describes the results of enabling an effective human-computer collaboration. In particular, guidance deals with providing a solution to the needs a user develops while performing analysis tasks. These needs are referred to as *knowledge gaps*. Ideally, the guidance process could provide a variety of supporting indications to the user, ranging from hints and recommendations, to step-by-step instructions, to foster a positive outcome of the analysis, solving the aforementioned knowledge gaps and a solution to the stalled analysis.

Although the definition of guidance is quite new, guidance approaches have been around for quite some time [Hor99]. Therefore, it does not surprise the number of approaches showing the benefits of providing guidance during the analysis process. Although the benefits of guidance are clear [CGM19a], what is still not clear is how the effectiveness of the guidance varies according to the user to whom it is provided, and to the task at hand. For instance, different types of guidance may be more effective to support exploratory analysis, while others to verify hypotheses. Furthermore, the effectiveness of guidance may also vary according to the previous knowledge of the user, for instance, if it is a novice user or he/she possesses some knowledge about the analysis domain and the visualization system. Thus, we conducted a user study investigating how different guidance degrees affect users with different levels of previous knowledge to solve different kinds of tasks. We pursued this aim, not only by investigating the repercussions of guidance on task performance, but also on how the provision of guidance may affect the user's mental state. We achieve this by analyzing for instance, how the provision of (or the lack of) guidance may induce frustration, feelings of being lost, improve the user's confidence in results, etc.. We think that this work is useful to designers who intend to create guided visual data analysis systems, fostering an increased awareness of users' needs and the development of mixed-initiative systems. In summary, our main contributions are:

- Investigating the interdependencies among guidance degree, user expertise, task performance, and mental state of the user.
- Describing the impact of different degrees of guidance on task performance and mental state of the user.
- Elaborating the impact of user expertise on task performance, and mental state.

3.3 Related Work

Our work elaborates concepts from two main research topics in literature: guidance in visualization and the dynamics of user's mental state during the analysis.

3.3.1 Guidance in Visualization

Guidance is a research topic that comprises Human-Computer Interaction, Information Visualization and Visual Analytics [Kei+08; Dix+04; SM86]. Guidance has its roots in mixed-initiative data analysis [Hor99] and it contemplates the assistance the user receives from the system, as well as the guidance the user gives to the system to steer the analysis [Cen+18]. Guidance describes what are the benefits deriving from a mixed-initiative analysis and how this collaborative analysis can take place. Formally speaking, guidance is defined as a "*computer-assisted process that aims to actively resolve a knowledge-gap during an interactive*" visual analysis session [Cen+17, p.2]. In simple words, the main goal of guidance is to solve a particular user need, namely a user's *knowledge gap*, which can be seen as the difference between the user's knowledge and the knowledge required to complete a given task. This gap may be related to different aspects of the analysis, like the lack of proper interaction means, or of specific domain-related concepts necessary to interpret the data. The output of the guidance process is an answer to the knowledge gap, that is provided to the user, in some visual form. Different degrees of guidance may be provided in order to meet the user's needs. Ceneda et al. [Cen+17] describe different guidance degrees resulting in different types of guidance. In practical scenarios, the same task can be supported with different guidance: In Figure 3.3, the search for specific data in a time-series can be supported with no guidance, but also with direct recommendations, or by prescribing actions.

The amount of works dealing with guidance is vast: Ceneda et al. [CGM19a] recently reviewed the literature of guidance approaches in visual data analysis. Guidance ranges from recommender systems [Won+16; GW09] to user's modeling [BM07; Maz+10]. In the following, we describe guidance approaches and differentiate our work from previous evaluation studies.

Willett et al. [WHA07] introduced *scented widgets*, which are common UI-elements enhanced with knowledge derived from other users' interaction choices. The authors underline that the introduction of such elements may flatten the difference in performances between expert and novice users. Gotz et al. [GW09] introduced behavior-driven recommendations, showing improvements in the completion time and correctness of

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

results. In the field of data mining, Bernstein et al. [BPH05] developed an intelligent ontology-based assistant that supports the choice of proper data mining algorithms with respect to the specific problem setting. Their results suggest that also expert users need guidance. Streit et al. [Str+12] generate an analysis model that is used for supporting analysts with their tasks. The advantage of this work is the provision of different degrees of guidance.

Similarly to these approaches, we want to evaluate whether the introduction of guidance leads to performance improvements among study participants. However, our aim goes beyond the evaluation of the effectiveness of a specific tool. In fact, in contrast to such approaches, *we aim to understand how such effectiveness varies according to the task, and what are the effects of different degrees of guidance in relation to different levels of user's expertise.*

3.3.2 User's Knowledge and Mental State

Chen [Che05] distinguishes between two main types required to make progress during an analysis: *operational* and *domain* knowledge, whether the user is able to *interact* in an effective way with the analysis tool, or possesses the necessary *domain* notions to interpret the context and the data. In our vision, different types of guidance may be necessary according to what kind of knowledge the user is missing. Thus, our aim is to investigate if similar guidance degrees have different effects according to the knowledge gap, i.e., lack of operational or domain knowledge.

A last research branch related to our work, is the one studying the relations between the visual analysis and the development of the user's mental state and sentiments, and the effects of such feelings on the analysis itself. Sacha et al. [Sac+16] point out that during an analysis, there is always a match between the uncertainty present in the data and the trust that users develop while proceeding with the analysis. The more the exploration advances, the more their trust grows. Although not explicitly mentioned, also the guidance may contribute to increase or decrease the user's trust. Many other psychological aspects are also connected to the analysis process. Celik et al. [Cel+13] point out that frustration and sadness are often connected to the inability to perform a task. Similarly, Kapoor et al. [KBP07] show that it is possible to automatically infer and predict the growth of frustration, during the execution of a task, and thus they identify possible thresholds for triggering guidance. However, the effects of guidance on user's frustration, and in general on users' mental state, have not been studied yet.

In order to understand how guidance affects the development of sentiments during data analysis, we asked participants of a user study to rate their degree of frustration, trust, and confidence after solving a set of tasks. We then relate these values to the provided guidance, to the user's expertise level, and to the analysis outcome.

3.4 Aims and Terminology

The aim of this work can be summarized by the following questions:

- 1) How do different guidance degrees affect the performance and the mental state of users with different degrees of previous knowledge?
- 2) Do the effectiveness and the effects of guidance vary according to the type of task the user has to solve?

Our assumption is that three dimensions play an important role in the design of guidance for visual data analysis, these are the task type, the knowledge of the user, and the guidance degree. We want to test how the variation of one of such dimensions influences the others, and ultimately the analysis outcome, in terms of user performance and mental state. We start describing these three dimensions, before formalizing our aim in terms of rigorous hypotheses.

Knowledge and Task types The first factor we describe, is the knowledge required to complete a task. Two kinds of knowledge are usually required to complete a visual data analysis: *operational* and *domain knowledge*. Our aim is to test whether the type of knowledge involved influence the effectiveness of the provided degree of guidance. According to the distinction between operational and domain knowledge, it is possible to discern two general types of tasks:

- *Exploratory tasks* relate to operational knowledge and therefore to the ability to interact with the tool. These tasks require basic interaction abilities, like choosing among different interaction means (e.g., filter, selection) and using them effectively.
- *Domain tasks* require domain specific knowledge to be successfully completed. These tasks are related to the ability to reason and connect a given domain concept to the task and data under analysis.

User Knowledge A second dimension we address, is the distinction between degrees of user's competence. Usually, a lack of user's knowledge, which can be also seen as the difference between the knowledge required to solve the task and the actual user's knowledge, is what we call a knowledge gap. When this happens, the user might have a hard time completing the task, and the analysis may stall. We hypothesize that different degrees of guidance would have different effects on users with different levels of knowledge relevant to the task. We distinguish between:

- *Knowledgeable users*, possessing the knowledge required to complete the task (i.e., operational or domain knowledge, see previous paragraph).
- *Novice users*, who may not possess the knowledge required to complete the task, with the exception of pregress expertise.

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

Guidance degrees Finally, we distinguish among three different degrees of guidance which we provided to the study participants to assist the completion of their tasks. Our assumption is that users with different knowledge may require different degree of guidance, and that the fact of having more or less support while solving the task would have variable effects on the task performance and the user's mental state. According to the guidance degrees described by Ceneda et al. [Cen+17]), we list the types of guidance we included in our study.

- *No guidance*: When no guidance is provided, users have to solve the tasks on their own. This is translated into the provision of simple visualizations, without any further support. According to Ceneda et al., the provision of additional aggregated values (i.e., min, max, outliers) does not constitute higher guidance.
- *Directing guidance*: This kind of guidance aims at providing different analysis options. Therefore, on top of the basic visualization, we indicate possible analysis paths. In the specific, interesting data subsets are recommended to the user, but the system may also recommend actions to proceed the investigation.
- *Prescribing guidance*: This is the highest degree of guidance. It aims at providing step-by-step instructions to reach the result. Among the different analysis paths and recommendations (see Directing guidance), the system picks one and provides it to the user, who must follow the indications (the different steps) to reach the final result.

The aforementioned three dimensions concur to the provision of effective support to the user. Considering all of them together allowed us to reason about the effectiveness of different guidance types in different situations. In particular, we investigate 1) if guidance can compensate for a lack of user's knowledge i.e., if there is a noticeable difference among novice and knowledgeable users supported with similar degrees of guidance. We examine 2) if some degrees of guidance are better suited than others for a given task type i.e., if some degrees of guidance are better suited for exploratory analysis or to complete domain tasks. Furthermore, by comparing the results of same users under different conditions of guidance, we study 3) if guidance can affect, in positive or negative, the performance of such users and the development of sentiments and mental map.

3.5 Hypotheses

We formalize the aim described earlier in terms of different research hypotheses, which we grouped into two hypotheses groups, H1 and H2. Hypotheses in H1 focus on the variation of user's performance metrics, while those in H2 consider the mental state and the feelings of the user, in response to the provided guidance.

Hypotheses group H1 Our first aim was to investigate the effects of guidance on task performance. At first, we analyze the effects of guidance on novice users and evaluate if the positive effect of guidance is mitigated in knowledgeable users. Our assumption

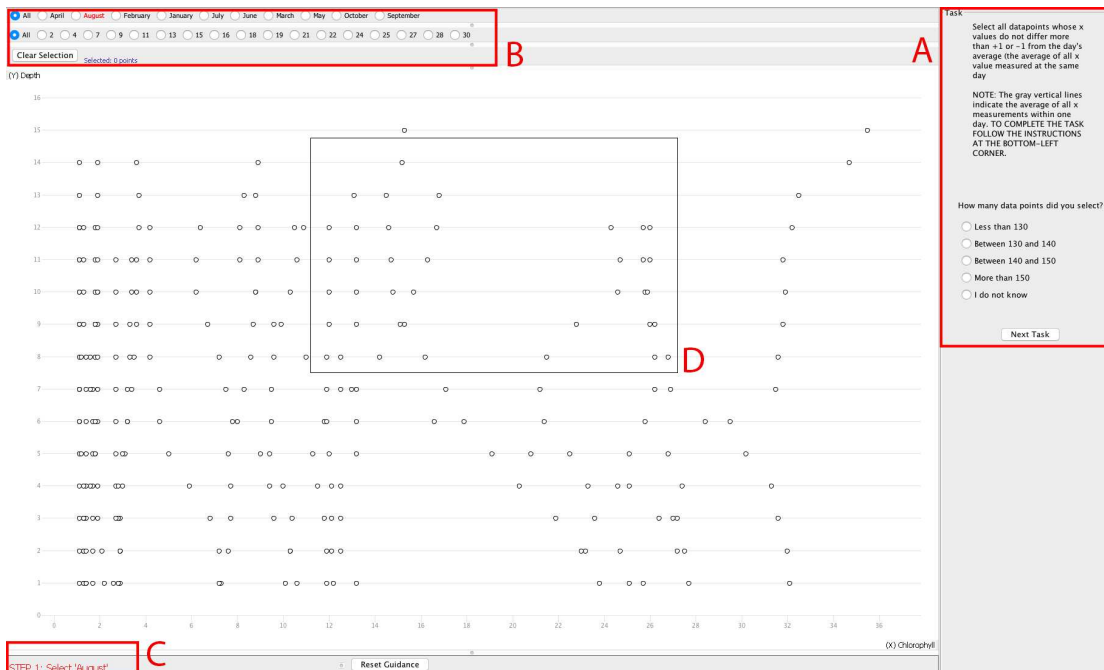


Figure 3.1: Interface of the evaluation environment. In the upper right corner (A), a text-box shows the current task. A further text-box, gives indications to interpret the guidance suggestions. On the top left (B), some combo-boxes give the possibility to filter the dataset. On the bottom of the visualization (C), a text-box shows the step-by-step instructions to reach the desired results (in case prescribing guidance was provided). At the center of the visualization (D), the rectangular selection tool is shown.

is that knowledgeable users may still benefit from guidance in terms of reduced task completion time. This is formalized in the following hypotheses:

- H1.1** A high degree of guidance causes significant improvements in task performance (timings, correctness, distance, total-steps) of novice users.
- H1.2** A high degree of guidance reduces completion time and amount of steps for knowledgeable users.

Hypotheses group H2 Our second aim was to evaluate the mental state of users when receiving different degrees of guidance. In particular, we wanted to understand if guidance causes a positive effect on the the user's mental state. Specifically, if guidance increases the user's confidence in the analysis results, or if in some situations the guidance can frustrate the user. This is formalized in the following hypotheses:

- H2.1** A high degree of guidance causes a significant improvement in participants' confidence in their results.

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

H2.2 A high degree of guidance causes more frustration for knowledgeable users than for novice users.

3.6 Study Design

To verify the hypotheses in H1 and H2, we designed a user study comprising six specific tasks (3 exploratory + 3 domain tasks) which we asked 65 participants under different conditions of expertise and guidance to solve. The hypotheses lead us in designing the evaluation environment and the evaluation procedure as follows.

3.6.1 Data

We use a dataset from the USGS program of research and observation in San Francisco Bay [CS16]. This dataset in combination with a careful task design allowed us to evaluate the effectiveness of guidance on both exploratory and domain tasks. The dataset contains multiple daily measurements of water samples collected along the 145 kilometer transect of the San Francisco bay. The whole dataset spreads over various decades (1969-nowadays), but for our study we selected only specific subsets, spanning roughly one year each. In particular, from the main dataset, we extracted *six* subsets. Each dataset was associated with exactly one task to avoid learning effects. Three datasets were used for domain related tasks, while the other three for exploratory tasks. Each dataset is equivalent to the others in terms of number of data dimensions involved. They just differ for the focus on a specific dimension of the original dataset. We complemented the six datasets with derived statistical values (e.g., average, max, min). We used these derived values as a base for *directing* guidance, to point the user to interesting data during the execution of the tasks.

3.6.2 Participants and Evaluation Sessions

We had 65 students at bachelor level participating in our study. They all are students in computer science and attended a course in information design and visualization, preceding the study, which implies a certain knowledge about the visual environment they were provided with. Nevertheless, we considered all the students as novice participants, since they never performed the analysis on the given dataset, nor possessed any domain knowledge about the topic. Before presenting the tasks to the students, we conducted a pilot testing with four participants to correct minor errors and fine tune the tests.

For the evaluation sessions we utilized EvalBench [AHR13], a software specifically designed to evaluate interactive visualizations (see Figure 3.1). The interactive visualizations were developed with Java, using the Prefuse library [HCL05], and TimeBench [Rin+13] to manage the temporal aspects of the data.

Study structure The user-study was divided into two subsequent evaluation sessions as detailed in Figure 3.2: one session dealing with exploratory tasks and the other session

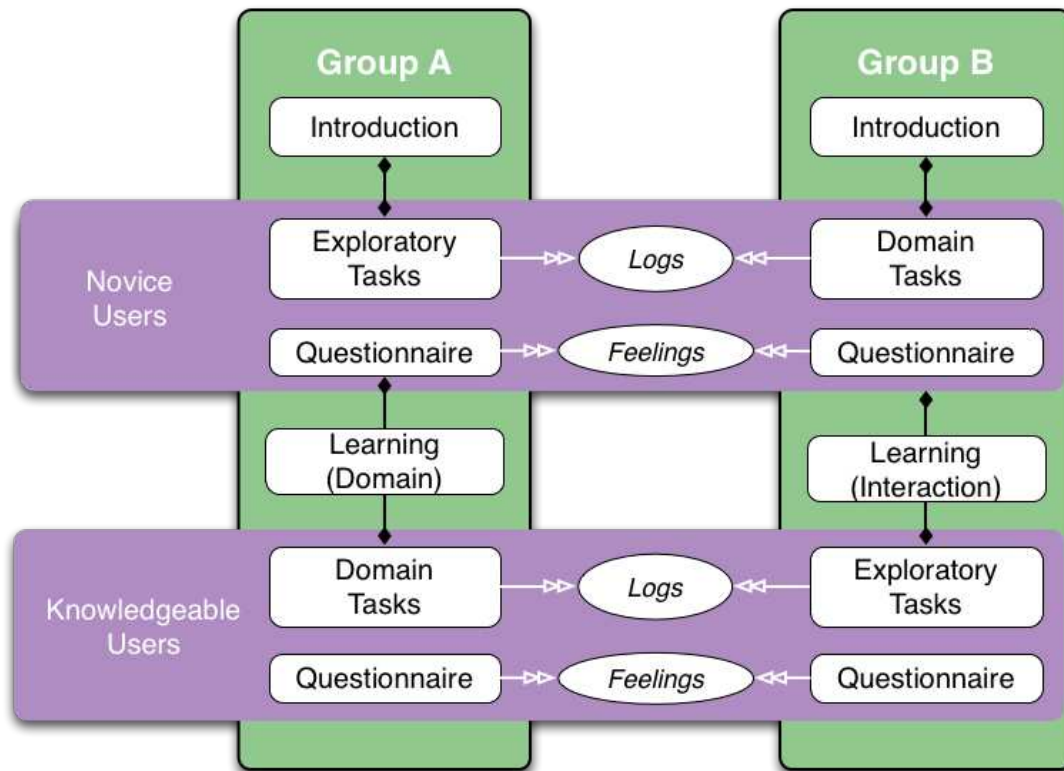


Figure 3.2: General structure of the user's study. We conducted two parallel evaluation sessions. After a short introduction, the participants performed two set of tasks. Each task was followed by a series of questions. Between the two task sessions, the participants had an active learning phase, where they were instructed either in interaction concepts or in domain concepts respectively. Subsequently, the participants completed the second set of tasks. A cross structure was chosen to minimize the learning effects on the participants.

dealing with domain tasks. We divided the participants into two groups, group A and group B, each of them executed both exploratory and domain tasks in the two task sessions, but group A performed exploratory tasks in the first session of the study and domain tasks in the second session, and group B did it the other way around. We did this to avoid learning effects of the participants and compare the execution of the same tasks with different levels of expertise.

At the beginning, both groups received an introduction to the main topics of the user-study. We told them that they were going to execute some tasks and that they were (possibly) going to receive guidance during this execution. They were not given any other information, except that the data regarded biological measurements extracted from water samples of the San Francisco bay, and that they were not allowed to use any external

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

help. We intentionally decided not to provide them with any further information about the interaction means, or about the specific domain concepts in order to simulate the behavior of novice users.

Thus, in the first session, the participants did not have any experience with the data, visualizations, and tasks. Furthermore, we did not provide them with any additional knowledge that might have been required to solve the tasks. For this reason, we considered them novice users. In the second session all participants had already some experience with the data, visualizations, and tasks. In addition, we added a learning session in between the two task sessions to train the participants in the concepts necessary to complete the following tasks. After undergoing the learning phase, we considered these participants knowledgeable users. We chose this study design to ensure that both types of users conduct both types of tasks while avoiding learning effects on the two groups.

Learning Session The first task session was followed by a briefing for the next session – a session in which participants were instructed in domain or in the operational concepts required to solve the three remaining tasks in the subsequent session. Thus, group A, after completing the exploratory tasks, was instructed with domain concepts, necessary to solve the domain tasks session. Group B, instead, received an education about interactive means and the exploratory session. In this learning phase the participants belonging to both groups had also the possibility to interact with a sample tool and further sediment the acquired knowledge. This allowed us to compare the performance of novice users (no prior knowledge) with that of users instructed in concepts relevant for solving the tasks (exploratory and domain tasks), mitigating at the same time the possibility of learning effects on the subsequent series of tasks. In fact, with such cross-structure, the expertise group A acquired while conducting the first session was not needed to complete the subsequent session of domain tasks. The same holds true for the domain knowledge group B acquired during the first session, which was not needed to solve the next tasks. We did not measure precisely the increase of knowledge, in terms of learned concepts, due to the learning session. However, from the results of the study we could see an increase in the number of participants who were able to solve the tasks without guidance after undergoing the learning session. In average, 10% more participants was able to solve exploratory tasks without guidance. This percentage increases to 20% for domain tasks. This means that 20% more participants could solve domain tasks without guidance after learning the appropriate domain concepts.

To collect the data necessary to test H1, the system recorded automatically timings, correctness and the number of operations required to complete the tasks, see Table 3.1. After the execution of each task, we asked the participants to answer ten questions about the visualizations and the interactive means (i.e., were they sufficient? were they useful?) the tool offered, as we wanted to test whether they interpreted the visual encodings correctly. We then asked them to evaluate the guidance they received. We encoded the possible answers as multiple choices, but we also let the participants add free text if they felt the options provided were not sufficient. To test H2, a further set of questions

asked the participant about their feelings while solving the task (see Table 3.2). All these subjective feelings [Cel+13; KBP07] were measured on a five-point Likert scale. At the end, we also collected the interaction logs (e.g., hovering a point, changing the selection, filtering the dataset) for evaluation and for extracting further metrics (see Section 3.7.2).

3.6.3 Task Design

We designed a total of six tasks: three focused on operational knowledge and three on domain knowledge.

Exploratory tasks These tasks are related to a user’s operational knowledge and his/her ability to interact with the analysis tool. We required the participants to perform a number of interactions to explore the dataset. We did not relate these tasks to any domain concepts, but rather asked the users to find and select specific data values, without any associated meaning. A typical exploratory task required the participants to isolate data points with certain characteristics by iteratively using the interactive means provided by the tool. In other words, an exploratory task consists of long sequences of selections and filter operations. The number of actions and the reasoning effort required to solve an exploratory task constitutes the main difference to domain tasks. We designed these exploratory tasks so that the only knowledge required to correctly and efficiently solve them was being able to interact with the visualization tool. In comparison, domain tasks required domain knowledge, while almost no interaction, besides simple selections. We designed exploratory tasks in such a way that it was possible for the participants without the advanced interaction means we introduced during the learning session. In total, around half of novice participants was able to complete correctly the exploratory tasks without guidance just by using the basic interaction means offered.

As mentioned earlier, after the first tasks session, we lead the participant through a learning session to let the participants acquire the knowledge needed to solve the following tasks. For what it regards exploratory tasks, we introduced the participants belonging to group B to the use of some advanced interaction techniques, like for instance the rectangular selection of multiple data points and the use of filters, etc. All the interaction means were available to all the participants from the beginning of the study. However, we assumed that the competences we taught to group B during the learning session would allow them to complete the tasks more efficiently, in respect to novice users. Furthermore, while knowledgeable users were presented and had time to experiment the different interaction means, novice users had to discover them while solving the tasks, marking another difference between the two groups. As a consequence, we expected a difference in the performance, as well as in the frustration level and confidence of these two types of users.

Domain tasks The same design principles led the design of domain tasks. We based these three tasks on three specific domain concepts: hyper-salinity of sea water, periods

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

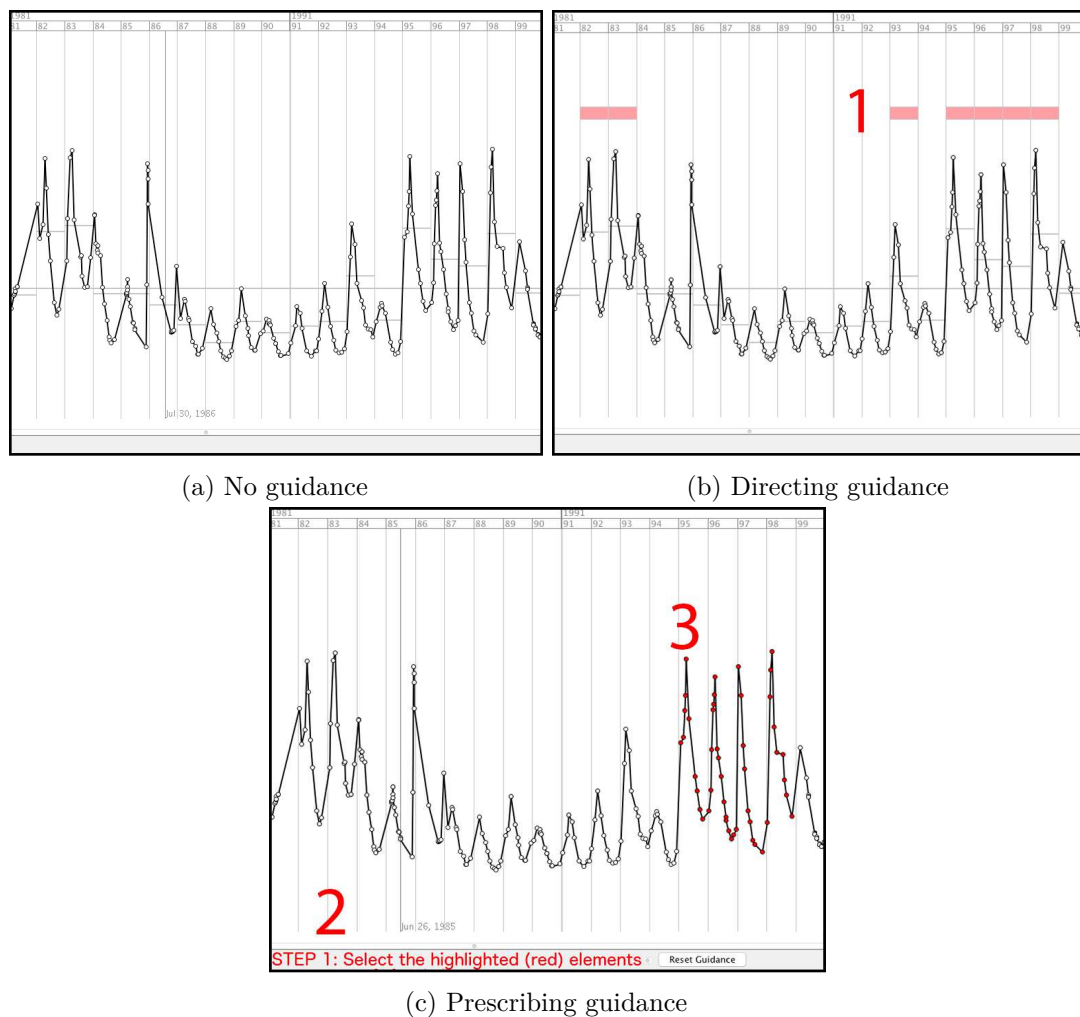


Figure 3.3: The same domain task is supported with different degrees of guidance. (a) no guidance: a time-series line chart shows the variation of water salinity in different years. (b) directing guidance: possibly interesting data is highlighted to address the analysis (e.g., data representing high salinity values). These periods are signaled in red at the top of the visualization (1); (c) prescribing guidance: step-by-step instructions are presented to the user (2) together with the highlighting of interesting data points (3).

of droughts in a given year, and dangerous low concentrations of nutrients in the bay water. In particular, we asked the participant to recognize a period of drought and a condition of hyper-salinity by analysing different time-series showing the development of water salinity in a given time period. In a third task, we asked the participants to reason about low concentration of nutrients by exploring a scatter-plot visualizing water nutrients at different depths in a specific region of the bay.

Also for these tasks, we worked to mark a difference between novice and knowledgeable participants, by providing the latter group, during the learning session, an introduction to these domain concepts, including exercises to consolidate the knowledge. For instance, in one of the domain tasks, the participants were requested to select all the data points corresponding to a period of drought. We explained how to recognize these periods just to the knowledgeable users, while the novice ones relied just on their individual idea of the concept and on the guidance suggestions, telling them that for instance, a clear sign of a period of drought is the raise of average water salinity in a given period. Thus, for novice users who did not receive any guidance it was sometimes not possible to find the correct answer to such tasks. In total, just one third of novice participants was able to complete correctly the domain tasks with no guidance. The learning session affected the participants' performance. In average, 20% more participants completed correctly the domain tasks with no guidance.

3.6.4 Concrete task examples

To give the reader a better idea about the task design, we describe one domain and one exploratory task in more detail. For completeness, we created two more tasks for each task type, for a total of six tasks. In the following, we describe just one of them, for each type.

Domain Task The participants had to solve the following domain task under different conditions of guidance. We asked them to "*select all the data points falling in the longest period of drought*". To solve this task, the study participants were presented a line-chart visualizing the fluctuation of salt concentration in a given period. On top of this basic visualization we added guidance. All the visualizations used for this task are illustrated in Figure 3.3. The figures show the encoding of the three guidance degrees. When solving the task, a user would either know directly (if knowledgeable) or possibly reason (if novice) that a period of drought affects the mineral composition of the water. For sea water, one of the most obvious results is that the concentration of salt increases. As a consequence, a user should have selected as the correct answer the longest period with the highest salinity values.

Aside the line chart, additional lines encoded the average salinity values of every visualized year. We shared this same visualization for all the different guidance degrees. On top of this visualization, we added additional visual clues to support increasing levels of guidance. For instance, when *directing guidance* was provided, we highlighted data points of years with particularly high average temperatures and salinity values (Figure 3.3b). These hints point users towards data regions/subsets that are helpful to solve the task. Directing guidance, per definition, does not give exact instructions to solve a task but rather recommends and directs the user towards interesting data regions. In the last scenario, *prescribing guidance* was provided. We led participants along a selected analysis path. While the users could freely interact with the tool, we provided them with precise

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

step-by-step instructions in textual form to follow this chosen analytical path and find the correct answer. Since the task outlined in Figure 3.3b) is a domain task it requires domain knowledge and reasoning to solve it rather than a operational knowledge. With the aim of limiting the effect of operational knowledge on the resolution of such tasks, we limited the required interactions to simple selections. In case of *prescribing guidance*), this meant that we highlighted the correct data points and asked users to select them by simply clicking on them.

For all the three guidance degrees, the correct answer was to select the data points highlighted in Figure 3.3c. The resolution of domain tasks relies mainly on the users' knowledge, and in case of novice users, on their ability to reason. Therefore, we expect novice users, especially without the guidance support, spending more time on reasoning and having rather approximated results. However, also when no guidance was provided, a percentage of novice participants were able to solve the domain tasks without guidance.

Exploratory Task We created a second set of tasks focusing on operational knowledge. When asking the participants to solve such tasks, we avoided any reference to domain concepts and just asked the participant to look for data with specific characteristics, without focusing on the meaning. In particular, as already mentioned, we structured such task as a long sequence of filtering and selections, to reach and select the desired data. In one task, we asked the participants to *select, for each measuring station, the FIRST data point such that, the value of Salinity (x axis) is greater than 2, but lower than 3 salinity units*. As it can be seen, no domain knowledge is requested except reading and understanding the graph (in this case, a scatter plot) and interacting with the tool performing selections and filtering. The visualizations used for this task, according to the provided guidance degree are portrayed in Figure 3.4. In this case, the participants were presented with a scatter plot representing values of salinity (x-axis) in relation to the change of water depth (y-axis). Since this task is focused on interaction, we provided the users with means to select and filter the data, for instance, filter according to the measuring station that captured the measurement. Other advanced interaction means were also provided, for instance the possibility to perform lasso selections and avoid the need of multiple clicks.

In a first scenario, some participants received no guidance. The participants dealing with such task were presented just the plain graph (see Figure 3.4a). In this first case (no guidance), we expected the participants to filter the dataset by selecting and exploring the measurement of all the different measuring stations (also the ones with no interesting measurements) and select the data with the requested characteristics. In a second scenario, other participants were supported by directing guidance (see Figure 3.4b). In such case, the participants could also rely on the highlighting of the measurements falling in the requested range [2 – 3], therefore being directed in the data retrieval. In addition, we also highlighted with a different color the filtering option i.e., the measurement stations that captured those data values, leading to the requested data points, in such a way to signal to the participants a set of possible filters to choose. In this way, we wanted to

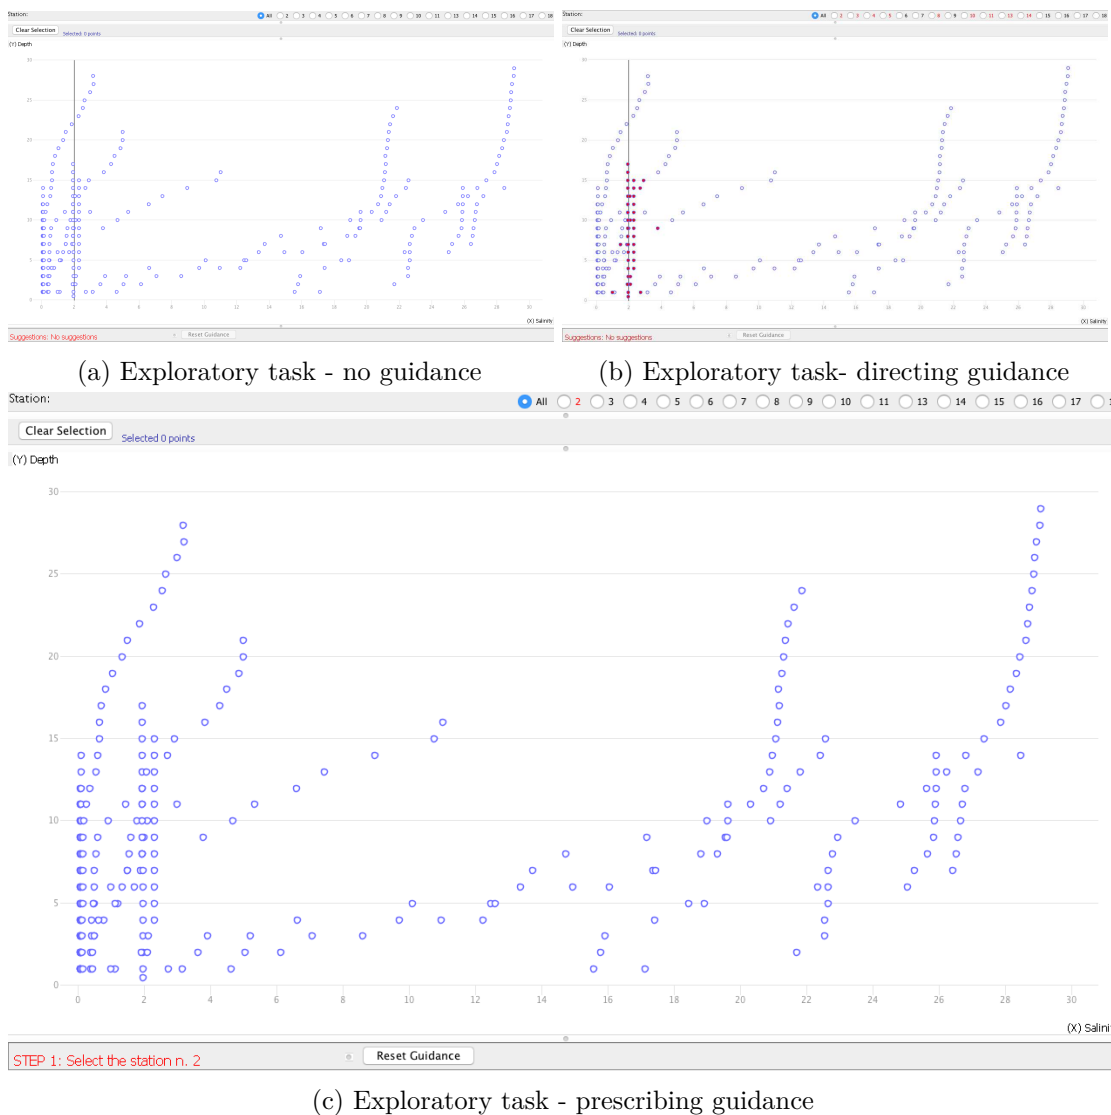


Figure 3.4: The same exploratory task supported with different guidance. (a) no guidance: a scatter plot shows values of salinity (x-axis) in relation to the change of water depth (y-axis). A widget allow users to filter the dataset in respect to the measuring station. (b) directing guidance: possibly interesting data points and filtering options are highlighted; (c) prescribing guidance: step-by-step instructions and highlighting of correct values as well as the filtering actions to be performed.

guide the interaction, by either highlighting the data and the filtering options necessary to make the requested data visible. Participants receiving directing guidance could also skip unnecessary actions, checking the highlighted data. Finally, in a latter scenario, a part of the participants had to rely on prescribing guidance (see Figure 3.4c). Similarly

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

to domain tasks, this kind of guidance consisted of a list of instructions, in addition to the highlighting of the data. However, while usually for domain task this list was composed by one or two actions, for domain tasks it consisted of a long sequence of filtering and selection steps to simulate a thorough exploratory analysis. With prescribing guidance, the user had to follow dutifully an average of twenty consecutive steps of alternate filtering and selections, to complete this task and select the required data.

Obviously, the correct answer to this task was the same despite the different provided guidance types. In the context of this task, we expect faster interactions with increased guidance. However, we also analysed the variation of frustration levels and confidence in users that had to strictly obey to the prescribed actions to complete the task, to see if they felt restricted by the guidance.

3.6.5 Visual Encoding Design

We chose basic visualization types for the study. We chose scatter plots, line charts, and temporally aggregated charts showing data values for each year side by side (Figure 3.5). We wanted to keep the visualization aspects of the study as general as possible, so to not interfere with the outcome of the analysis and the effectiveness of the provided guidance. At the same time, we chose these visualizations also because the participants were familiar with them, but also effective for the given tasks.

Scatter plots represented data values as dots with one of three variables (either water salinity, chlorophyll, or suspended solids) on the x-axis and water depth on the y-axis. Line charts represented dots connected by lines with time on the x-axis and salinity on the y-axis, and temporally aggregated charts juxtaposed yearly (x-axis) oxygen values (y-axis). The visual encoding we used are portrayed in Figures 3.3, 3.1, and 3.5.

Interaction means For all chart types, we provided basic interactive means such as details on demand when hovering a data point, selection of single data points when clicking on them, rectangular selection (by dragging the mouse to span a rectangle) for selecting multiple data items, and deselection of all data points with a right mouse-click. The successful selection/deselection of a data point was visualized by a change of fill-color. For exploratory tasks, we also provided radio-buttons to filter the displayed data points according to the time-stamp of the measurement (month and day). We did not provide filtering for domain tasks, as it was not required. This difference in the interaction means did not influence the outcome of the study, as domain and exploratory tasks were never compared directly. As a final remark, all the participants were provided from the beginning of the study with the possibility to use all the interaction means. The participants belonging to Group B were introduced during the learning session to the use of all the available interaction means.

3.6.6 Guidance

In their work, Ceneda et al. describe three degrees of guidance: orienting, directing, and prescribing [Cen+17]. However, since common practices in visualization such as axis labels could also be seen as a very low level of guidance, the border between no guidance and orienting (giving some hints for orientation) becomes blurred. Thus, to avoid confusion and have a clear baseline for comparison we implemented just three of them: (1) *no guidance*, (2) *directing guidance*, and (3) *prescribing guidance*. By design, the participants received all three guidance degrees, one time in each task set. In total, each participant received the same degree of guidance twice: once while executing exploratory tasks, and once while performing domain tasks.

When no guidance was provided, we presented the participants a common visualization (e.g., a line chart) showing one of the data subsets, with some additional data visualized, like average or minimum values (see for instance Figure 3.3a and Figure 3.4a). When directing guidance was provided, participants received an additional indication about possible interesting data or actions to consider. Figure 3.3b and Figure 3.4b show the encoding chosen for directing guidance. Interesting interaction options were highlighted for exploratory tasks (upper side of the interface), while interesting data-points were




Figure 3.5: A chart representing aggregated values of nutrients (y-axis) according to the measurement station.

highlighted for domain specific tasks within the visualization. Finally, participants receiving prescribing guidance were provided with step-by-step instructions to reach the desired results as shown in Figure 3.3c and Figure 3.1. The instructions were given as red text, in the bottom-left of the visualization. Prescribing guidance produces mandatory actions [Cen+17]. Hence, although participants could perform any other action and deviate from the analysis path, the instructions proceeded only after the user conducted the required steps. Moreover, we provided them with the possibility to restart the guidance process. We motivate the introduction of this extreme degree of guidance to explore the full range of guidance possibilities. It is worth clarifying that this high guidance degree does not correspond to a simple presentation of results, and it also differs from a pure automated data analysis [Cen+18]. The user is always required to interact and confirm the different steps and moves. Moreover, as already pointed out, the participants always had the possibility to deviate from the suggested analysis path for further analysis.

3.7 Results

Sixty-five participants submitted their results and the interaction logs.

3.7.1 Analysis Approach

We analyzed the logs and the results of the user-study using the R environment for statistical computing [Tea14]. Our aim was to spot significant differences between subgroups. In our study, we mostly compared three groups (i.e., the three guidance degrees) among each other, for each task type and expertise condition. Hence, we used the Kruskal-Wallis test [KW52], which is a non parametric test similar to ANOVA which can be used with more than two groups and is also well-suited for comparing results obtained from Likert scales [DD10].

In a few tests, we compared the variation of single metrics (e.g., frustration) in users with different expertise. For instance, the tested variation of correctness in novice and knowledgeable users, for the same type of tasks (e.g., exploratory tasks). In such cases, since we had to compare just two groups (i.e., novice vs knowledgeable) we applied the Mann-Whitney-Wilcoxon test [MW47]. We never compared directly exploratory and domain tasks among each other.

Since we performed many tests and hence to account for the probability of a false positive discovery, we applied to all the tests the correction technique by Benjamini and Hochberg [BH95]. This choice implies that, while usually a common threshold value is chosen for all the tests (usually set to $p \approx .01$), in our study it varies according to the test. In the specific case, for each test two p-values are calculated, p which is resulting from the test, and $p_{corrected}$ which is calculated from p . The corrected p-values are calculated considering the total number of tests performed and an initial significance level of 0.05. Hence, when we report on the acceptance of a test, we will also report the

correspondent corrected p-value. When a significant difference was detected then we performed post-hoc tests to compare the different groups among each other and evaluate the pairwise differences. As a final step to the analysis, we manually inspected the data and further analyzed the results with box-plots and scatter plots.

Performance	Description
<i>Completion Time</i>	A timer measured the interval between the start of the task and the submission of an answer.
<i>Correctness</i>	A real number in [0,1]. This value is a weighted ratio between correctly selected data items and all selected data items.
<i>Distance</i>	A real number in [0,1], measuring the semantic distance of the selected data items from the correct ones.
<i>Total Steps</i>	The total number of actions (clicks, filter, etc.) required by a user to complete a task.

Table 3.1: Performance Metrics. We recorded these metrics while the participants executed the tasks. The completion time was provided directly by the evaluation environment EvalBench [AHR13]. The others (correctness, distance, steps) were calculated from the interaction logs (see Section 3.7.1).

3.7.2 Users' statistics

Performance The system automatically extracted a set of measures to understand task performance (see Table 3.1). These measures include the total number of actions conducted by the user to complete a given task: number of clicks, rectangular selections, and applying filters. We also computed a correctness value to reflect the ratio of correctly selected data items to the total number of correct data items weighted by the total number of selected data items. We included this measure to account for cases in which participants select huge numbers of data items which makes it likely that they also select some correct ones. Another measure, the distance, was computed to quantify the semantic distance of the answer, in terms of selected data items, to the correct answers. We calculated this metric by averaging the temporal distance of the selected points from the solution.

$$distance(avg) \approx \frac{\sum(temporal_dist(x,solution))}{total_data_selected}$$

Since all the tasks comprised temporal aspects, measurements falling in different time periods were considered distant. We did this to understand if wrongly selected data items are semantically close to the correct ones (e.g., they are in the same month) or if they are completely wrong (e.g., they are in different years).

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

Mental State	Description
<i>Lost</i>	We asked the participants how lost they felt while executing the task.
<i>Frustrated</i>	We asked the participants how frustrated they felt while executing the task.
<i>Confident</i>	We asked the participants how confident they felt about the correctness of the submitted result.
<i>Easy</i>	We asked the participants to evaluate how easy the task was.
<i>Guidance Appropriate</i>	We asked the participants if they considered the guidance they received appropriate to solve the task.

Table 3.2: Indicators of users' mental state. We asked participants to answer some questions regarding their feelings after each task. Each variable was rated on a five-point Likert scale. We then related the users' feelings to the degree of guidance they received, and to their knowledge level.

Feelings Besides performance measures, this study comprises a set of measurements dealing with user's feelings. Guidance approaches inherently deal with users. Hence, it is important to understand how guidance affects the development of user's psychological aspects. These are listed in Table 3.2. Similarly to user's knowledge, such psychological aspects are difficult to measure and quantify. However, for our purpose of deriving correlations and tendencies rather than quantitative values, we use a simple qualitative scale to measure the participant's own assessment of their feelings. Usually, this method may be influenced by the personality of the participants, who may present extreme/average input styles. However, such drawbacks were mitigated and averaged by the number of participants involved. Therefore, we did not apply any further correction to those tests.

3.7.3 Outcome

The tests indicate that guidance has an overall positive effect on users' performance and mental state. Guidance is particularly successful for novice users solving exploratory tasks and can easily compensate for a lack of operational knowledge. Instead, the tests highlight that for domain tasks, at least a minimum of knowledge should be possessed by the users, not only to understand the tasks and the context, but also to interpret correctly the guidance.

Our study highlights that guidance is important in complex scenarios: We show that the benefits are particularly pronounced when domain knowledge and reasoning are needed: for knowledgeable users solving domain tasks, the results obtained with directing

guidance were in line with those obtained by prescribing guidance. However, our study reveals that guidance may even have a bad impact on the analysis if the guidance degree does not match the knowledge gap and users' expectations. From our results we can see that novice users, tend to trust excessively the guidance suggestions, and that the prescribing guidance degree may sometime frustrate knowledgeable users. The tests revealed that directing guidance is beneficial for knowledgeable users who are able to interpret and judge correctly the suggestions. When assisted with this kind of guidance, participants obtained performances similar to prescribing guidance. On the other hand, this degree produces no improvements (same results as no guidance) when provided to novice users. Thus, for novice users the prescribing degree of guidance seems better suited.

In the following, we discuss the results in relation to our hypotheses. We then outline observations and interesting additional findings. A summary of the study outcome can be found in Table 3.3.

H1.1 We investigated if guidance positively affects the performances of *novice users*. The results reveal differences in the performances of novice users receiving guidance, and those who did not receive any guidance. The box-plots in Figure 3.6 show that novice users perform significantly better with prescribing guidance (for both, exploration and domain tasks), in respect to the other guidance degree (directing) and to the scenario in which guidance was not provided at all. Hence, H1.1 can be accepted. However, directing guidance shows no significant improvement of performance of novice users compared to no guidance.

Task Completion Time Novice users solved *exploratory* and *domain tasks* faster when supported by prescribing guidance. In fact, we found a significant difference between prescribing and no/directing guidance ($p \approx 0.01$, $chi - sq = 23.4$, $df = 2$ for both task types, see Figure 3.6a). For both task types, no significant differences of timings were reported between directing and no guidance, while in general, completion times resulted higher for domain tasks, in respect to exploratory tasks. We noticed some cases, in which the participants who did not receive any guidance solved their tasks faster than those receiving directing guidance. We guess that this difference can be explained with the additional time required to interpret the guidance suggestions, especially for novice users. Hence, we imagine that while the participants who received no guidance started immediately to look for the correct answer, the users who received directing guidance (i.e., pointed to possible interesting subsets of the data) lost some time judging the applicability of the suggestions.

Correctness Correctness values are also influenced by the guidance degree. Figure 3.6b shows the box-plots for *exploratory* and *domain tasks*. Looking at the exploratory tasks, tests reveal significant differences between prescribing and no/directing guidance for novice users ($p \approx 0.005$, $chi - sq = 27.4$, $df = 2$). Similarly, for novice users solving domain tasks, we found very significant differences in correctness values between prescribing

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

guidance and the other two degrees ($p \approx 0.007$, $chi - sq = 27.1$, $df = 2$). Moreover, correctness values showed that when we provided directing guidance to novice users, they did not answer more correctly to questions compared to no guidance. In average, the correctness increased with increased guidance, but the guidance itself could not replace the lack of knowledge, in novice users. This is particularly true for domain tasks, where novice users had similar results with both no guidance and directing guidance. For exploratory tasks, the charts show an increased correctness between no guidance and directing guidance, but the difference was not significant.

Half of the novice users (49.6%) completed correctly the exploratory tasks without any suggestion (no guidance), this number increases to 66% for those guided by directing guidance, and finally, the majority of participants receiving prescribing guidance (> 90%) completed correctly these tasks. On the other hand, the results highlight that guidance cannot completely overcome the lack of knowledge, in case of domain tasks. Just 32% of the novice users completed the tasks correctly, this percentage raises to 37% for directing guidance, and 85% for prescribing guidance.

Distance For novice users, and similarly to the other measures, we noticed a significant difference in the distance measures between participants assisted with prescribing guidance, and participants assisted with no guidance. This holds true for both *exploratory* ($p \approx 0.0002$, $chi - sq = 16.9$, $df = 2$) and domain tasks ($p \approx 0.0003$, $chi - sq = 15.7$, $df = 2$). For interaction tasks, the tests did not highlight any difference in the distance measure between directing and prescribing guidance. Moreover, novice users receiving directing guidance had results closer to the correct values (smaller semantic distance), compared to those who received no guidance (see Figure 3.7). For domain tasks, the lack of domain knowledge may have had nullified the effectiveness of directing guidance, as the tests did not highlight any significant improvement in respect to no guidance.

Total Steps *Novice users* performed an average of 42 actions to complete a task: 13 filters, 6 multiple selections, and 23 single selection clicks. For *exploratory tasks*, the users receiving directing guidance performed similarly to those who did not receive any guidance (approx. 83 steps each). However, users provided with prescribing guidance needed on average only half the amount of steps (45 steps) which presents a significant difference. For *domain tasks*, the influence of different degrees of guidance is even more significant. On average, participants provided with no guidance completed a task with 24 actions. Directing guidance lowered this number to 15 actions, while novice users supported with prescribing guidance, took on average 8 actions. The statistical tests reported a significant difference in the number of steps required by novice users performing *domain tasks*, between prescribing guidance and no guidance; there was no significant difference, on the other hand, between prescribing and directing guidance (see Figure 3.7).

H1.2 We hypothesize that a high degree of guidance may reduce completion time and the number of steps needed for *knowledgeable users*. H1.2 can be accepted partially. The tests did not show significant differences in the number of steps. However, we noticed

a significant difference in completion times of knowledgeable users, in particular only between prescribing guidance and the other degrees ($p \approx 0.02$, $chi - sq = 29.8$, $df = 2$, see Figure 3.8) when solving *domain tasks*. Same results were obtained for *exploratory tasks*: guidance affected completion times but not the number of total steps. Completion time was significantly better with prescribing guidance in respect to no guidance and directing guidance ($p \approx 0.002$ and $p \approx 0.02$ respectively, $chi - sq = 15.7$, $df = 2$). These results are in line with our assumption that knowledgeable users may still benefit from guidance. The tests reveal reduced completion times for these users: the guidance allows them to focus on the supervision of the analysis, alleviating the burden of focusing on minor details.

Besides completion times, high guidance also had significantly positive effects on correctness, distance values, and mental state of knowledgeable users. For knowledgeable users solving domain tasks the tests highlighted a significant difference between prescribing and no guidance. However, no difference was detected between prescribing and directing guidance. This may indicate that some knowledge may allow users to correctly interpret directing guidance. Hence, this degree should be considered when designing guidance for knowledgeable users, as it still leaves the users a certain degree of freedom, which has a positive impact on the mental state of the users (the participants commented that they do not feel restricted), and may lead them to discover the unexpected.

H2.1 We hypothesize that guidance may influence positively the confidence of participants and the tests showed that confidence levels were significantly higher with higher guidance. *Novice users* rated their confidence in their results significantly higher when receiving prescribing guidance, if compared to no guidance ($p \approx 0.002$, $chi - sq = 12.1$, $df = 2$, see Figure 3.9), for exploratory tasks. The same comparison is significantly different ($p \approx 0.006$, $chi - sq = 10.15$, $df = 2$) also for knowledgeable users solving the same task type. In this test, although the charts report increased confidence associated with the provision of directing guidance, the tests did not report any significant difference if we compare the confidence levels obtained with no guidance. For novice users solving domain tasks, the tests revealed a significant difference ($p \approx 0.000025$, $chi - sq = 21.1$, $df = 2$) between prescribing and the other two guidance degrees. A still significant, but lower result ($p \approx 0.02$, $chi - sq = 6.8$, $df = 2$) is reported also for knowledgeable users solving domain tasks with prescribing and with no guidance. Conversely to the results obtained with exploratory tasks, where we noticed an increased confidence with increased guidance, for domain tasks the confidence values obtained with directing guidance are absolutely comparable to those obtained with no guidance. It is clear the influence of a proper knowledge on those users. Comparing general confidence levels of novice with those of knowledgeable users, for exploratory tasks, the tests did not show any significant difference. For domain tasks, the results show however a significant difference. Furthermore, for these tasks, the confidence related to no guidance is comparable to confidence levels with medium guidance, for novice and knowledgeable users. Finally, when comparing *domain tasks* to *exploratory tasks*, we noticed that the average confidence resulted much lower for domain tasks than for exploratory tasks.

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

H2.2 We hypothesized that different guidance degrees influence how much novice and knowledgeable users feel frustrated when solving their tasks. In particular, we thought that novice users would be significantly less frustrated by a high degree of guidance than knowledgeable users.

Novice users Like stated by Celik et al. [Cel+13] the lack of knowledge is proportional to the users' frustration. Guidance, in this sense, represents a compensation for the lack of knowledge. Our results indicate that they feel less frustrated when receiving a higher degree of guidance, both for *exploratory and domain tasks*. Frustration decreases with increasing guidance, but in our tests it is prescribing guidance that marks a significant difference with the others degrees. In fact, prescribing guidance significantly reduces frustration in novice users compared to no or directing guidance. Figure 3.10 shows a box-plot representing the level of frustration with respect to the provided guidance. In the figure, the results represent both *exploratory* and *domain tasks*. The total average frustration for exploratory tasks (avg: 1.65) and domain tasks (avg: 1.55) is comparable. If we consider individually the single task types, we notice that *domain tasks* evoke frustration in novice users: the tests indicate a significant difference between prescribing guidance and no guidance ($p \approx 0.001$, $chi - sq = 25.6$, $df = 2$), while no significant difference is reported with directing and between directing and no guidance. For *exploratory tasks*, we also noticed a significant difference between prescribing and no guidance ($p \approx 0.02$, $chi - sq = 13$, $df = 2$). Although directing guidance reduced by the half the perception of frustration (avg:1.0) in respect to no guidance (avg: 2), the tests reported no significant differences between prescribing guidance and directing guidance as well as between directing guidance and no guidance.

Knowledgeable users We hypothesized that high degrees of guidance may cause frustration in *knowledgeable users* who already know how to conduct the analysis. Although we showed that the frustration of novice users decreases while the guidance degree is increased, the opposite is not true for knowledgeable users. The tests did not show an increased frustration in correlation with increased guidance (for both, exploratory and domain tasks). The test did not show any significant difference in frustration levels for different guidance degrees. Hence, H2.2 cannot be accepted.

For *domain specific tasks*, the frustration of knowledgeable users decreased to almost zero, both for prescribing and directing guidance. The increased knowledge enabled participants to correctly interpret the suggestions, producing low frustration levels also for directing guidance. For *exploratory tasks*, some participants reported increased frustration when receiving prescribing guidance, since they already knew how to interact with the visualization. However, this was mentioned by a small number of participants, and did not affect the overall results. Knowledgeable users show lower levels of frustration as novice users, also when provided with prescribing guidance. Some of them reported that they felt frustrated by the restrictions that come with this high degree of guidance. Moreover, knowledgeable users rated the tasks easier to solve when receiving prescribing guidance.

In summary, frustration is related to the inability of users to complete the tasks. The tests suggest that prescribing guidance reduces significantly the frustration of novice users. Moreover, our results show that domain tasks were more stressful than exploratory tasks, for novice users. Besides frustration levels, prescribing guidance also had significantly positive effects on all variables related to their mental state.

3.7.4 Observations and Interpretation of the Results

In this section, we present and discuss findings and observations that are not directly connected to the hypotheses we formulated beforehand, but were apparent from our results and that are worth mentioning.

Effects of Knowledge We observed that providing knowledge to participants had a significant influence on their performance and mental state when solving *domain tasks*, but not so pronounced in participants solving *exploratory tasks*. This may be explained by the fact that participants, without knowing what they were actually doing, but only knowing how to do it (exploratory tasks with operational knowledge), still did not feel like having control. However, receiving high guidance mitigated these strong differences between novice and knowledgeable users. This means that domain knowledge has a significant impact on performance and mental state, but a high degree of guidance could also have the potential to compensate for a lack of domain knowledge.

On the other hand, domain knowledge may also compensate for missing guidance. After the learning session, the participants did not feel lost when receiving no guidance. Furthermore, they felt that even directing guidance made solving the tasks as easy as when receiving prescribing guidance i.e., the tests did not show any significant differences between the two groups.

Difference between Domain and Exploratory Tasks A positive effect of guidance may also be found when comparing mental states of participants dealing with domain or exploratory tasks, respectively. Regardless of the degree of guidance provided, knowledgeable users felt significantly more lost, more frustrated, less confident, and thought that the tasks were harder to solve when solving exploratory tasks than when solving domain tasks. Novice users, on the other hand, felt significantly more lost and less confident when dealing with domain tasks. Furthermore, novice users found domain tasks harder to solve than exploratory tasks. We reason that this is due to the fact that having no knowledge about how to interact with the visualization may be compensated with trial and error, but having operational knowledge did not enable participants to control what they are doing semantically. Missing essential domain knowledge to solve a given task cannot be easily compensated by trial and error.

Confidence, Correctness, and Frustration Another finding from this study is that participants' confidence in their answers was justified. We observed that their confidence levels correlated with correctness levels. Furthermore, a negative correlation can be found

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

between correctness and frustration level. While these findings are not surprising, they foster our trust in the reliability of our results.

Misleading Hints In a handful of cases, providing directing guidance resulted in even worse performance (times and correctness) than providing no guidance at all. However, the tests did not show significant differences here. It may be that in some cases novice users trusted the hints provided by directing guidance too much. In fact, some participants selected all data points within the highlighted regions of interest, without reasoning about their effective meaning. This means that a vague kind of guidance – e.g., providing recommendations to the user, etc. – should be used with caution because it may mislead novice users. For expert users, in fact, this behavior could not be observed.

Appropriateness of Guidance Another interesting finding is that when participants felt they received an appropriate degree of guidance they also completed the tasks with a positive outcome in all other variables: They had better performance in terms of time and correctness, they felt less lost, less frustrated, and more confident about their answer. Finally, they had also the impression that the task was easier to solve. While this outcome was to be expected, it stresses the importance of providing an appropriate degree of guidance with respect to the expertise of the user. While novice users considered prescribing guidance to be much more appropriate than the other two degrees, this effect was mitigated for knowledgeable users. For real expert users this may be even more true and too high degrees of guidance could lead to frustration.

3.8 Discussion and Future Work

Although we considered carefully each and every design aspect of our study, there are also limitations to this work.

Knowledge One of our main concerns, while designing the study, was how to ensure different knowledge levels of participants. Usually, knowledge is hard to judge, evaluate, or measure precisely, as there are many factors influencing the way it is acquired. However, we did our best to ensure that our novice users had no additional information to solve the tasks – they had, in fact, no experience with the data. On the other hand, we taught our knowledgeable users what they needed to know and some learning effect from the first session also added up to that knowledge. To further consolidate the acquired knowledge, the participants had to exercise, and revisit the concepts before proceeding with the remaining tests. However, as already mentioned, we did not measure precisely the increase of knowledge, but the results of the study show clearly that knowledgeable users had better performance than the novice in the *no guidance* condition. Around 10% more participants was able to solve exploratory tasks after the learning session, and an average of 20% more participants could solve the domain tasks after learning the required domain knowledge.

Key Findings

1. While it is no surprise that a high degree of guidance had positive effects on the performance of novice users, it is remarkable that guidance, especially the prescribing degree, had significant positive effects on performance and mental state also of knowledgeable users for almost all combinations of task types (H1).
 2. Guidance was particularly effective to account for the lack of operational knowledge. For domain tasks, the users should possess at least a minimum of knowledge to interpret correctly the suggestions. This indicates that missing operational knowledge is easier to compensate by guidance than missing domain knowledge (H1).
 3. Knowledgeable users were not frustrated by high degrees of guidance while there was a positive effect on confidence and the subjective assessment of the difficulty of the task (H2).
 4. Participants' subjective assessment of appropriateness of guidance degree was reflected in better performance, and more positive mental state, which reflects the importance of providing an appropriate degree of guidance for the given user (H2).
 5. Knowledge plays an important role for positive performance and mental state especially when solving domain tasks. However, prescribing guidance may compensate for the lack of knowledge in many aspects (additional finding).
 6. Knowledge may also compensate for a lack of guidance. Knowledgeable users with no guidance obtained similar performances to novice users provided with directing guidance, for both exploratory and domain tasks (additional finding).
 7. Domain tasks evoked more frustration than exploratory tasks in novice users, since trial and error can compensate for a lack of operational knowledge while not for a lack of domain knowledge (additional finding).
-

Table 3.3: Key findings. In this table we summarize the results obtained in our study. We provide references to the hypotheses where we discussed these results in more detail. *Additional finding* refers to results that were not taken directly from the hypotheses, but inferred from them. Please refer to Section 3.7.4 for the details.

Since our study participants were familiar with standard interaction techniques, we had to design exploratory tasks with less obvious interaction techniques. This was backed up by the interaction logs which showed that usually just simple clicks were used by novice users. Just a few of them used other interaction means, and many reported that they learned about all the different interaction options just during the learning phase. We found significant differences in novice users and knowledgeable users in terms of mental state and performance, which furthermore confirms the distinction of their knowledge levels.

Knowledgeable Users Another limitation is presented by the fact that our knowledgeable users cannot be considered real experts yet. A real expert would be someone who was working in the given domain and with the provided interactive visualizations for a long time. Our study design did not include this type of user. The used visualizations

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

were specifically designed to include different degrees of guidance into different basic types of visualization to be able to test our hypotheses, and there is just no real life scenario with real expert users that would be suited to test these hypotheses. Thus, our results reflect only the behaviour of novice users contrasted with the behaviour of knowledgeable users, who both benefit the guidance received. However, we see a tendency of knowledgeable users to feel frustrated by prescribing guidance when they felt that the tasks were very easy. We can only hypothesize that real expert users may have found prescribing guidance disturbing or restricting, but this is left for further investigation.

Directing Guidance Another interesting point regards the representation of the hints provided by directing guidance. We did our best to visually distinguish the hints given by directing guidance and the actual instructions given by prescribing guidance in order to not mislead users. We chose simple highlighting to indicate interaction options and interesting data regions (directing guidance), while for prescribing guidance we chose precise textual instructions in combination with highlighting specific data items. We consistently used this encoding in all tasks and visualization types, and furthermore, informed participants about these differences. We wanted to provide guidance in a way as general as possible, and we found that this simple representation was quite useful and effective for our scope. However, other encodings would also be feasible and may lead to different results. Thus, it would be interesting to investigate further encodings of guidance and their effects in future work.

A curious outcome of our study was that often our tests did not report *significant* differences between directing guidance and no guidance. From the box-plots (see for instance Figures 3.6-3.9) it is clear that guidance introduced some differences, but the test did not highlight it as significant. We tried to explain this situation reasoning about the fact that in some situations, like at the beginning of the test, the novice users had not sufficient means to understand the guidance hints. However, in some other cases, we could not fully explain this result. In particular, it was unexpected encountering this lack of significant differences in knowledgeable users. We could think that in those cases the acquired knowledge was then sufficient to fill the differences between the two guidance degrees. The cause may also be related to the possibly misleading hints given by this guidance degree. However, neither the logs nor the participants' comments gave us a better understanding of the real cause. For this reason, we reserve the possibility of further investigations in this direction.

Visual Encoding In line with ensuring basic visual encodings of guidance, we also chose a small number of basic visualizations to represent the dataset. The chosen visualizations represent pretty standard choices in visual data analysis, and are well suited to solve the tasks we proposed to the participants. Another motivation for choosing them was that the participants were already familiar with them. The participants, in fact, reported that in the vast majority of the cases the visualizations were well understood.

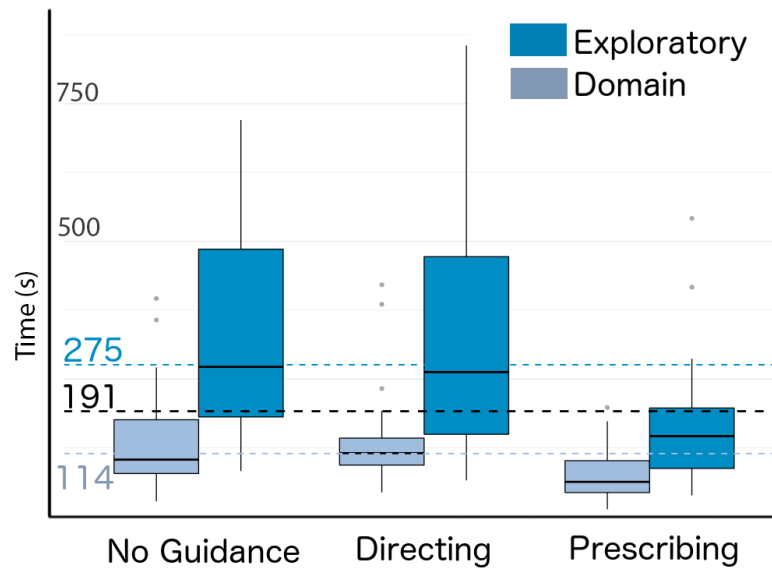
However, a consequence of our choice, is that our findings may not be generalized to more sophisticated methods. This again would be an interesting topic for future investigations.

Tasks Finally, we constructed our study on a limited number of different tasks. In particular, we focused these tasks on some specific domain concepts, and simple operational procedures. Although we designed them with special respect to keeping them simple (i.e., basic look-up tasks, simple interactions, so to simulate a general exploratory analysis) we cannot guarantee the results we obtained may be generalized to other types of tasks. Hence, our results should be seen as initial insights, how guidance works for these and similar tasks, how different guidance degrees work for different users, possible effects on a user's mental state and critical aspects that need to be considered. However, we think that these results could be extended and consolidated for other tasks and domains.


3.9 Conclusion

We presented a user study about guidance, which constitutes a first step towards a scientific understanding of the effects of guidance in different analysis scenarios. In this context, we consider a number of different aspects that interact with guidance. We relate the effects of different degrees of guidance to a user's expertise level, we consider different types of tasks, and we measure task performance as well as the user's mental state. Our study suggests that guidance has positive effects on both knowledgeable and novice users. On the other hand, the study reveals that guidance must be designed carefully to meet the user's needs and that novice users may also be misled by medium guidance. We conclude that our work describes the value and effectiveness of having guidance while conducting a visual data analysis.

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE

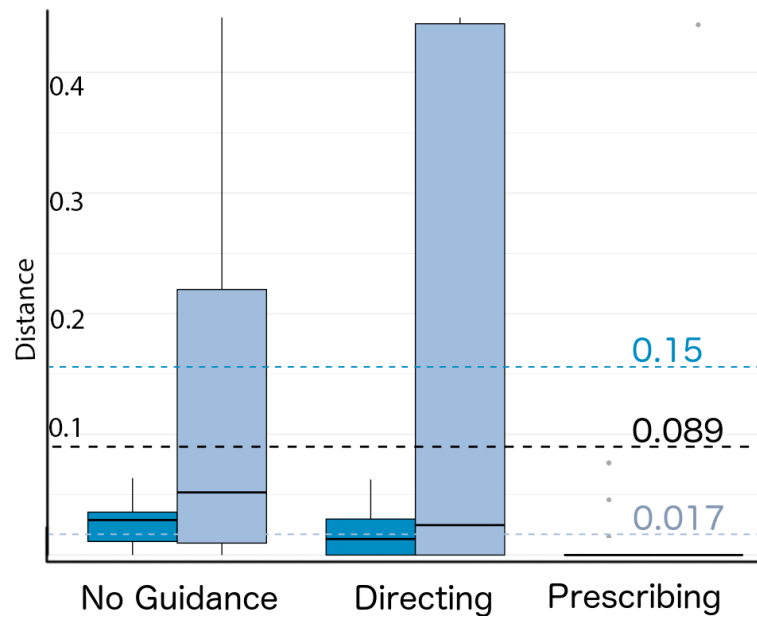


(a) Completion time (seconds) of novice users.




(b) Correctness metric of novice users.

Figure 3.6: Box-plots for H1.1: We report time and correctness performance metrics for novice users (blue tones), according to different guidance degrees (x-axis). Dashed lines encode the total average (exploratory and domain tasks combined; in black) and individual average values for exploratory and domain tasks (in light and darker blue).



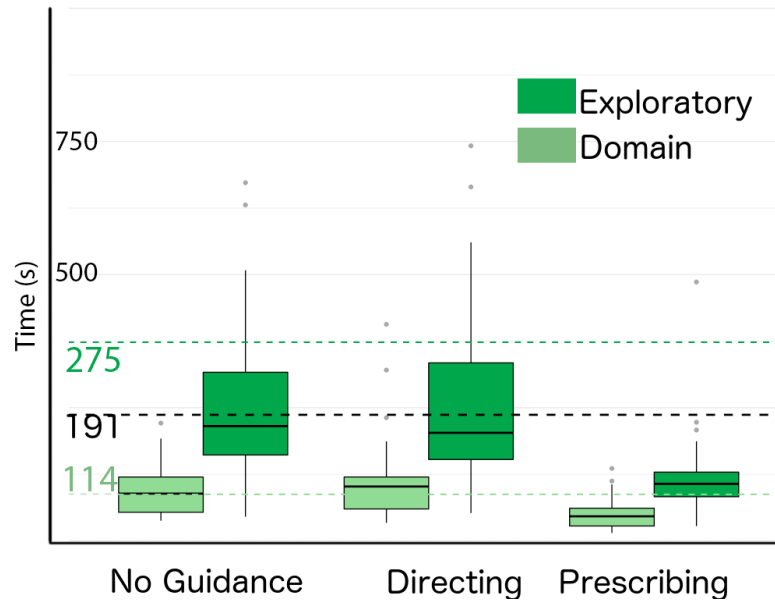
(a) Distance metric of novice users.




(b) Number of steps/operations of novice users.

Figure 3.7: Box-plots for H1.1: We report distance metric and the number of steps for novice users (blue tones), according to different guidance degrees (x-axis). Dashed lines encode the total average (exploratory and domain tasks combined; in black) and individual average values for exploratory and domain tasks (in light and darker blue).

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE




(a) Completion time (seconds) of knowledgeable users.




(b) Number of steps/operations of knowledgeable users.

Figure 3.8: Box-plots for H1.2: completion time and total steps of knowledgeable users (green tones). Dashed lines encode the total average (exploratory and domain tasks combined; in black) and individual average values for exploratory and domain tasks (in light and darker green).




(a) Confidence of novice users.



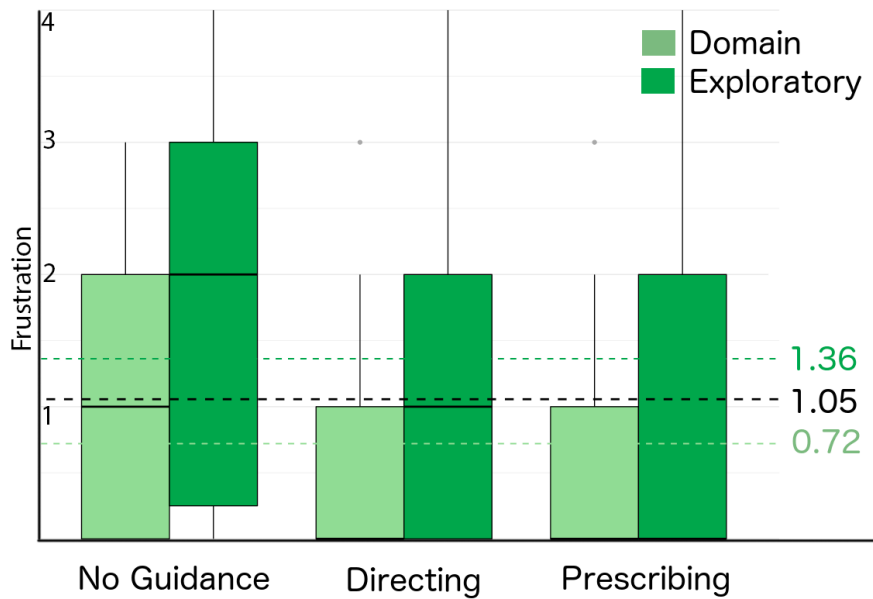
(b) Confidence of knowledgeable users.

Figure 3.9: Box-plots for H2.1. Guidance influences positively the user's confidence. Confidence was measured on a five-point Likert scale, where 0 encodes no confidence. Novice users are represented with blue tones, and knowledgeable users with green tones. Dashed lines encode the total average (exploratory and domain tasks combined, in black) and individual average values for exploratory and domain tasks.

3. YOU GET BY WITH A LITTLE HELP: THE EFFECTS OF VARIABLE GUIDANCE DEGREES ON PERFORMANCE AND MENTAL STATE



(a) Frustration of novice users.



(b) Frustration of knowledgeable users.

Figure 3.10: Box-plots for H2.2. Frustration of novice users (blue tones) and knowledgeable users (green tones), according to different guidance degrees. The level of frustration is ascending: values closer to 0 indicate less frustration. We also report, with dashed lines, the total average frustration (for exploratory and domain tasks combined, in black), as well as the average value for individual domain and exploratory tasks.

Bibliography

- [AHR13] W. Aigner, S. Hoffmann, and A. Rind. „EvalBench: a software library for visualization evaluation“. In: *Computer Graphics Forum*. Vol. 32. Wiley Online Library. 2013, pp. 41–50.
- [BH95] Y. Benjamini and Y. Hochberg. „Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing“. In: *Journal of the Royal Statistical Society. Series B (Methodological)* 57.1 (1995), pp. 289–300. ISSN: 00359246.
- [BL09] E. Bertini and D. Lalanne. „Surveying the complementary role of automatic data analysis and visualization in knowledge discovery“. In: *Proc. of the Workshop on Visual Analytics and Knowledge Discovery: Integrating Automated Analysis with Interactive Exploration*. ACM. 2009, pp. 12–20.
- [BM07] P. Brusilovsky and E. Millán. „User models for adaptive hypermedia and adaptive educational systems“. In: *The Adaptive Web*. Springer, 2007, pp. 3–53.
- [BPH05] A. Bernstein, F. Provost, and S. Hill. „Toward intelligent assistance for a data mining process: An ontology-based approach for cost-sensitive classification“. In: *IEEE Transactions on Knowledge and Data Engineering* 17.4 (2005), pp. 503–518. DOI: 10.1109/TKDE.2005.67.
- [BS03] B. Bederson and B. Shneiderman. *The craft of information visualization: readings and reflections*. Morgan Kaufmann, 2003.
- [Cel+13] P. Celik et al. „Not all rejections are alike; competence and warmth as a fundamental distinction in social rejection“. In: *Journal of Experimental Social Psychology* 49.4 (2013), pp. 635–642.
- [Cen+17] D. Ceneda et al. „Characterizing Guidance in Visual Analytics“. In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (Jan. 2017), pp. 111–120. ISSN: 1077-2626. DOI: 10.1109/TVCG.2016.2598468.
- [Cen+18] D. Ceneda et al. „Guidance or No Guidance? A Decision Tree Can Help“. In: *Proc. of the International Workshop on Visual Analytics (EuroVA)*. Eurographics Digital Library, 2018, 19–23. DOI: 10.2312/eurova.20181107.

- [CGM19a] D. Ceneda, T. Gschwandtner, and S. Miksch. „A review of guidance approaches in visual data analysis: A multifocal perspective“. In: *Computer Graphics Forum* 38.3 (2019), pp. 861–879. DOI: 10.1111/cgf.13730.
- [CGM19b] D. Ceneda, T. Gschwandtner, and S. Miksch. „You get by with a little help: The effects of variable guidance degrees on performance and mental state“. In: *Visual Informatics* 3.4 (2019), pp. 177–191. ISSN: 2468-502X.
- [Che05] C. Chen. „Top 10 unsolved information visualization problems“. In: *Proc. of IEEE Computer Graphics and Applications* 25.4 (2005), pp. 12–16.
- [CS16] J. Cloern and T. Schraga. „USGS Measurements of Water Quality in San Francisco Bay (CA), 1969-2015“. In: *Sci Data* (2016). DOI: <https://doi.org/10.5066/F7TQ5ZPR>.
- [DD10] J. C. De Winter and D. Dodou. „Five-point Likert items: t test versus Mann-Whitney-Wilcoxon“. In: *Practical Assessment, Research & Evaluation* 15.11 (2010), pp. 1–12.
- [Dix+04] A. Dix et al. *Human-Computer Interaction*. 3rd. Pearson Education, 2004.
- [Gib77] J. Gibson. „The theory of affordances“. In: *Perceiving, Acting, and Knowing* (1977), pp. 67–82.
- [GW09] D. Gotz and Z. Wen. „Behavior-driven visualization recommendation“. In: *Proc. of Conference on Intelligent User Interfaces*. ACM, 2009, pp. 315–324. DOI: 10.1145/1502650.1502695.
- [HCL05] J. Heer, S. K. Card, and J. A. Landay. „Prefuse: a toolkit for interactive information visualization“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM. 2005, pp. 421–430.
- [Hor99] E. Horvitz. „Principles of mixed-initiative user interfaces“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM. 1999, pp. 159–166.
- [KBP07] A. Kapoor, W. Bursleson, and R. W. Picard. „Automatic prediction of frustration“. In: *International Journal of Human-computer Studies* 65.8 (2007), pp. 724–736.
- [Kei+08] D. A. Keim et al. „Visual analytics: Scope and challenges“. In: *Visual Data Mining*. Springer, 2008, pp. 76–90.
- [KW52] W. H. Kruskal and W. A. Wallis. „Use of ranks in one-criterion variance analysis“. In: *Journal of the American Statistical Association* 47.260 (1952), pp. 583–621.
- [Maz+10] M. A. Mazurowski et al. „Individualized computer-aided education in mammography based on user modeling: Concept and preliminary experiments“. In: *Medical Physics* 37.3 (2010), pp. 1152–1160.
- [MW47] H. B. Mann and D. R. Whitney. „On a test of whether one of two random variables is stochastically larger than the other“. In: *The annals of mathematical statistics* (1947), pp. 50–60.

- [Rin+13] A. Rind et al. „Timebench: A data model and software library for visual analytics of time-oriented data“. In: *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), pp. 2247–2256.
- [Sac+16] D. Sacha et al. „The role of uncertainty, awareness, and trust in visual analytics“. In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), pp. 240–249.
- [SM86] S. Smith and J. Mosier. *Guidelines for Designing User Interface Software*. Tech. rep. ESD-TR-86-278. Mitre Corporation, Bedford MA, 1986.
- [Str+12] M. Streit et al. „Model-driven design for the visual analysis of heterogeneous data“. In: *IEEE Transactions on Visualization and Computer Graphics* 18.6 (2012), pp. 998–1010. DOI: 10.1109/TVCG.2011.108.
- [TC05] J. J. Thomas and K. A. Cook. *Illuminating the path: The research and development agenda for visual analytics*. Tech. rep. Pacific Northwest National Lab.(PNNL), Richland, WA (United States), 2005.
- [Tea14] R. C. Team. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. 2013. 2014.
- [WHA07] W. Willett, J. Heer, and M. Agrawala. „Scented widgets: Improving navigation cues with embedded visualizations“. In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (2007), pp. 1129–1136. DOI: 10.1109/TVCG.2007.70589.
- [Won+16] K. Wongsuphasawat et al. „Voyager: Exploratory Analysis via Faceted Browsing of Visualization Recommendations“. In: *IEEE Transactions on Visualization and Computer Graphics* 22.1 (2016), pp. 649–658. DOI: 10.1109/TVCG.2015.2467191.



Die approbierte gedruckte Originalversion dieser Dissertation ist an der TU Wien Bibliothek verfügbar.
The approved original version of this doctoral thesis is available in print at TU Wien Bibliothek.

Guide Me in Analysis: A Framework for Guidance Designers

This paper appeared in: Computer Graphics Forum 3(4); 177-191, 2020. ©2020 Wiley and sons. Reused with Permission (See [Cen+20])

Full List of Authors Besides Davide Ceneda, the following authors contributed and provided valuable input to this paper: , Natalia Andrienko, Gennady Andrienko, Theresia Gschwandtner, Silvia Miksch, Nikolaus Piccolotto, Tobias Schreck, Marc Streit, Josef Suschnigg and Christian Tominski.

4.1 Abstract

Guidance is an emerging topic in the field of Visual Analytics. Guidance can support users in pursuing their analytical goals more efficiently and help in making the analysis successful. However, it is not clear how guidance approaches should be designed and what specific factors should be considered for effective support. In this paper, we approach this problem from the perspective of guidance designers. We present a framework comprising requirements and a set of specific phases designers should go through when designing guidance for visual analytics. We relate this process with a set of quality criteria we aim to support with our framework, that are necessary for obtaining a suitable and effective guidance solution. To demonstrate the practical usability of our methodology, we apply our framework to the design of guidance in three analysis scenarios and a design walk-through session. Moreover, we list the emerging challenges and report how the framework can be used to design guidance solutions that mitigate these issues.

4.2 Introduction

Visual Analytics (VA) approaches can be effective tools for making sense of large datasets and perform complex tasks. Their strengths come from a tight integration of automated analysis methods and visual interactive interfaces [TC05]. In recent years, many VA approaches have been proposed to solve data analysis problems in a wide set of scenarios. However, usually benefits come at a price: Automated analysis methods and visualization techniques need to be configured and meaningful parameters need to be set to obtain high-quality results. Despite the development of guidelines and the adoption of well-established design patterns [Dix+04; Die+18], using interactive interfaces may present many challenges to analysts.

Given these premises, the research community started to develop approaches and techniques to support data analysts during the analysis process. These are known as guidance [Cen+17]. The main aim of guidance is to ease problematic situations and mitigate issues that might hinder the analyst from achieving results, generating insights, and in the end producing new knowledge. Recent studies show that including guidance in the analysis may be beneficial for the user [Cen+18b; May+12; HB05; Joh+05; CGM19; Col+18]. Interest in guidance is quite recent and research has just started scratching the surface of this field. Previous work on guidance, in fact, explores and describes just the characteristics of the guidance process [Cen+17; Col+18]. Only little research exists describing general procedures to implement guidance in practical scenarios.

In this paper, we provide an initial steppingstone to close the aforementioned gap by reasoning about the process of designing effective guidance in VA. In contrast to previous works, we study the problem from the perspective of designers. We describe a framework comprising a list of steps for guidance designers and a set of qualitative requirements to guide the whole design process.

Guidance is a context-dependent process. Therefore, it is hardly possible to create an algorithm for guidance design consisting of concrete instructions while being applicable to any analysis scenario. Instead, we provide a framework that points designers to important considerations in the context of guidance design and guides them through this process in a step-wise manner. We complement the framework with a set of design requirements that should be satisfied to make the design effective and obtain a user-tailored solution.

To demonstrate the applicability of our general methodology, we describe three design examples and a comprehensive design walk-through by a VA expert who was not involved in the development of the framework. The first example is about guiding users in the exploration of cyclic patterns in univariate time-series data. The other two examples are set in the application domains of engine testing and financial fraud detection.

In summary, the value of this work comes not only from the development of a general framework, but also from the discussion of threats, risks, and possible countermeasures thereof that could arise during the design. Our contribution is thus threefold:

- We provide a general framework comprising a step-wise procedure for designers aiming at designing *effective* guidance in VA approaches.
- We describe possible countermeasures to risks and threats that could arise during the design process and support an effective implementation of design requirements.
- We demonstrate the value of our framework by applying it to the design of guidance for VA. In this context, we describe challenges and combine them with an appropriate design of guidance solutions, thus showing the applicability of our framework.

4.3 Related Work

This work is mainly focused on guidance in VA [CGM19]. The research of guidance has quite a long story. Its roots are in Human-computer Interaction and mixed-initiative visual data analysis [Sil91; Hor99]. Guidance in VA was born from the need to assist and support users during interactive analytical work. It has been defined as "*a computer-assisted process that aims to actively resolve a knowledge gap encountered by users during an interactive visual analytics session*" [Cen+17, p.2]. This definition contains three key aspects: First, guidance is a continuous effort that runs alongside the regular VA activities. Second, guidance addresses a knowledge gap, which captures the discrepancy between what needs to be known to make analytical progress and what is actually known by the user, such as which visual/analytical methods to use, how to set parameters, or how to explore and get insights from the data. Third, guidance is not static, but it reacts to a dynamically changing interactive analysis session.

If done properly, guidance can support the VA process in different regards. Guidance can help to inform, mitigate bias, reduce cognitive load, and it can be beneficial for training, engagement and verification [Col+18].

In the past, a number of characteristics of guidance approaches have been identified [Cen+17; Col+18; Sch+13]. These characteristics primarily cover aspects of why guidance can be provided and how it should be enacted (e.g., the degree of guidance, the input based on which guidance is generated, and the way it is communicated). In this regard, the concept of *knowledge gap* acquires the most important role in designing and implementing guidance methods:

In fact, many issues, i.e., knowledge gaps, may arise for the user during the whole analysis process. Likewise, multiple kinds of guidance can be envisioned. To support the user in solving such knowledge gaps, the design of guidance certainly includes the choice of an appropriate user interface, but is not limited to it [Dix+04]. It also includes the design of an intelligent, knowledge-based system which possibly encompasses the creation of a knowledge base and a reasoning mechanism, determining what kind of knowledge should be provided to fill the gap and let the user continue the analysis. The user interface design concerns the way *how* to provide the necessary knowledge to the user, whereas the intelligent system design focuses on *what* and *when* to provide guidance to the user.

There are several specific examples where guidance has been applied successfully to assist users [GST13; Lub+12; May+12; Str+12; Str+14; Cen+18b]. For instance, Kandel et al. [Kan+11] designed an approach that guides the user towards the selection of appropriate data transformations based on the type of data under analysis. Gotz and Wen [GW09] developed a behavior-driven approach that supports the analyst in selecting the most appropriate visualization for a given analytical task. Bernard et al. [Ber+17] provide guidance to the process of labeling human motion data through the use of unsupervised algorithms. Gladisch et al. [GST13] supports the exploration of hierarchical graphs by using a flexible degree-of-interest function. May et al. [May+12] guides the user towards interesting regions of large graphs.

Guidance can be done in many ways and for different scopes. Among the possibilities, often user feedback (either implicitly, or explicitly) is considered. Also, guidance can be informed heuristically, based on data and view quality measures. To date, many quality measures have been introduced. For a review, please see [Ber11; Beh+18].

In summary, the literature seems to be in agreement on the potential benefits and the general requirements of guidance. However, there is no unified framework for designing effective guidance in VA. Based on a thorough inspection of existing work [CGM19], our research aims to narrow this gap in the literature by proposing a design framework built around questions that a guidance designer has to address when developing guidance for VA.

4.4 Motivating Example

To further motivate the need of a framework for guidance designers, let us consider an example of analyzing multiple time series describing a football (soccer) game [And+17]. At first, we describe briefly the scenario and its challenges, focusing on data, users (i.e., analysts), and tasks [MA14]. Then we change perspective and discuss the example from the point of view of designers by listing factors to be considered when designing guidance.

The dataset under consideration already offers challenges originating from the relations among the multiple variables and the temporal axis. We list various time-variant attributes derived from the trajectories of the players and the ball (see Figure 4.1):

- Attributes of the players: speed, movement direction, distances to the two goals, distance to the centre of the own team, etc.
- Attributes of the ball: in play or out of play, which team possesses it, speed, direction, distances to the goals, etc.
- Attributes of the teams: dimensions along and across the pitch, distances from the centre to the goals, mean distance between the team members, mean distance to the nearest opponent player, area of the intersection with the opponent team, etc.

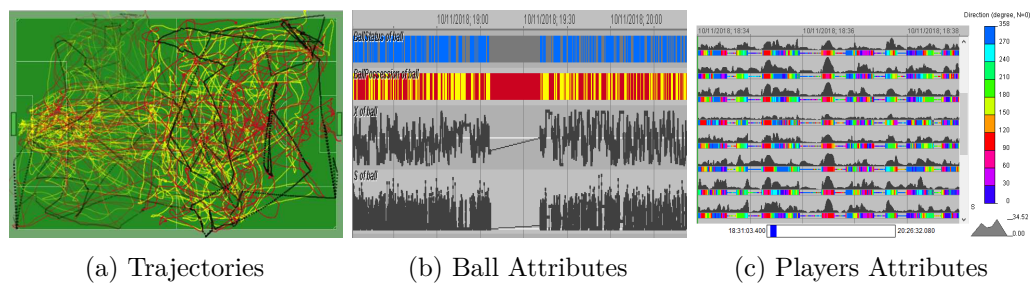


Figure 4.1: Analysis of Soccer Matches [And+17]. We motivate the need of a proper guidance design: (a) Fragments of trajectories of the players and of the ball for a selected time interval of 5 minutes length. Finding coordinated movement behaviours by looking at trajectories is complicated due to the overcrowded visualization. (b) Time series representing ball attributes, for instance, ball status (in play or out of play), ball possession (which team), X-coordinate, and speed. Due to the high frequency of the data, it is difficult to spot patterns at first sights. (c) Time series of attributes of the players. The need to consider multiple attributes makes the discovery of behaviors and patterns complicated.

We consider analysts investigating the tactics of the teams in terms of coordination between the movements of the team players. This task consists of detecting correlations between multiple time series of the players-related attributes. The specific challenge is that different patterns of coordination may be used in different kinds of situations. For example, after a team gains the ball in the central part of the pitch, the adopted tactics may be expansion of the team to the sides of the pitch, but the team may behave quite differently after gaining the ball close to the own goal or close to the opponents' goal. Hence, the analysis requires extraction of subsets of situations with particular characteristics (regarding the attributes of the ball and the teams) and applying correlation analysis to these subsets.

We assume that analysts have sufficient domain knowledge to select the relevant attributes to investigate specific tactics. However, we also suppose that the analysts may not know the following things.

- Data selection: how to set query conditions to select situations with particular characteristics?
- Methods: what is the class of techniques that can be helpful to detect and analyze coordinated behaviours? What specific methods are suitable and how to set their parameters?
- Interpretation: how to interpret the results of the methods, i.e., translate numbers into concepts? How to see and explore the coordinated behaviours corresponding to these results?

- Evaluation and validation: how to assess the coherence of a pattern derived from a group of situations across the individual situations (i.e., how much variance exists)?
- Comparison: how to compare the patterns derived for different groups of situations or for different teams?

Multiple issues might occur during the analysis. From the perspective of designers, these are the possible knowledge gaps that need to be anticipated and addressed when designing the guidance. The benefits of an effective design resides in a positive solution of such gaps and, as a consequence, in an easier time for the analysts using VA tools. Therefore, the main question arises:

How do we design *effective* guidance [CGM19] to support a positive analysis outcome? What questions and what criteria should guide the development of an effective guidance solution?

To provide a satisfactory answer to these questions, we start by making a couple of considerations about the given analysis scenario (see Section 4.5.1).

The above list suggests that different knowledge gaps may affect the analysis at different moments. Analysts wishing to analyze the data should be aware of such issues and know how to overcome them. If this is not the case, then analysts should at least be able to ask for, or rely on, guidance. Hence, we assume that a first important aspect designers should consider is how to incorporate mechanisms to detect problematic situations and let the analysts ask for guidance. The presence of biases also confirms that guidance needs to be tailored to the needs of specific users.

Sticking to the described scenario, considering the specifics of the data (i.e., multivariate time series) is especially relevant for supporting data selection and for suggesting the necessary methods. An important aspect guidance designers should consider are the issues related to the choice of appropriate parameters for the analytical methods involved. Analysts should be supported in setting conditions for time steps preceding or following a given time step (e.g., “team B possessed the ball in the previous time step”) and specify the minimal length of the time interval in which a given combination of conditions must hold (e.g., “the ball possession duration of team A must be at least 5 seconds”). Guidance may support this process by proposing appropriate parameter settings. However, analysts might not be fully satisfied with them. Hence a further aspect a guidance designer should consider is the provision of proper means to steer the guidance process in the eventuality that the guidance provided does not satisfy completely the needs of analysts.

Finally, as designers, we could expect the user to be a domain expert, which implies the use of appropriate visual means to visualize and conduct the analysis. For instance, the use of specific means to visualize the field, the trajectories of the ball and the players, without interfering with the analysis dynamics. In this context, another important aspect

to address is how to encode the guidance in the analysis process without distracting the user and disrupting the analysis flow.

Analyzing this example shows that multiple factors are to be considered when designing guidance. So far, we listed just some of them that should be taken into account when designing guidance for the analysis of time-varying soccer data. In the following, we formalize, abstract, and complement these aspects in a general design framework. At first, we describe a list of requirements that should guide designers when dealing with guidance. We then illustrate a step-wise procedure to design guidance supporting these requirements.

4.5 A Framework for Guidance Designers

Guidance is described as a closed loop [Cen+18a]. On the one hand, the system provides possible solutions to mitigate problems arising during the analysis (i.e., guidance for the analyst). On the other hand, the analyst guides the system by steering the process along the desired analysis path (i.e., guidance for the system). Our framework considers both sides of this process.

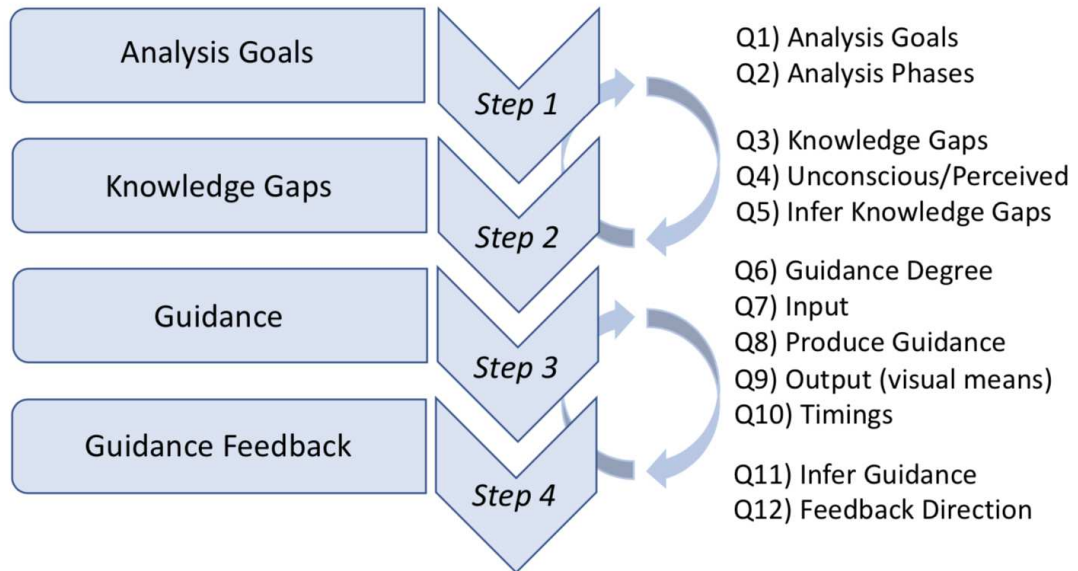
Methodology Before introducing the framework, we want to provide a brief overview of the procedure we followed to derive it. The framework was built from multiple sources followed by an iterative refinement process. Initially, we performed a deep investigation of the literature. By performing an analysis of the literature on the topic, we developed an initial understanding of the factors that should be considered when dealing with the process of guiding visual analysis. As a result of this initial phase, we produced a raw list of steps and a description of their interdependencies. We later confronted such initial output with our experience in the field. Having already developed guidance approaches and dealt with similar challenges in the past helped us in refining the initial output into an ordered list of steps and design requirements.

4.5.1 Requirements

In general, our framework is aimed at facilitating the design of *effective* guidance. To support effectiveness, designers should check not only that the guidance is effective for accomplishing the analysis objectives, but that it is also communicated to the analyst at the right moment by appropriate means. Therefore, in the following, we list a set of requirements we think should be met during the design process, to support effectiveness, and make the guidance accepted and user-tailored. Designers are usually confronted with multiple design alternatives during this process. *We believe that the identified requirements could serve as a guideline to choose among such alternatives.*

We based the requirements on an initial discussion by Ceneda et al. [CGM19], who identified a set of characteristics that should concur to the goal of effective guidance. We

Guidance Design Framework



Effective Guidance




Figure 4.2: The guidance design framework. The framework aims to support the design of effective guidance (R0, see Section 4.5.1). We list a set of steps (Step 1–4) as well as quality criteria (R1–R5) that should guide designers during the design process. The arrows going back and forth between the steps illustrate the iterative nature of the design.

complemented and further elaborated these characteristics to derive the following design requirements.

R0 - **Effective** guidance is what we see as the end goal of the design process. The designed guidance should be effective for a given task and a given user. To obtain such result, a number of requirements have to be supported:

R1 - **Available**: *Guidance is there for you* - Users should be aware that guidance is available and that support can be provided or requested at any time. Designers

should make available interactive means to request guidance and appropriate visual means to convey guidance.

- R2 - **Trustworthy**: *Guidance will help you* - Any generic data analysis task includes a certain degree of variability. Guidance should be regarded as a support to overcome the uncertainty involved and not being a source of further confusion. Designers should take care of specific ways to encode and provide guidance to make it trustworthy and accepted by users, and not an additional source of misinformation. Trust, once lost, is hard to restore.
- R3 - **Adaptive**: *Guidance will adapt to the situation* - Usually, as the analysis evolves so do the problems users encounter. The guidance system must know what the actual state of the analysis is, in order to deal with dynamically changing knowledge gaps. Designers should implement mechanisms to capture the analysis phase, provide interfaces for inferring the knowledge gap and provide guidance accordingly.
- R4 - **Controllable**: *Guidance can be tuned if necessary* and the user needs to be in control of the analysis - Guidance is a mixed-initiative process [Hor99]. Therefore, the designed solution should enable users to steer the analysis, choose between alternative recommendations, turn off the guidance if not needed, or provide means to ask for assistance in the first place.
- R5 - **Non-disruptive**: *Guidance will not annoy or mislead you* - A final quality that we expect to be supported by the guidance process is that it should not disrupt the analysis flow and the analysts' mental map. The guidance should be provided without having users to exit their state of flow.

In the following, we introduce four major general steps that should be completed in order to come up with an appropriate design of system and user guidance (see Figure 4.2). To make the framework easier to understand, we complement the description of the individual steps with examples from our introductory soccer scenario and with a description of risks that might arise during the design, as well as possible countermeasures.

4.5.2 Step 1: Analysis Goals

When designers approach the problem of providing guidance, they should start with identifying the analysis goal. The following questions should be answered:

- Q1 - *What are the analysis goals?*
 Q2 - *In which analysis phases issues might occur?*

Different analysis goals may require different guidance solutions. In the case of the soccer match analysis (Section 4.4), the goal is to identify coordination patterns in the tactical movements of players.

The process of pursuing a given analysis goal can be divided into a sequence of analysis phases that the analyst needs to go through, e.g., explore the data, evaluate findings, document results, etc. Therefore, guidance designers should not only consider the analysis end goal, but also examine the different sub-tasks that the analyst has to deal with in order to reach the end goal. Designers should then implement strategies to infer such phases, which is crucial to design adaptive guidance (R3). Breaking down the analysis process into single atomic tasks will allow the design of guidance for isolated problems and compose them to solve more complex analysis tasks.

According to Andrienko et al. [And+18] the data analysis process is composed of data preparation, data analysis and model development, and evaluation. A further subdivision of the preparation phase includes: 1) understanding the data and 2) preprocessing the data before the analysis. A subdivision of the analysis phase includes: 3) exploring the data and 4) developing a model. Finally, the data model is 5) evaluated and tested against the work hypotheses. Guidance may be needed at each of these phases, as many issues might arise along them, as shown also by the motivating example (see list of knowledge gaps in Section 4.4).

Step 1 – Risks, Threats, and Countermeasures

Possible risks derived from a non-satisfactory execution of this design step are **overestimation**, **underestimation**, and **misunderstanding** of the analysis goals. This means that the designer may identify too many or too few activities/tasks/goals requiring guidance. Even worse, a wrong design can also be a consequence of human errors, for instance, designers misunderstanding the analysis goals. To mitigate the latter risk, adaptive guidance mechanisms can be devised. For a detailed discussion, see Section 4.5.3. In case of overestimation, the end user (i.e., the analyst) could be bothered by the excess of support, which may lead the analyst to ignore the provided guidance and nullify its benefits, thus going against R1. The threat deriving from underestimation instead is the design of insufficient guidance. The underestimation could lead to a lack of proper support for critical tasks. In general, underestimation and overestimation run against the general aim of designing trustworthy (R2) and non-disruptive (R5) guidance. In such situations, an effective strategy for supporting a proper design requires close collaboration with domain experts who could provide crucial information about how to structure the tasks and support the identification of analysis objectives. Furthermore, the implementation of means for the user to control (R4) and to fine-tune the guidance might be a viable alternative solution to counteract such threats.

4.5.3 Step 2: Knowledge Gap

After identifying the analysis goal and the different analysis phases, guidance designers need to understand possible knowledge gaps [Cen+17] arising in the completion of those goals.

At this stage, the following question has to be asked:

Q3 - What knowledge gaps might hinder the analyst from proceeding the analysis?

A knowledge gap refers to the lack of knowledge or information that makes it difficult to complete the analysis or a certain phase of it, as identified in the previous step.

Structure or Execution?

A first general distinction we should make as designers, is whether we think the analyst may need help to reach the analysis objective, or to define a sequence of operations to reach it in the first place. In other words, designers should reason whether the knowledge gap is a problem of *structure* or *execution* [Sil91].

In the first case, the knowledge gap relates to finding the correct operators (e.g., algorithms, visualizations) and a combination thereof, in order to obtain the desired results. At this regard, a viable solution may consist of listing the available operators, as well as informing the analyst about alternative options that might serve the same purpose.

In the second case, the knowledge gap is related to the execution of the conceived plan, for instance, a structured sequence of operators, as detailed before. This could include the choice of parameters for each step. The execution of a given sub-task is related to the decisions taken by the analyst in the previous analysis steps. Therefore, a designer has two alternatives: directly guide the execution of such steps, for instance, guide the choice of proper inputs and support the analysis of the obtained output, or opt for an informative solution: provide the analyst with all the important information about the input as well as details about the possible expected outcomes, and let the analyst take the decision. In the end, it is a matter of giving the analyst more or less freedom.

Types of Knowledge Gaps

Another way of reasoning about the knowledge gap is considering its type. Four different types of knowledge gaps, which may appear at any phase of the analysis, can be identified [Cen+17].

1) Data: the lack of knowledge about the data. This kind of problem generally affects the pre-processing phases. In our soccer example, we may see the analysts having problems with understanding the relationships between the variables (e.g., the ball position) and the time axis (see Figure 4.1b). A knowledge gap in the data domain may also affect other analysis phases, for instance, data exploration. In this case data issues may be related to identifying specific data cases that are helpful to validate hypotheses, for instance, finding data subsets that describe certain known tactics.

2) Tasks: the lack of procedural knowledge (i.e., what are the steps) to complete a sub-task or to reach the analysis goal. For instance, finding sequences of interactions

(i.e., selections, filtering) in relation to different features, like ball possession, shots, and foul events etc. with the goal of analyzing team behaviors.

3) VA Methods/Algorithms: the lack of knowledge about what visual and analytical methods to apply, what algorithms to choose, and how to set their parameters. For instance, the analyst might have problems to compose a visual summary of the soccer match, since stacking trajectories may lead to visual occlusion and clutter (see Figure 4.1a). Methods and algorithms are needed at multiple analysis phases, such as data-preprocessing or model building (e.g., setting the parameters for clustering the positions of different soccer players). Other issues might occur when selecting appropriate features during the segmentation of the time series, to abstract, for instance, the data into events of the soccer match (see Figure 4.1b).

4) Knowledge and Insights Management: the lack of skills in interpreting patterns or the organization of the knowledge itself. This knowledge gap is related to merging, interpreting, labeling, and managing the findings to generate insights and new knowledge. For instance, translating patterns perceived from the visualization or discovered by algorithms into the domain concepts.

At this stage, we have considered analysis phases and knowledge gaps thereof. Related questions a designer should consider are how to identify knowledge gaps and understand if the analyst perceives them during the analysis. These considerations might in the end lead to completely different guidance solutions.

Perceived and Unconscious Knowledge Gaps

Having identified the knowledge gaps, the next question to be asked is the following:

Q4 - Are analysts aware or unaware of their knowledge gaps?

As designers, we should think whether the knowledge gap is perceived by the analyst. A perceived knowledge gap is one that analysts are aware of. The opposite is considered an unconscious knowledge gap. An unconscious knowledge gap related to the *data*, might be, for instance, that the analysts are unaware of missing values, noise, or outliers in the data. This may affect the analysis outcome. There might be unknown biases, and analysts might observe false patterns in the data. Unconscious knowledge gaps might lead to wrong interpretations and conclusions. If the unconscious knowledge gap is related to *tasks*, analysts may use wrong procedures to pursue the analysis. In the case of an unconscious knowledge gap related to *VA methods and algorithms*, the analyst might simply use the wrong parameters or select parameters unsuitable for the task. If the unconscious knowledge gap is related to *knowledge and insights management*, analysts might be confident about specific observations made and conclusions reached when instead the initial hypotheses would need to be re-evaluated, e.g., by looking at the data from another perspective. An unconscious knowledge gap should be treated with special care,

since it may reduce the acceptance of the guidance. To prevent this, guidance designers should consider ways to make analysts aware of problematic situations before providing the guidance. Solving issues at this stage is a way to increase trustworthiness (R2) and to achieve non-disruptive solutions (R5).

Identification of the Knowledge Gap

The correct identification of the knowledge gaps is related to supporting the design of non-disruptive guidance (R5). A wrong or incomplete identification may lead to wrong or sub-optimal guidance that can lead to unexpected analysis outcomes.

Q5 - How can potential knowledge gaps be identified during the analysis?

The easiest solution is to let the analysts enter the knowledge gap directly. This solution works well when analysts are aware of the knowledge gap, and know that guidance is available (R1). Conversely, the knowledge gap could be indirectly inferred from the analysts' actions when working with the visualization by analyzing the interaction behavior. For example, a user that fiddles quite a while with the user controls of a parameter might indicate that the user has difficulties setting a suitable parameter value. To summarize, two main mechanisms can be identified:

- Knowledge Gap Interface: Enables *conscious* analysts to communicate the knowledge gap to the system. This is useful only if analysts are aware of their knowledge gap.
- Knowledge Gap Inference: Enables the system to derive the knowledge gap from analysts' behavior. This is particularly useful when they are not aware of their knowledge gaps.

Step 2 – Risks, Threats, and Countermeasures

Similarly to the previous step, possible risks emerging from a non-satisfactory execution of this design step are the underestimation and overestimation of the possible knowledge gaps that may arise during the analysis. A viable solution to this problem would be the identification of critical analysis scenarios. These correspond to those moments in the analysis in which it is mandatory for the analyst to take decisions. If the end user is not required to take decisions and to reason about alternatives, guidance is not needed. The identification of these critical moments is crucial to avoid such threats, since these are the situations in which knowledge gaps might occur.

Underestimation and overestimation are related to the **completeness** of the designed guidance solution. This means that a major threat to the design comes from a mismatch between the analysts' needs and the guidance conceived to solve such situations, which conflicts with R3. Although this represents a formidable challenge for research, from the practical point of view, this risk needs to be minimized when designing guidance. While

it is hard to guarantee that a guidance solution is complete, in the following we provide suggestions to minimize this threat:

Design for the top-N knowledge gaps To improve design completeness, as an initial step, designers could start thinking of guidance to cope with the most problematic knowledge gaps, and thus, design guidance for the majority of crucial cases.

Design adaptive guidance In a second step, designers could aim for *adaptive* guidance mechanisms that could learn as the system is being used [Sil91]. In this way, designers should not worry about incorporating all the predefined content and rules (i.e., what to do when X or Y happens), but just define the boundaries in which the guidance can be provided. Machine learning techniques might be a good choice for such learning mechanisms.

Let the analysts guide themselves The generation of dynamic content cannot always be pursued. A learning system might of course also fail in some situations, which may result in the provision of incomplete, or worse, wrong guidance. Therefore, designers need a backup solution for such cases. To avoid the aforementioned problem and at the same time improve the completeness of the design, we solicit the design of mechanisms to help analysts guide themselves. In practical scenarios, this corresponds to providing analysts with all the necessary information they might need to make a legitimate choice. This could be helpful for instance, during exploratory analysis, when analysis goals cannot be precisely defined. Although this solution puts a large part of the burden on the analyst, we argue that, in case of doubt, it is better than providing the analyst with imprecise recommendations.

4.5.4 Step 3: Guidance Generation

This step deals with designing the appropriate guidance needed to narrow or resolve the knowledge gaps (step 2) and get closer to the analysis goal (step 1). Designers have to consider the characteristics of the guidance required as well as the moment to provide it.

Guidance Characteristics

We structure the guidance characteristics according to the work by Ceneda et al. [Cen+17].

Guidance Degree Designers should decide how much support the analyst needs. Proper mechanisms are needed to adapt the guidance degree to the current analysis situation. The questions are:

Q6 - What degree of guidance is needed? What mechanisms can be employed to switch among different degrees?

The choice of the guidance degree is mainly influenced by the analysts' prior knowledge. Too much or too little guidance might be detrimental, depending on the user knowledge

and experience. Consequently, a dynamic degree is preferred, since user knowledge can vary from task to task.

There are three guidance degrees [Cen+17]: orienting, directing, and prescribing. *Orienting* guidance provides users with hints so that they can orient themselves and maintain their mental map. It usually makes use of auxiliary means, such as highlighting or transitions between states, enabling users to seamlessly switch the analysis context and pursue different exploration goals. Orienting guidance could, for our soccer example, highlight interesting values or players (e.g., the player who completed the most passes) for further analysis. *Directing* guidance provides more assistance to the analyst than orienting guidance, usually in the form of an ordered list of suggestions. For instance, automatically suggesting and ranking the most prominent events in the soccer game, based on some interestingness measure. Finally, the last degree of guidance aims to *prescribe* a set of actions analysts should take to overcome their knowledge gap. The system could even carry out the actions autonomously.

From a designer's point of view, the provision of the most appropriate guidance degree is mandatory for an effective analysis (R0). Designers should consider all degrees, as well as the design of mechanisms to seamlessly switch between them. The employed guidance degree should match the analysts' knowledge, in order to not be too restrictive or leave too many unknowns. Providing the most appropriate guidance degree (at the right time) is important to meet requirements R2, R3, and R5.

Guidance Input Considering the guidance input means considering all the different sources that might be useful to produce guidance. The designer has to ask:

Q7 - What input is available?

Usually, different types of inputs are available at the time of designing guidance. The *data* under analysis might be used to extract statistics about the team players. An input that is commonly exploited is a *knowledge* base, for instance a catalog of labeled soccer events. Another factor that needs to be taken into account is who the end *user* is. The user knowledge—both operational and domain knowledge—has a direct impact on the type of guidance needed. A coach, for instance, might be interested in the tactics of the next opponent. In contrast, fans might seek more information on their favorite club. The *history* of user actions and information about *provenance* may also be useful inputs to generate guidance. Finally, user *preferences* and possible subjective *biases* should be taken into account as well.

Algorithms and Procedures to Calculate Guidance At this point, knowing about the possible *inputs* and the *degree* of guidance, it is important to identify suitable algorithms to compute the guidance output. The following question has to be answered:

Q8 - What algorithms and procedures are needed to generate guidance?

Algorithms for producing guidance vary according to the scenarios in which guidance is needed and according to the knowledge gaps. Algorithms for producing guidance refer to *how* guidance is generated and might be different from the algorithms used to identify the knowledge gaps (Step 2).

Guidance Output Once produced, the guidance output must be provided to the analyst. Usually we consider visual means, but also acoustic or even haptic output might be helpful.

Q9 - What are appropriate means to communicate the guidance output?

In order to support R1 and R5, appropriate means need to be selected. The guidance designer may choose to provide suggestions/hints in the form of simple *text*. Other frequently used expedients to convey guidance are *highlighting and changing color* of interesting data items [HB05]. *Motion and animation* could also be used to communicate guidance [Joh+05]. However, *glyphs and visual artifacts* are the most common way for encoding guidance suggestions [CGM19].

Identification of the Moment to Provide Guidance

Finally, the last question related to the design of appropriate guidance is identifying the correct moment or time frame to provide it:

Q10 - When should the guidance be provided?

One might think that the instants that immediately follow the detection of the knowledge gap are the best option in every situation. However, this may depend, for instance, on the task and on the analysts' behavior. The choice of the wrong moment to provide guidance may negatively affect the acceptance of guidance (R1) and disrupt the analysis flow (R5).

Step 3 – Risks, Threats, and Countermeasures

Similarly to the previous design steps, we mention threats affecting the design and discuss available countermeasures. Possible risks deriving from a non satisfactory execution of this step are: the introduction of **biases**, the choice of a **wrong guidance degree**, and the choice of **wrong timing** for providing guidance. A wrong realization of step 3 would counteract the implementation of trustworthy (R2), adaptive (R3), and non-disruptive (R5) guidance.

It is well-known that we, as humans, are affected by cognitive biases [HNM15]. These biases represent a systematic deviation from what is generally recognized as a *rational judgment*. Types of biases can be, for instance, the *confirmation bias* where users tend to stick to hypotheses that comply their way of thinking, or the *repetition bias* in which a user trusts, and thus, remains anchored to repeated procedures. Guidance can help to

solve biases but it can also introduce new biases itself. Hence, it is essential for designers to understand users' biases, take them into consideration during the design and think of guidance mechanisms to break systematic, wrong cognitive patterns. On the other hand, it is necessary that the designed guidance does not itself introduce further (unwanted) biases. If we stick to the mentioned example, if a system provides the same guidance suggestions in similar analysis scenarios, then as a consequence the user may learn that in such situations a predefined set of actions can be used to exit a stalled situation. However, if the scenario changes just slightly, the assumption that those actions are still useful may no longer be valid (bias of repetition). Such biases should be recognized by the system and their introduction should be carefully considered by the designers. As a solution, designers could think, for instance, of mechanisms to warn the user about the changed context. Unfortunately, these biases are subjective by nature and a generalized solution cannot be devised for each and every guidance solution. As a general advise, it is recommendable to conduct the design in collaboration with the end-users. Considering iterative cycles alternating design and evaluation phases could also help mitigate such problems.

Further threats affecting this step derive from an inappropriate provision of guidance, i.e., a wrong timing is chosen, and from the selection of an unsuited degree of assistance. It is easy to imagine that providing guidance at the wrong moment may sway the analyst. In the same way, the choice of a wrong guidance degree may frustrate users, limit their actions and nullify the benefits of guidance. Although guidance theoretically might be required at any time, it is worth mentioning that in practical analysis scenarios, when discrete interaction is involved, it is indeed likely that guidance is needed only at distinct points in time. These moments, in fact, correspond to those situations in which the user is required to take a decision or make a judgment [Sil91]. In the absence of these cases, the opportunities to offer guidance are minimal. Therefore, to avoid the aforementioned problems, the role of designers is to identify these decision points in the analysis, to define an order of such moments, and define critical decisions. Providing guidance and limit the alternatives available at such points can make a difference in a successful analysis.

In general, we recognize that identifying the precise moment to provide guidance is not always possible. Exploratory analysis is a viable example of this, since its goals as well as the whole process are affected by a great degree of uncertainty. However, also when this is not possible, providing the analysts with orienting guidance, i.e., providing all the necessary information about possible actions so that the analyst is able to make an informed decision, can be a suitable baseline solution.

4.5.5 Step 4: Guidance Feedback Loop

When designers know how to identify the analysts' knowledge gaps and possible guidance solutions to close or narrow them, they need to design means that allow analysts to fine-tune the provided guidance. Guidance is a mixed-initiative approach [Hor99], and proper methods to steer the process must be identified. With the aim of designing such feedback mechanisms (i.e., guidance for the system), designers should think of two main

aspects: 1) Mechanisms to derive guidance for the system from analysts' actions (usually in the form of feedback). We will refer to this aspect as *guidance inference*. 2) The direction of such guidance: guidance can be directed towards the past or the future. This step is aimed at assuring that the provided guidance is controllable (R4).

Inferring Guidance for the System

Interaction is the most common way for analysts to fine-tune the guidance, for instance, its degree [CGM19].

Q11 - How can the system derive guidance from the analyst's actions?

At this design stage, the designer should decide whether sequences of direct actions, or indirect signals, or both should be considered to infer the analysts' feedback about the provided guidance. Two kinds of feedback can be identified:

- Direct feedback: the analyst moves *sliders* or uses other controls for changing the guidance parameters directly.
- Indirect feedback: the analyst acts on the *data*. Analysts move the data, group the data, label the data, which affects the guidance algorithms indirectly.

The literature on using *direct* interaction in visual analysis is vast [AES05; Shn96]. Interaction can be used to provide feedback to the guidance process, too. In other words, the analyst fine-tunes the guidance parameters by means of interaction with user interface elements, such as widgets, buttons, etc. For instance, if analysts are not satisfied with the data grouping suggested by the guidance system, they may use sliders to adjust the results. The guidance system should hence adapt future guidance results. Usually, single actions are considered. In other cases, the history of actions is contrasted with a knowledge base to extract useful usage patterns [Fuj+97].

The second interaction method is what we refer to as *indirect* feedback [End+14]. This is the case when analysts do not directly communicate their feedback to fine-tune the guidance system, but the feedback is derived from their interaction with the data (and not with the widgets). For instance, the analysts' intention to change the data grouping may be indirectly derived from the action of moving specific data points closer to each other, in contrast to the direct use of sliders or widgets. Although direct feedback is the most common method, indirect interaction might open the door for more natural feedback, since it allows a direct contact with the data, which could lead to better user acceptance (R2).

Direction of Feedback

In the previous step, we identified the analysts' feedback and ways to infer it. In this step, designers have to identify the direction of the feedback.

Q12 - What is the direction of the analysts' feedback?

As mentioned, the guidance directions can be *past* and *future*. Following the literature in cognitive sciences [Dow99], we refer to actions towards the past as *feedback*, and actions taken to call for future guidance as *feedforward* actions. Our idea of feedback, is similar to the one used in cognitive sciences [Dow99], and is related to the concept of relevance feedback. With relevance feedback, relevant items, for instance, the results of a query, are used by the system to provide further guidance to the user. However, in this case it is the user that guides the system and steers the guidance process. As designers it is important to specify the quality of such evaluation: *positive* and *negative*. Positive and negative feedback are meant to provide a positive or negative evaluation of the guidance the system has provided in the previous analysis loop. Feedforward actions, either positive or negative, should enable analysts to provide hints how they want the guidance to look like in the next guidance loop, and thus, steer and refine the generation of appropriate guidance suggestions.

Step 4 – Risks, Threats, and Countermeasures

A major threat to the design is an unsatisfactory realization of R4: **controllable** guidance, and thus, the provided guidance cannot be controlled by the user. In other words, there is an imbalance between the possibilities offered by the system and the requests of the user: the system guides and forces the choices of the analyst, but the analyst cannot guide and steer the analysis. In some situations, limiting the available alternatives is desirable, for instance, when is recommended to perform a limited set of actions. However, this cannot be assumed as a general design pattern, as the analyst may need a larger set of analytical options and be enabled to deviate from the current line of inquiry. The literature about the science of interaction is vast [Dix+04]. Designers should choose and design the interaction flow considering the analysis requirements and find a suitable balance between restricting and guiding the analyst.

4.5.6 Iterative Design of Guidance

Having described the different steps and the requirements, we want to make a short digression discussing the iterative nature of our framework. It is common practice in computer science but also in visualization and VA to consider iterative cycles of design, in which a product or a process is cyclically refined with respect to user feedback, in order to obtain a satisfactory result. It is also common that the number of design goals increases or the goals change in the course of this iterative design process. Our framework follows the same strategy by providing the possibility to move back and forth between the steps. For instance, the understanding of the analysis goals might change (Step 1) as guidance mechanisms are defined (Step 3). Our framework proposes a set of qualitative requirements and provides a list of easy-to-use design questions for each step of the process. These qualitative requirements and design questions help users to design comprehensive guidance mechanisms even when refining the design multiple times.

4.6 Designing Guidance: Three Scenarios

The framework can be applied to a wide range of scenarios in the context of guidance design for VA.

To make the framework easier to understand for the reader, we illustrate it by describing three examples and a comprehensive design walk-through. The three design examples are taken from literature. Some of the authors of this paper collaborated to their development in various ways. Instead, the design walk-through describes a complete design which we performed from scratch using our framework. While the examples should be useful to understand the different aspects considered by the framework, the walk-through should illustrate a way to instantiate it and make it actionable. Table 4.1 complements the examples by summarizing the answers to the questions posed in the previous section.

Questions	Cyclical Patterns	Condition Monitoring	Fraud Detection
Q1 What are the analysis goals?	explore cyclical patterns	anomaly and failure detection	frauds exploration and confirmation
Q2 In which analysis phases might issues occur?	exploration	exploration	model building
Q3 What is the knowledge gap?	parameters	data/VA methods	parameters
Q4 Is the knowledge gap perceived/unconscious?	perceived	perceived, but bias may occur	perceived
Q5 How can knowledge gaps be identified?	case study interviews	case study interviews	case study interviews
Q6 What guidance degree is needed?	orienting	orienting/directing	directing/prescribing
Q7 What input is available?	data	data/thresholds	data/ domain knowledge
Q8 Algorithms to produce guidance?	cycle detection algorithms	correlation/classification algorithms	neighborhood exploration
Q9 Appropriate means to encode guidance?	glyphs in sliders	overview/marks	forbid certain queries
Q10 When should guidance be provided?	throughout	throughout	throughout
Q11 How can guidance for the system be derived?	direct feedback	direct feedback	direct feedback
Q12 What is the direction of the feedback?	n.a.	forward and backwards	forward

Table 4.1: Questionnaire summarizing the design of guidance in three application scenarios. The questionnaire is based on the guidance design framework. (1) Guidance for Cyclical Patterns Exploration [Cen+18b]. The knowledge gap refers to the lack of knowledge regarding the length of the cycles in the univariate time series. The questionnaire shows that while many aspects are considered in this guidance design, question Q12 is not fully answered (n.a.). (2) Condition Monitoring and Failure Detection: The focus is on high dimensional multivariate time-series data. Guidance is needed to correctly set the parameters of the algorithms to detect anomalies and correlations across events. (3) Fraud Detection in Financial Systems [LGM19]. Guidance is needed to support the analyst in analysing a financial transaction graph and discerning whether such transactions are frauds or regular money movements. The knowledge gap refers to finding parameters and form non-empty and meaningful queries to the system. Also in this example, the designed guidance is quite comprehensive, as all the questions are answered.

4.6.1 Exploration of Cyclical Patterns in Time Series

In our first example, we address the visual analysis and exploration of cyclical patterns in univariate time series [Cen+18b]. For unknown data, it is typically not clear beforehand if and where cycles and patterns exist in the data. This leads to time-consuming phases of trial-and-error searching, where analysts have to spot a possible pattern and then verify its existence in the whole dataset. A purely algorithmic solution to find cyclic patterns is not feasible either. Algorithms to automatically detect cycles are difficult to select and configure. Thus, guidance is designed to mitigate these problem by reducing time consuming tasks.




Figure 4.3: Guidance for exploring cyclical patterns [Cen+18b]. Analysts are supported in finding cycles. Suitable cycle length values are encoded in the sliders (see the gray bars on the left-hand side of the image) that control the visualization of patterns. By choosing the suggested values, cycles appear in the visualization.

The idea is to support the detection of cycles by indicating possible instances of cyclical patterns. This information will guide the user towards configurations that will potentially make cycles visible in the visualization (see Figure 4.3).

Step 1: Analysis Goal This design example is limited to a specific, yet important, research challenge, that is identifying patterns in cyclical data. We specify this design problem as a sequence of two sub-problems: The first problem concerns the **exploration** of the dataset. We want to support users in finding cycles and recurrent patterns. The second problem is that, once analysts have explored the data, they have to **build a model** (for instance, understand the regularity of the discovered cycles) and formulate appropriate hypotheses. In this scenario, we imagine that issues might appear in such analysis phases.

Step 2: Knowledge Gap We can imagine the analysts being experts in the field. This means that they possess sufficient domain knowledge to interpret the data correctly. However, patterns and cycles might not be known in advance. Hence, the knowledge gap can be framed as a execution problem. Analysts do not know in advance what **parameters** will make such cycles appear. In the design of the guidance solution, methods to infer the knowledge gap during the analysis are not mentioned. Since the design of guidance is limited to a specific research problem, the risks of over/underestimation of the possible knowledge gaps do not exist, in this case.

Step 3: Guidance Generation In order to provide an answer to the aforementioned knowledge gap and support a successful data exploration, the guidance provides suggestions of possible parameter settings that would make cyclical patterns visible. Thanks to such suggestions, analysts can configure the visualization in a way that lets them explore the most promising cyclical patterns.

Guidance Degree - Analysts are supported with **orienting** guidance. The choice of this guidance degree is due to the fact that the analyst needs to perform an exploration analysis. Since the importance of the detected patterns is not known in advance, it is a better option not to guide the analysts directly by providing recommendations (i.e., directing guidance), but rather enable them to make an informed decision. The designed guidance shows the automatically detected patterns but does not enforce any order and provides the analyst with statistics about these patterns. Hence, in this scenario, analysts were allowed to formulate and test their own hypotheses, without being influenced by the provided guidance.

Guidance Input - The only input needed for the guidance process is the *data* itself. Together with the aforementioned algorithms, the data is used to calculate a list of possible parameters settings that can make the cycles appear, if present.

Algorithms and VA Methods - In order to produce the suggestions, the guidance process exploits two algorithms (i.e., the Chi-Square Periodogram (CSP) and the Discrete Fourier Transform (DFT)), which are commonly used for finding cyclical patterns in time series. The DFT provides precise indications of patterns, while the CSP complements them with a probability, which constitutes a source of guidance for the analyst. During the design, we chose these algorithms because they complement each other. The algorithms produce a list of long and short patterns, giving the analyst a nice overview of the data.

Guidance Output - Once the suggestions are computed they must be visualized. In this scenario, the suggestions are encoded directly using the sliders to modify the visual appearance of the visualization. This choice was made to avoid distracting users from the exploration activities (supporting R5). Furthermore, this choice also reduced the risk of introducing biases, since the suggestions are integrated in the normal analysis workflow. The idea is to assign and visualize the output of the algorithms to the place where analysts have the opportunity to identify them. This implicitly makes also the guidance solution immediately available (R1).

Guidance Timing - In this specific example, we did not consider time frames to provide guidance, since there are no critical judgment moments. Analysts are required only to judge the different alternatives provided by the system and formulate hypotheses. Therefore, providing the analysts with a detailed list of patterns at the beginning of the analysis was considered a sufficient source of guidance. In a more complex design scenario, we could imagine the system supporting also the choice among alternatives with a higher degree of guidance (e.g., directing guidance). In this case, critical moments to provide guidance could be the moments preceding the choice of a specific pattern, after the analyst has already filtered out the less promising patterns. At those points,

directing guidance and recommendations could be effectively provided.

Step 4: Guidance Feedback Loop In the current iteration, the described guidance solution does not allow fine-tuning, meaning that it does not allow the algorithms nor their parameters to be changed. However, direct feedback could help users to decide how many, as well as what kind of patterns are suggested. This kind of solution, would constitute a *feedback* to the system and could also work to evaluate which algorithm may provide the better results.

4.6.2 Condition Monitoring and Failure Case Detection in Engine Testing

In automotive engineering the analysis of test data obtained from an engine in a test-bed is a common task. Engine testing is a key phase in engine development, and serves as verification and validation of engine designs. Typically, engines go through repeated, programmatic test cycles in the test-bed. For analysis purposes, numerous sensors are equipped, which record characteristic properties of the engine over time. Typically, multiple timelines are used to represent sensor measurements. The primary goal of engineers and analysts is the detection of anomalies as well as their root cause, which may be related to design errors. In this scenario [Sus+20], guidance could be used to reduce the burden on the user to detect anomalies. However, since anomalies may vary, also the knowledge of the analyst is of great importance to detect relevant abnormal events. Hence, a proper balance needs to be found when designing guidance, between user freedom and system restrictiveness.

Step 1: Analysis Goal Based on design study interviews, two main goals have been identified for this use case (see Table 4.1). Under the **exploration** goal, analysts want to test if the engine behaves as expected or is affected by anomalies. Also, identification of correlated and uncorrelated measurement data is important. Depending on the engine design, measurements may influence each other, or be independent from each other. **Model building** involves finding a description for regular and anomalous test states, eventually rules to install for an automatic monitoring. Hence, analysis phases include verification and falsification during exploration and monitoring. A guidance system should be available (*R1*) and adaptive (*R3*) to these tasks. As the tasks include both data analysis for model building and monitoring for failure cases, the guidance may necessarily be disruptive at times (*R5*) but should exercise the disruption only when needed.

Since engine certification tests are standardized routines, the risk of underestimating (see Section 4.5.2) the analysis goals almost does not exist. The analysis of test cycles can be easily identified as an exploration task, with the aim of identifying anomalies. As for model building, this goal was introduced in the attempt to partially automate and ease the detection of anomalies, reducing the burden of the analysis. However, a fully automatic monitoring is not possible due to the changing conditions of each test.

Step 2: Knowledge Gap Analysts are trained automotive and mechanical engineers. They possess domain knowledge on expected engine characteristics under varying loads and effects of wear over time during the set of test cycles. A first knowledge gap can be framed as a **data** problem. Patterns in the data may represent normal and abnormal engine states. Some are known from experience and training, but for newly developed engines, new patterns may occur during verification and validation. Also, normal and abnormal states can be described not only by single variables, but by combinations of variables and their interplay. There can be abrupt but also smooth transitions between normal and abnormal conditions. This is a large search space. In addition, it may occur that sensor readings become imprecise or erroneous due to failing sensors, which may not immediately be apparent. A second knowledge gap is represented by the choice of the *algorithms* and VA methods. Not all the statistical algorithms are suited to detect a given anomaly. Therefore, analysts should be also guided to choose among alternative detection algorithms. Experts in general are **aware** of the knowledge gap, but may be biased to look for expected variables and at the expense of new variables or their combinations. Trustworthiness of the guidance ($R2$) will be especially important if unknown or unexpected parameters are suggested for analysis. In this step, threats to the completeness of the designed solution could be avoided with the implementation of learning mechanisms for guidance. Well known patterns must be taken into consideration. However, a simple rule-based guidance is not enough and fully automatic analysis is not possible either, due to the changing conditions of each test. The nature of the task asks for the introduction of learning mechanisms, to adapt to new anomaly patterns, in which new rules are dynamically added to complement the existing knowledge base.

Step 3: Guidance Generation A key task is to learn what normal and abnormal conditions are. This can be supported by guidance approaches based on showing a suitable degree of similarity between engine measurements over time. Assuming that most of the time, the engine test is in a normal state, abnormal states could show large differences when compared to reference data. This should be done for large amounts of data recordings and many variables. Because data is large, the idea is to support the analysis goals by adding a further level of abstraction on top of the analysis workflow. Instead of analyzing hundreds of timelines, analysts will be provided with glyphs that would point to possible problematic situations, reducing in this way the search space.

Guidance Degree - Two main guidance degrees apply in this use case. The exploration task can be supported with **orienting** guidance. The guidance system continuously evaluates the measurement data for abnormal behaviors and reports occurrences to the users. The approach is designed to point the user to adjusting the thresholds if needed, and hence is controllable ($R4$). Orienting guidance was chosen since a ground truth does not exist for this task. The variable nature of the anomalies makes the task hard to solve in an automated way. Hence, higher degrees of support are not suitable. Providing recommendations (as in directing guidance) is in certain cases not possible and even detrimental.




Figure 4.4: The SignalLens approach indicates detected time series anomalies by level-of-detail and markers [Kin10]

Directing guidance is instead desired for model building activities. In this, the guidance could learn from user feedback and subsequently help analysts in choosing the most appropriate algorithms and statistical methods for a given scenario and test cycle.

Guidance Input - This will be the **data** to be monitored. An initial set of **parameters** need to be set, specifying e.g., thresholds and intervals for the initial anomaly detection. We can assume that rules exist from engineering knowledge and best practices, but they will need to be adapted during the long-lasting test runs.

Algorithms and VA Methods - As the data set is large, the designed solution requires the application of data reduction techniques. This can imply reducing the frequency of the data, e.g., by sampling time steps. Also, feature selection methods must be applied to reduce the number of variables. Still, the amount of data may be large. For this reason, in a first analysis step, a measure of the anomalies is calculated [Zha+17; CBK09]. Analysts could use this measure as a first indication about the presence of possible anomalies in the data.

The designed solution also requires that analytical methods are applied to compare current data with historic data and report larger differences as possible anomalies.

To support this task, a regression model is used [LW+02] and the detected features are visualized by level of detail and markers (see Figure 4.4). However, these algorithms represent just an initial step into the analysis. Since anomalies vary, the system allows for an easy interchange of the algorithms to use in a give scenario.

Guidance Output - Various visual methods can be used. The analysts, in a normal workflow, are used to simple time-series visualizations of the recorded sensor measurements. Due to the data dimensionality, the designed solution introduces visual glyphs to give the analysts an overview of possible detected anomalies.

In the following phase, when measurements have to be compared among each other, a scalable approach is based on the visualization of difference matrices to compare the linear relations between sensors of a known normal cycle with those of an unknown cycle. Large differences in the correlation of certain variables hint to possibly anomalous parameters [Zha+17] and are visible as prominent rows and columns (see Figure 4.5). A possible risk is represented by the introduction of biases in the analysis: the additional

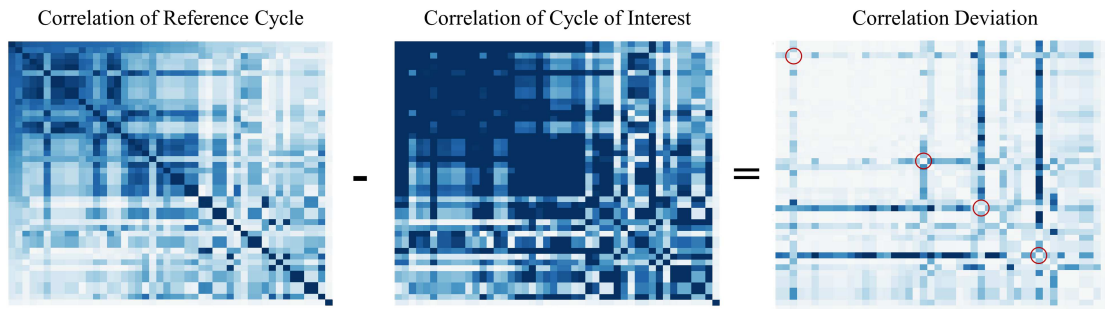


Figure 4.5: Detecting anomalies by comparing sensor data correlation matrices of a reference test cycle (known normal) with an unknown test cycle. Larger differences in columns or rows hint at anomalous values, in comparison to the reference cycle. By means of visual pattern-driven exploration of cross patterns deviating sensors can be identified. Four examples of such cross patterns are highlighted in the figure by red circles at the cross's intersection on the diagonal element.

step introduced to help the detection of anomalies, should not put additional burden on the user. Hence, to reduce even more such burden on the user, only parameters that are relevant for the current analysis are visualized. Furthermore, guidance has been designed to be included in the normal workflow of the engineers, so as not to disrupt (R5) the reasoning process, but providing support to it.

Guidance Timing - Timing is relevant. According to the analysis workflow of the engineers, it is possible to identify two main moments when guidance has to be provided. If the engine stops working during the tests, analysts need to be guided to the root cause of the malfunction. The second moment is at the end of the test cycle, when data is analyzed and anomalies have to be identified. The two scenarios require similar actions: Analysts should be guided to compare the detected anomaly with the reference model of the engine function and detect design issues.

Step 4: Guidance Feedback Loop Feedback modalities have been selected to make the analysis controllable (R_4). User **feedback** on thresholds can help to refine anomaly detection. In particular, not all the detected anomalies are relevant. When the guidance reports such an event, the analyst can fine tune the anomaly detection algorithm, by excluding irrelevant sensor measurements, steering the analysis. Refinement should be possible by the user continuously. Since the analysis is in real time, the knowledge may change anytime but needs to be reflected immediately by the system. **Direct** input is commonly used in this scenario.

4.6.3 Visual Detection of Frauds in Financial Systems

Financial institutions are interested in ensuring that illicit operations, i.e., fraudulent money transactions, are detected and prosecuted in short time. Fraudulent schemes

have nowadays a huge impact on the financial system, impacting the economy and the trustworthiness of the institutions [Ko+16; Lei+18]. To tackle such incidents, financial institutions analyze on average millions of transactions (money movements) per year, the majority of which are legitimate, to detect possibly unlawful schemes and behaviors. The amount of data being analyzed does not permit for manual exploration of all cases, and a first labeling of the data is made by automatic algorithms. Afterwards, financial fraud analysts are in charge of making the final decisions, i.e., should a customer be accused of fraud?, by analysing a subset of transactions. Fraudulent schemes are, usually, complex: This requires to analyse particular structures (patterns) in the transaction graph. The cost and the implications of possible false positives, i.e., accusing an innocent person, are high. The usual analysis process, hence, involves a computational system searching for candidate patterns (i.e., possible frauds) and a data analyst who is responsible for analyzing the critical cases and deciding final verdicts.

Step 1: Analysis Goal Working in close collaboration with financial fraud analysts, a number of tasks and goals was defined. The focus of the guidance in this example scenario is to support specialized analysts in the detection and analysis of possibly fraudulent money movements and understand if they comprise criminal actions. In particular, given a specific bank account, analysts want to understand the flow of money in a quick and effective way. It should be also possible to perform the same tasks considering multiple accounts and their relations, which in this case corresponds to understanding how money is moved through a network of selected accounts. The second task consists of understanding the structure of a transactions network, i.e., to understand if the considered transactions constitute fraud or not. To do so, analysts possess the required domain knowledge to judge individual cases. However, they still need support to detect possibly hidden patterns.

The first task is an **exploration task**, i.e., exploring the transaction network (compare Section 4.5.2). The second task is about **building and evaluating a model**, i.e. understanding if the given network represents a fraudulent scheme (see also Table 4.1).

Step 2: Knowledge Gap The analysts are bank employees who are experienced in the domain and, therefore, know the data and tasks very well. Analysts can form and send queries to an internal scoring system developed by the bank, looking for suspicious patterns, but they do not know if such queries are meaningful and represent an actual pattern in the transaction graph. This results in long and often unsuccessful trial-and-error analyses in which the analyst has to build and refine the queries in multiple stages. In other words, when performing analytical tasks, the analysts' knowledge gap relates to finding meaningful *parameters* and combinations thereof, which will not yield empty or contextually irrelevant queries' results, so to foster an effective *exploration* of the transaction network and detect financial frauds. The designed guidance addresses such issues in that it supports the analyst in forming meaningful and non empty queries. The identification of possible knowledge gaps and analysis goals was pursued in collaboration with financial fraud analysts in terms of design study interviews, to minimize the risk of

underestimating possible knowledge gaps and to have a clear view of the goals of the project. Designers were able to frame all the possible transaction schemes to a finite set of basic cases. Using such cases as building blocks, analysts are able to construct complex queries without limitations to the query expressiveness and fraud detection capabilities. On top of this, the guidance was designed to avoid the formulation of queries outside these building blocks or cases which are not present in the data.

Step 3: Guidance Generation *Guidance Degree* - The guidance degree needed in the context of this scenario is a combination of **directing** and **prescribing**. As analysts often face the problem of finding meaningful patterns and formulating appropriate non-empty queries' results, the designed guidance **indicates** parameter combinations that produce contextually relevant queries and **prohibits** the formulation of queries that lead to an empty result.

Guidance Input - The first input to the guidance mechanism is **money transaction network data**. Accounts/customers are represented as nodes, while the flow of incoming/outgoing money is represented by edges between nodes. In a first step, an automatic fraud detection system flags suspicious transactions. The analyst is then responsible for delving into these specific cases and confirm or reject their criminal intent. A further input is the **domain knowledge**, which is used to formulate queries, which are meaningful and non-empty to avoid irrelevant results.

Algorithms and VA Methods - The manual analysis of the transaction graph is not feasible, as these graphs contain hundreds of thousands of nodes (the accounts) and millions of edges (the transactions). After suspicious transactions have been identified by the automatic fraud detection system (provided by the bank, but we are not allowed to describe the algorithm due to bank regulations), the exploration of the subset of suspicious accounts and transactions is supported by a VA solution. When the analysts form queries to explore the transaction graph, algorithms are used to conduct a preemptive exploration of the neighborhood of the user-selected nodes. This exploration allows our VA solution to detect meaningful transaction patterns and consequently support appropriate query formulations. As the analysis makes progresses, the network's neighborhood is continuously updated and only relevant actions are allowed.

Guidance Output - The provided guidance approach supports the exploration of the whole transaction network by restricting the parameter space and allowing only for the formulation of meaningful queries resulting in non-empty output. A prescribed set of queries (also called building blocks) is hard-coded in the VA solution. These restrictions do not hinder a comprehensive analysis, as they allow analysts to cover all cases present in the data, and thus, to effectively solve their tasks. Therefore, the risk of missing possible frauds is avoided (which supports R3). While fraud detection is always affected by some degree of uncertainty, analysts are trained and aware of it. Hence, the risk of misinterpreting the recommendations is considered low.

Fraud analysts are used to working with visualizations, however, not as their primary means of investigation. Thus, an expressive visual encoding for comprehensive visual

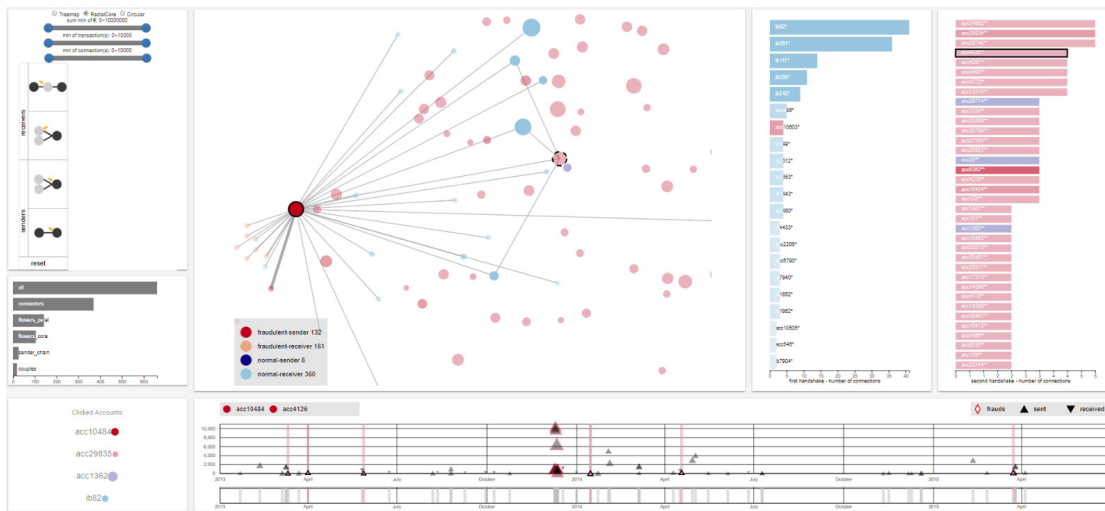


Figure 4.6: Overview visualization for finding frauds in a money transactions network [LGM19]. The analytical exploration of the network data is supported by allowing analysts to formulate only semantically relevant queries and make the analysis effective. The money transaction graph is displayed in the middle. The left-hand side of the interface is where the guidance takes place. It shows the building blocks, i.e., predefined query components, that the analyst can use to formulate queries. As the exploration proceeds, the list of building blocks is updated, some of them are removed, so that just meaningful and non-empty queries can be formed.

analysis was designed: The transaction graph is represented as a node-link diagram in the center of the visualization (see Figure 4.6). Another constraint was the ease of use, so to put no additional burden on the analysts (R5). Thus, the guidance suggestions, i.e., the allowed operations in a specific time frame were encoded as draggable building blocks to allow formulating queries in a visual way. Whenever a selector would lead to an irrelevant result, it would be grayed out and made unusable. This makes the guidance suggestions immediately available and visible (R1).

Guidance Timing - When the analyst selects a bank account for exploration, the guidance mechanism explores preemptively network's neighborhood searching for known patterns. As the exploration proceeds and new nodes are selected, the guidance mechanism updates and provides a new set of meaningful actions. In this sense, the guidance mechanism anticipates possible future actions of the analyst.

Step 4: Guidance Feedback Loop As mentioned, when analysts select a certain node for exploration, the guidance mechanism automatically explores the network's neighborhood for detecting potentially interesting money movements. In this sense, the analyst's actions, i.e., the selection of nodes to explore, influence the way the system expands the neighborhood graph. This can be considered as a direct feedback to the guidance. However, the number and type of patterns the VA solution is able to find

cannot be directly modified or fine tuned. To overcome this issue, in a current update of the guidance mechanisms, in addition to predefined queries, analysts are allowed to specify user-defined patterns, providing a finer grained feedback to the guidance mechanism.

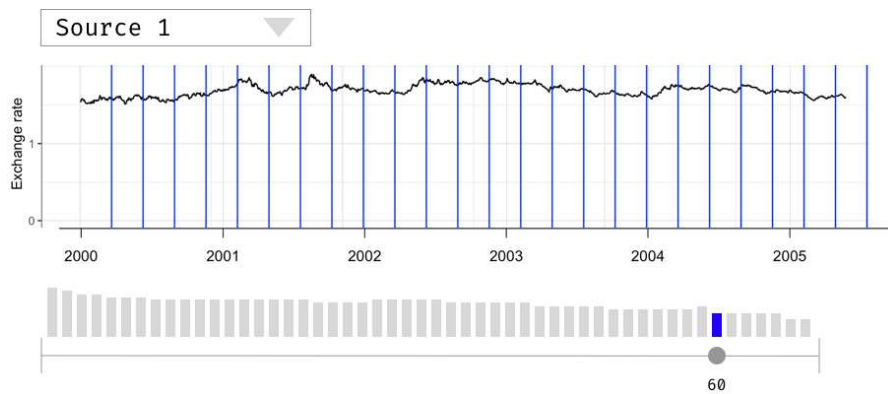
Having described three design examples, we can start drawing some conclusion: As shown in Table 4.1, not all the questions are covered by the examples we describe. The guidance feedback loop is often overlooked or considered just for minor parts of the guidance mechanism. Moreover, the examples we describe offer guidance for selected tasks of the analysis process only. Thus, our examples demonstrate practical solutions to specific issues, rather than guidance designs for comprehensive systems. This is a common issue in the literature: Existing works do not provide any comprehensive guidance solution nor complete design examples [CGM19]. In the following, to overcome such issue and to ease the application of our framework to other contexts, we describe a comprehensive and complete step-by-step design process.

4.6.4 Design Walk-Through: Guidance-enriched Blind Source Separation


In the following, we describe a design walk-through which should provide additional means to designers aiming at using our framework. The actual design walk-through was carried out by a designer, who is expert in VA but was not involved in defining the framework. After an introduction to the framework, we asked him to perform a complete design. We describe how the designer iterated over all the design steps considering risks, countermeasures, and design requirements.

Problem Description - We ground the design walk-through in the field of statistics. It happens often nowadays, that a variety of measurements are collected at different locations and times. The reader can imagine various sensors collecting, for instance, temperature measurements, or the fluctuating price of a given good on the stock market, in different regions of the world. These datasets are collections of multivariate time-series. Statistically speaking, one of the problems arising when dealing with such data is that it is hard to separate the actual measurement from other signals composing the time-series. In general, we can think of these other measurements as noise, but a wide variety of signals can add up to form the final value. This is a topic of interest in many domains. For example, physicians need to analyze and detect possible anomalies in an electric cardiogram (ECG) of a pregnant woman in which the heartbeats of the mother and baby are mixed together. This problem can be generally formulated as *separating a signal into its components without any assumption about the characteristics of the original signal*, and it can be shortly called a blind source separation problem (BSS). The task of the designer is to provide guidance and support to solving this problem.

Step 1: Analysis Goal The designer bootstrapped the design and approached Step 1 by performing a thorough literature research gaining confidence with the topic. At the same time, the designer had several discussions with the statisticians involved in BSS, which supported the understanding of the BSS problem and analysis goals of the statisticians.



(a) Guided selection



(b) Machine learning assisted exploration

Figure 4.7: Guidance-enhanced blind source separation (BSS). We enhance with guidance the task of separating a signal into its basic components. (a) The area above the sliders used to input parameters is used for showing the impact of different parameter choices to users, hence informing them about the possible outcome of a parameter selection. (b) Thanks to classification algorithms the output of the BSS is classified to enhance the exploration and interpretation of the results.

In this phase, he reiterated over Q1 (see Section 4.5) several times and gave an initial answer to Q2 too. Thanks to the interviews, he learnt that usually statisticians would tackle the problem by exploiting mainly the functionalities and packages of R [Tea14] for statistical analysis. By using this iterative method he was also able to keep low the risks associated with this first step (see Section 4.5.2), settling in the end with a compromise between over- and underestimation of the analysis goals.

A typical BSS task is usually approached as follows: By using R and other statistical tools, the data, i.e., the collection of multivariate signals, is analyzed by means of algorithms appositely created to the scope. The output of such algorithms are a set of signals representing the components of the original measurements. In a subsequent step, still through R, the statisticians produce a static visual representation of the results for inspection: They analyze statistics of the output signals and visually inspect them to understand if the result is sufficiently precise and if any interesting pattern is present. If that is the case, the analysis can be considered concluded. However, what often happens is that the analysis of the data has to be repeated multiple times varying the parameters of the BSS algorithms each run, confronting the results with the ones obtained in the previous iterations, and finally interpreting these results to understand if they make sense. Statisticians are usually not interested in finding the optimal solution, which would be anyhow unfeasible to calculate due to the large parameter space, but in calculating one solution that represents a very good, yet not perfect, estimation of the original signal. This statisticians' workflow resembles typical VA tasks [Sac+14]. However, very little emphasis is given to use visualizations or a comprehensive VA methodology. They use static images, which do not provide an easy overview and comparison. Hence, the designer decided, in line with R5, to give further emphasis to the visual means and support of the analysis, which were already a part of the workflow, but enhancing it by using a VA methodology and including guidance in the loop.

Step 2: Knowledge Gap While understanding the BSS problem and the statisticians' workflow, the designer moved towards the fulfillment of Step 2. Hence, in parallel to questions regarding the single analysis phases, the statisticians were also interviewed about possible problems they might encounter while completing their tasks. This led to defining a list of possible knowledge gaps, as required by Q3. In total, the statisticians mentioned three main knowledge gaps (defined as KG-1, KG-2, and KG-3 below). All of them can be framed as problems of execution. These will be addressed by the guidance. The analysis workflow is instead already very well structured and defined.

Statisticians have to modify several times the parameters of the BSS algorithms in order to complete the task. Although they can be considered as experts in that they know exactly the meaning and the influence of different parameter combinations, still the *parameter space* is huge (KG-1), which hinders the possibility of an exhaustive search. The second and third knowledge gaps relate to the *exploration* (KG-2) and *interpretation* (KG-3) of the results produced in the first phase. As the algorithms compute the components of a given signal, the results have to be compared with previously obtained components, and their statistical properties have to be compared among each other to judge the goodness

of the new obtained output. This requires to maintain and/or remember a collection of previously computed results, which in the original workflow required long back-and-forth exploration of statistics and static visualizations in R. Additionally, issues arise when analysts have to consider the computed results in the light of a specific context. Analyst must, in fact, possess not only knowledge about statistics, but also about the data domain. Similar results might be considered useful or totally useless according to the domain of the data. These domain competences concur to determine if a computed signal is a good representative of the original one. The same consideration applies also to the choice of parameters. For instance, some parameter combinations might make more sense than others according to different domains.

After some meetings, the designer had the impression that the statisticians were pretty aware of the knowledge gaps (Q4). In this regard, in earlier iterations of Step 1 and 2, the designer discovered that some of them often relied on some sort of rule-of-thumb methodology to solve the tasks. Therefore, from that point on, the designer decided to proceed in this promising direction, and shed more light on such established but implicit practices to see if it was possible to formalize and exploit them, solving KG-1 and KG-3.

Step 3: Guidance Generation As the designer gained some insights regarding possible knowledge gaps affecting the analysis, the designer looked into possible solutions to solve them. Step 3, which aims at designing guidance for the detected knowledge gaps, was initialized by analyzing the types of input available as well as the types of guidance that could be produced. Three inputs are mainly available to the guidance process (Q7): the data, the implicit knowledge of the statisticians, and the domain knowledge, depending on the application scenario. The data can be directly exploited: statistics can be extracted and used to give the statisticians an early idea of which parameters to choose or how to interpret the results obtained. The domain and the statisticians' implicit knowledge, as the word suggests, is not readily available. Hence, the designer looked into ways to formalize and make this knowledge explicit so that it could be used to support the exploration and the interpretation of results. In particular, this will come useful to support KG-1: the guidance will suggest possible parameters based on the data domain.

The visual analysis tool was designed to easily integrate the design framework into the normal design flow of any visualization environment. Its step-wise structure allows designers to integrate the two processes. Hence, the designer sketched the visual appearance of the tool to accommodate both the guidance and the user interface widgets. A multiple coordinated views approach was chosen, which supports the two main perspectives: 1) data selection and parameters imputation which should address KG-1 and partly KG-2, and 2) the results visualization which should allow easy exploration (KG-2) and interpretation (KG-3) of the results. The designer had a further interview round to gain a deeper understanding of the tasks. A new iteration over Step 1 was carried out and a list of the operations necessary to solve the tasks was defined. In total two critical operations were identified: We describe them as well as the guidance designed to support them, next.

The issues for the statisticians start at an early point of the analysis, since they are immediately called to determine the granularity of the data and the input lags separating the input signals. However, they know nothing about the data in such early phases, which usually reduces the analysis to long iterations of random parameters selections and results inspections. To help them, the designer thought of guidance which can be framed as **orienting** support exploiting the data input (Q6-7). As the user loads the data, the tool automatically calculates statistics about the loaded signals and immediately visualizes them to inform the parameter selection. The auto-correlation of the different input signals can be, for instance, displayed as a line chart (Q8) to determine a proper lag value that separates the input measurements and an appropriate granularity. Hence, following the statisticians' workflow, the tool was designed to calculate automatically such statistics and arrange them in the visualization to facilitate their work.

In the current design iteration, the guidance hints were integrated directly in the sliders used to select the parameters of the BSS algorithm (see Figure 4.7a), in a way that does not distract the user and makes the guidance readily available (R1-5). A small area above the sliders is reserved to visualize such hints (Q8). The guidance embedded in the sliders acts by showing to the user characteristics of the parameter space. In addition, the analyst can easily modify what statistics should be considered for the guidance, enforcing R4. The guidance degree can be framed as orienting support (Q6).

The designer considered an additional, more advanced, solution for supporting parameter selection. One of the problems left open is in fact the selection of parameters at the beginning of the analysis, when the statisticians have no indication which parameter(s) to select first. This solution considers the domain knowledge, to directly suggest possible parameter combinations, depending on the data domain. By exploiting the metadata and the provenance information of the input signals, it is usually possible to automatically trace the data back to a specific domain (Q8). In such a way, besides the orienting support that comes from the integration of hints in the sliders, the statisticians can be directly guided towards parameters that make more sense in certain scenarios. Such guidance can be framed as **directing** support (Q6) and it was appreciated by the statisticians especially to initiate the analysis.

Once the parameters have been chosen, the system launches the BSS algorithms which in average may take some seconds to produce the output, i.e., the signals composing the original measurements. The second part of the workflow is dedicated to compare, explore, and interpret the obtained results. The analysis cycle can then be repeated and better parameters be chosen.

The last operation the designer aims to support by means of guidance is the interpretation and comparison of the obtained results. At this point of the analysis, analysts must discern whether the results obtained make sense, also in light of the previous parameters combination and different output statistics. To support this task, the guidance mechanism makes use of a machine learning algorithm (the random forest algorithm is used) (Q9). It automatically classifies, groups together and presents to the users the results, which are usually composed of an unordered set of components. This is visualized in Figure 4.7b

(Q8). Thanks to this support, the analyst is immediately presented with a reasonable classification of the results that helps him/her to interpret the output. At the same time, the designer provided interaction facilities and visualizations (superimposition of outcomes of different runs as well as the parameters associated with them) which allow the statisticians to easily compare the obtained signals with those of previous runs. To this end, the system stores the history of interaction and output results. This support can be framed as directing guidance (Q6).

Step 4: Guidance Feedback Loop The last part of the design dealt with the definition of feedback mechanisms to let the statisticians steer the guidance process. The described mechanisms make assumptions, e.g., on the domain and on the results classification, that the analyst might need to fine tune and refine. Therefore, when providing guidance, the designer also decided to offer means to modify parameters of the guidance. In general, the system stores the produced results and reuses them for future analysis. For instance, if the analysis reaches a positive conclusion, the results of the classification algorithms are added to the knowledge base to improve the classification algorithms. The same happens when a correct data domain is selected and proper parameters are suggested in the first phase of the analysis. All these small details add to the support of adaptive and trustworthy guidance (R3-4).

The design took about one month and a half. For the sake of clarity, we described the resulting design in a linear fashion, although, in this temporal period, the designer iterated many times over all the steps and produced many design alternatives which were presented to the statisticians. The statisticians appreciated some ideas, rejected others which were not in line with their analysis workflow, and in the end settled for the design we described earlier. Finally, we asked the designer to pay attention to possible flaws and gaps in the guidance design framework, but no major issues were raised during the design. This shows how the framework considers carefully all major aspects involved in the design of guidance, and can be considered rather complete in this respect. As it has been written, the design of guidance poses many challenges and requires designers to foresee issues that may arise during the analysis. Our framework helps in this respect as it points designers to consider thoroughly all these aspects in a step-wise process.

This section laid the basis for the evaluation and discussion of the benefits of our framework and showed how the design of guidance could be easily integrated in the VA process. In the following section, we summarize our observations.

4.7 Discussion

Completeness of the Framework. When creating the framework for guidance designers, we took care that all important aspects needed for designing appropriate guidance are included. We cannot guarantee, however, that the output of the design process is complete. One important point we want to raise is, in fact, that our framework requires designers to think carefully, consider, and *foresee* all possible issues and problems that

might occur during the analysis, and subsequently think of possible guidance solutions to overcome such situations. This can require a lot of effort from designers, which might make the design difficult to complete. One possible solution to this problem would be the implementation of learning mechanisms so that the guidance system can learn and improve over time, as new knowledge gaps arise.

Design Feedback and Evaluation. We see the design of guidance as a finite sequence of (reiterated) steps. This iterative process allows designers to keep design risks low. Still, the proposed framework is not an algorithm that can be applied automatically and not a formal procedure that can be followed thoughtlessly. As discussed previously, we see no practical way to guarantee in advance that the output of the process is complete with respect to the analysts' needs. We discuss this risk, as well as others, and show how to minimize the possible risks in Section 4.5.3. However, like with any design framework, proper evaluation of the design can be done only after the implementation of the designed system. What we can support, with our framework, is the consideration of the major design aspects that concur to an effective guidance solution.

The output of the design process could be evaluated in practical use-case scenarios, where the analyst is faced with real analysis problems, thus testing the effectiveness of the designed guidance. Since this evaluation is taking place in a different time frame in respect to the design process, we consider it separate from the design itself. Hence, we did not include it explicitly in our framework.

Integrating Guidance and VA Design. We envision our guidance design framework to become an integral part of any VA design process. Our framework already integrates references to VA design in Step 1. However, it is still unclear how a tight integration can be accomplished. We see this as an important question for future research.

Methodology and Design Procedure. We point out that instead of exploiting a single source to build this framework, we derived best practice of guidance design from multiple sources. Our design framework is based on a careful analysis of existing VA process models and a characterization of guidance functions in related work. Moreover, we enriched our findings by our own research in guidance for applied VA methods. Thus, our design framework represents an integrated best practice of methods, and desirable properties for effective guidance. We believe that this framework will help researchers and practitioners in VA to approach the design of guidance solutions step by step and to consider critical aspects that are easily overlooked otherwise.

4.8 Conclusion and Future Challenges

In this work, we present a design framework and a set of qualitative requirements to guide the design of effective guidance in VA. Conversely to previous research that focused on describing the characteristics of the *process of guiding* [CGM19], our goal has been to describe the *process of designing guidance* (from the designers' perspectives) and to present it as a sequence of steps applicable to a wide variety of analysis scenarios.

To show the usefulness of our framework, we demonstrate its application to the design of three guidance approaches in different application domains and a design walk-through in the domain of statistics. We list the challenges emerging in such scenarios and report how the framework can be used to design guidance solutions to mitigate them.

Finally, although we propose a comprehensive framework to design effective guidance, there are a number of challenges that remain unsolved:

- C1 - *How can we know/evaluate that guidance is effective?* Guidance should be considered effective, if it can solve the knowledge gap of the analyst. A qualitative study with the actual end users of a guidance-enriched VA approach might be a suitable means to shed light on this aspect and should be an integral part of any evaluation of guidance methods.
- C2 - *How can a knowledge gap be inferred by the system?* We propose two general mechanisms, direct and indirect. However, these are rather abstract. Since guidance requires a context dependent solution, there is no general answer to this question. A combination of user modeling and the integration of expert knowledge seems promising.
- C3 - *How can a knowledge gap be conveyed to the system?* The research on mechanisms for the analyst to communicate knowledge gaps is still far away from providing a definite answer to this question. This involves finding ways to encode knowledge and communicate it effectively to a computational system.
- C4 - *How can we infer what degree of guidance is needed?* What degree of guidance is needed depends on the user knowledge, the tasks, and possibly the users' behavior. Again, there is a need for further research in this direction. Similar to the inference of the knowledge gaps, a combination of user modeling and expert knowledge might help in this respect.
- C5 - *What methods exist to generate guidance?* An answer to this question requires the consideration of the task and scenario-specific aspects. In our design scenarios, methods to generate guidance were chosen to reduce design risks, maximize effectiveness and support a seamless integration of guidance in the analysis process.
- C6 - *How can we decide the right moment for guidance?* Being on time is necessary for guidance to be effective. A designer needs to understand the scenario to understand when guidance should be provided: right away, after detecting a knowledge gap or even at a different point in time. It mainly depends on the how critical is the decision the analyst has to take.



Die approbierte gedruckte Originalversion dieser Dissertation ist an der TU Wien Bibliothek verfügbar.
The approved original version of this doctoral thesis is available in print at TU Wien Bibliothek.

Bibliography

- [AES05] R. Amar, J. Eagan, and J. Stasko. „Low-level components of analytic activity in information visualization“. In: *Proc. of IEEE Symposium on Information Visualization*. IEEE. 2005, pp. 111–117.
- [And+17] G. Andrienko et al. „Visual analysis of pressure in football“. In: *Data Mining and Knowledge Discovery* 31.6 (2017), pp. 1793–1839. DOI: 10.1007/s10618-017-0513-2.
- [And+18] N. Andrienko et al. „Viewing Visual Analytics as Model Building“. In: *Computer Graphics Forum* 37.6 (Jan. 2018), pp. 275–299.
- [Beh+18] M. Behrisch et al. „Quality Metrics for Information Visualization“. In: *Computer Graphics Forum* 37.3 (2018), pp. 625–662.
- [Ber+17] J. Bernard et al. „Visual-interactive semi-supervised labeling of human motion capture data“. In: *Electronic Imaging* 2017.1 (2017), pp. 34–45.
- [Ber11] E. Bertini. „Quality Metrics in High-Dimensional Data Visualization: An Overview and Systematization“. In: *IEEE Transactions on Visualization and Computer Graphics* 17.12 (2011), pp. 2203–2212.
- [CBK09] V. Chandola, A. Banerjee, and V. Kumar. „Anomaly Detection: A Survey“. In: *ACM Computer Surveys* 41.3 (2009), pp. 1–58. ISSN: 0360-0300.
- [Cen+17] D. Ceneda et al. „Characterizing Guidance in Visual Analytics“. In: *IEEE Transactions on Visualization and Computer Graphics* 23.1 (Jan. 2017), pp. 111–120. ISSN: 1077-2626. DOI: 10.1109/TVCG.2016.2598468.
- [Cen+18a] D. Ceneda et al. „Guidance or No Guidance? A Decision Tree Can Help“. In: *Proc. of the International Workshop on Visual Analytics (EuroVA)*. Eurographics Digital Library, 2018, 19–23. DOI: 10.2312/eurova.20181107.
- [Cen+18b] D. Ceneda et al. *Guided Visual Exploration of Cyclical Patterns in Time-series*. Visualization in Data Science (VDS at IEEE VIS 2018). 2018.
- [Cen+20] D. Ceneda et al. „Guide Me in Analysis: A Framework for Guidance Designers“. In: *Computer Graphics Forum* 39.6 (2020), pp. 269–288. DOI: 10.1111/cgf.14017.

- [CGM19] D. Ceneda, T. Gschwandtner, and S. Miksch. „A review of guidance approaches in visual data analysis: A multifocal perspective“. In: *Computer Graphics Forum* 38.3 (2019), pp. 861–879.
- [Col+18] C. Collins et al. „Guidance in the human-machine analytics process“. In: *Visual Informatics* 2.3 (2018), pp. 166–180.
- [Die+18] A. Diehl et al. „VisGuides: A forum for discussing visualization guidelines“. In: *Proc. of the Eurographics Conference on Visualization*. Eurographics Association. 2018, pp. 61–65.
- [Dix+04] A. Dix et al. *Human-Computer Interaction*. 3rd. Pearson Education, 2004.
- [Dow99] P. W. Dowrick. „A review of self modeling and related interventions“. In: *Applied and Preventive Psychology* 8.1 (1999), pp. 23–39.
- [End+14] A. Endert et al. „The human is the loop: new directions for visual analytics“. In: *Journal of Intelligent Information Systems* 43.3 (2014), pp. 411–435.
- [Fuj+97] I. Fujishiro et al. „GADGET: Goal-oriented application design guidance for modular visualization environments“. In: *Proc. of IEEE Conference on Visualization*. IEEE, 1997, pp. 245–252. DOI: 10.1109/VISUAL.1997.663889.
- [GST13] S. Gladisch, H. Schumann, and C. Tominski. „Navigation Recommendations for Exploring Hierarchical Graphs“. In: *International Symposium on Visual Computing*. Rethymnon, Crete: Springer, 2013, pp. 36–47. DOI: 10.1007/978-3-642-41939-3_4.
- [GW09] D. Gotz and Z. Wen. „Behavior-driven visualization recommendation“. In: *Proc. of Conference on Intelligent User Interfaces*. ACM, 2009, pp. 315–324. DOI: 10.1145/1502650.1502695.
- [HB05] J. Heer and D. Boyd. „Vizster: Visualizing online social networks“. In: *Proc. of IEEE Symposium on Information Visualization* (2005).
- [HNM15] M. G. Haselton, D. Nettle, and D. R. Murray. „The evolution of cognitive bias“. In: *The Handbook of Evolutionary Psychology* (2015), pp. 1–20.
- [Hor99] E. Horvitz. „Principles of mixed-initiative user interfaces“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM. 1999, pp. 159–166.
- [Joh+05] J. Johansson et al. „Revealing structure within clustered parallel coordinates displays“. In: *Proc. of IEEE Symposium on Information Visualization*. IEEE. 2005, pp. 125–132.
- [Kan+11] S. Kandel et al. „Wrangler: Interactive visual specification of data transformation scripts“. In: *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM. 2011, pp. 3363–3372.
- [Kin10] R. Kincaid. „SignalLens: Focus+Context Applied to Electronic Time Series“. In: *IEEE Transactions on Visualization and Computer Graphics* 16.6 (2010), pp. 900–907.

- [Ko+16] S. Ko et al. „A survey on visual analysis approaches for financial data“. In: *Computer Graphics Forum*. Vol. 35. Wiley Online Library. 2016, pp. 599–617.
- [Lei+18] R. A. Leite et al. „Visual analytics for event detection: Focusing on fraud“. In: *Visual Informatics 2.4* (2018), pp. 198–212.
- [LGM19] R. Leite, T. Gschwandtner, and S. Miksch. „NEVA: Visual Analytics to Identify Fraudulent Networks“. In: *Submitted to Computer Graphics Forum* (2019).
- [Lub+12] M. Luboschik et al. „Heterogeneity-based Guidance for Exploring Multiscale Data in Systems Biology“. In: *Proc. of IEEE Symposium on Biological Data Visualization*. IEEE, 2012, pp. 33–40. DOI: 10.1109/BioVis.2012.6378590.
- [LW+02] A. Liaw, M. Wiener, et al. „Classification and regression by randomForest“. In: *R news* 2.3 (2002), pp. 18–22.
- [MA14] S. Miksch and W. Aigner. „A matter of time: Applying a data–users–tasks design triangle to visual analytics of time-oriented data“. In: *Computers & Graphics* 38 (2014), pp. 286–290. DOI: 10.1016/j.cag.2013.11.002.
- [May+12] T. May et al. „Using Signposts for Navigation in Large Graphs“. In: *Computer Graphics Forum* 31.3pt2 (2012), pp. 985–994. DOI: 10.1111/j.1467-8659.2012.03091.x.
- [Sac+14] D. Sacha et al. „Knowledge generation model for visual analytics“. In: *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), pp. 1604–1613.
- [Sch+13] H.-J. Schulz et al. *Towards a Characterization of Guidance in Visualization*. Poster at IEEE Conference on Information Visualization (InfoVis). Atlanta, USA, 2013.
- [Shn96] B. Shneiderman. „The eyes have it: A task by data type taxonomy for information visualizations“. In: *Proc. of IEEE Symposium on Visual Languages*. IEEE, 1996, pp. 336–343. DOI: 10.1109/VL.1996.545307.
- [Sil91] M. S. Silver. „Decisional guidance for computer-based decision support“. In: *MIS quarterly* (1991), pp. 105–122.
- [Str+12] M. Streit et al. „Model-driven design for the visual analysis of heterogeneous data“. In: *IEEE Transactions on Visualization and Computer Graphics* 18.6 (2012), pp. 998–1010. DOI: 10.1109/TVCG.2011.108.
- [Str+14] M. Streit et al. „Guided visual exploration of genomic stratifications in cancer“. In: *Nature Methods* 11.9 (2014), pp. 884–885. DOI: 10.1038/nmeth.3088.

- [Sus+20] J. Suschnigg et al. „Exploration of Anomalies in Cyclic Multivariate Industrial Time Series Data for Condition Monitoring“. In: *Proceedings of the EDBT/ICDT 2020 Joint Conference Workshops*. CEUR Workshop Proceedings, 2020.
- [TC05] J. Thomas and K. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Computer Society, 2005.
- [Tea14] R. C. Team. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. 2013. 2014.
- [Zha+17] P. Zhao et al. „Advanced correlation-based anomaly detection method for predictive maintenance“. In: *Proc. of IEEE International Conference on Prognostics and Health Management*. June 2017, pp. 78–83.

List of Figures

1.1	The Information Visualization Reference Model by Card et al. describes how data can be transformed into an interactive visualization. Reproduced after [CMS99]	7
1.2	The VA reference model by Keim et al. aims at combining the strengths of computational systems and human perceptual abilities to make the best of visual data analysis. (Figure taken from [Kei+08b], ©2008 Springer. Used with Permission)	9
1.3	The Knowledge Generation model by Sacha et al. describes how knowledge is generated thanks to VA. (Figure taken from [Sac+14b], ©2014 IEEE. Used with Permission)	10
1.4	The design-triangle by Aigner and Miksch describes the fundamental aspects influencing the design of VA solutions. (Figure taken from [MA14], ©2014 Elsevier. Used with Permission)	11
1.5	The nested model for visualization design and evaluation described by Tamara Munzner. (Figure taken from [Mun09], ©2009 IEEE. Used with Permission)	11
1.6	The conceptual model of guidance in VA, discussed in Chapter 2. The model was created extending the van Wijk’s [Wij06] model (in gray) for visualization with guidance-related blocks (blue). As in the original model, the system aspects of guidance are visualized on the left of the scheme while user aspects (U) are on the right. Figure reproduced with permission after [Cen+17a]	15
1.7	Guidance can be characterized in terms of the main aspects: knowledge gap, input and output, as well as guidance degree. (Figure taken from [Cen+17b], ©2017 IEEE. Used with Permission)	17
		197

1.8	<i>Orienting guidance</i> can be provided by encoding the information the user needs by using different <i>perceptual properties</i> . Highlighting (a) can be used to make visible information that could have some interest for the user, for instance data columns (highlighted in red). Changing the layout and using forms for different data subsets may stimulate the user to explore them. (b) Different data subsets are connected with visual ribbons (highlighted in red) to signal interesting relationships. Motion can be a way to convey guidance too. (c) Motion is used to signal that the data under analysis has a specific characteristic. The reader can imagine the blue lines highlighted in Figure (c) moving up and down. (d) In general, the end-goal of orienting guidance is provide suitable <i>suggestions</i> to proceed the analysis. The system suggests interesting graph regions to explore next (the arrow highlighted in red). (Figure taken from [CGM19a], ©2019 Wiley. Used with Permission)	23
1.9	Jürgen Bernard [Ber+12b] provide <i>directing</i> guidance to the <i>data preprocessing step</i> . The approach aims to support the user in choosing the parameters of each preprocessing step. User guidance is derived from user’s <i>direct</i> interaction with the tool. This <i>feedback</i> is also used to refine the system guidance. (Figure taken from [CGM19a], ©2019 Wiley. Used with Permission)	28
1.10	Wongsuphasawat et al. [Won+16] provide <i>directing</i> guidance to the <i>visual mapping</i> step. The approach aims to provides the user with recommendations of suitable visual mappings. Recommendations are ranked based on some objective visual criteria as well as thanks to user’s direct feedback (Figure taken from [CGM19a], ©2019 Wiley. Used with Permission)	29
1.11	Bernard et al. [Ber+17] provide <i>orienting</i> guidance to <i>model building</i> activities. Through the use of unsupervised (top red box) and supervised (bottom red box) methodologies, the approach is able to provide the user with suitable suggestions of labels for different captures of human motion data. (Figure taken from [CGM19a], ©2019 Wiley. Used with Permission)	31
1.12	Gladisch et al. [GST13] provide <i>orienting</i> guidance to support the exploration of large graphs. In particular, the system guides the user towards interesting areas of the graph. (Figure taken from [CGM19a], Courtesy of Christian Tominsky, ©2019 Wiley. Used with Permission)	33
1.13	Guidance is a mixed-initiative process. On the one hand, the user explicitly or implicitly expresses his/her analysis target and a possible knowledge gap that hinders progression, by interacting with the system. On the other hand, the system reacts to the user’s actions and gives cues that help to decide which steps to take to reach the target. (Figure taken from [Cen+18a], ©2018 ACM. Used with Permission)	35

1.14	An example exploratory task users had to solve during the user study we set up. The task can be supported with different guidance. (a) no guidance: a scatter plot shows values of salinity (x-axis) in relation to the change of water depth (y-axis). User have to select and filter data points by themselves. (b) directing guidance: The system highlights interesting data points and filtering options; (c) prescribing guidance: the system provides step-by-step instructions to solve the task (bottom left of the image). (Figure taken from [CGM19b], ©2019 Elsevier. Used with Permission)	43
1.15	Spiral visualization (center) and user interface with sliders (left). The spiral visualizes a subset of 1,825 temperature values from a dataset with more than 25,000 days worth of data. 365 segments per spiral cycle are visualized to emphasize the yearly temperature fluctuation. (Figure taken from [Cen+18b], ©2018 IEEE. Used with Permission)	47
1.16	Guidance suggestions represented in a slider. The statistical output of CSP algorithm is used to delineate areas of interest while the different shades of gray/black encode probabilities. The output of the DFT algorithm is encoded with small triangles pointing to precise points within these areas. A further triangular shape handle, is used to show precisely the currently selected value. In the bottom part of the image, a zoomed view of the selection allows for a fine-grained value selection. (Figure taken from [Cen+18b], ©2018 IEEE. Used with Permission)	48
1.17	Our guided solution for the exploration of cyclical patterns in time-series. A classical spiral plot is enhanced with data-driven guidance mechanisms to support the identification of patterns. We statistically determine cycle lengths that reveal strong patterns and visually indicate these interesting cycle lengths while the user interacts with the slider. (Figure taken from [Cen+18b], ©2018 IEEE. Used with Permission)	49
1.18	The distance to a suggested cycle length candidate is encoded as a glowing ring surrounding the spiral. a) A wide glow indicates that the current configuration is far off from a suggested cycle length. b) A narrow glow suggests that number of segments per cycle matches the cycle length in the data. (Figure taken from [Cen+18b], ©2018 IEEE. Used with Permission)	50
1.19	The procedure of the user study # 2. We asked the participants to perform the same task under different conditions (guidance/no guidance). After each execution we asked questions regarding their trust and confidence, to understand if guidance improved the analysis. (Figure taken from [Cen+18b], ©2018 IEEE. Used with Permission)	52
1.20	The guidance design framework. The framework aims to support the design of effective guidance. We list a set of steps (Step 1–4) as well as quality criteria (R1–R5) that should guide designers during the design process.	56
2.1	Guidance can be characterized in terms of the main aspects: knowledge gap, input and output, as well as guidance degree.	88

2.2	Components of guidance (in blue) attached to van Wijk’s [Wij06] model (in gray). Aspects of guided VA are shown to the left, while user aspects (U) are on the right. Guidance considers the user’s knowledge (or lack thereof) and may build upon various inputs, including data, interaction history, domain conventions, and visualization images. Different degrees of guidance are possible. Orienting uses visual cues to enhance perception. Directing supports exploration by providing alternative options. Prescribing directly operates on the specification. Guidance positively affects the user’s knowledge in a dynamic process that eventually converges to zero knowledge gap.	91
2.3	Orientation by means of visual cues [Lub+12]. (a) The lineplot shows clear spikes among millions of data points. The heterogeneity bands below the plot suggest that there is more to these spikes hidden at higher levels of granularity. (b) Zooming in on one of the spikes in confirms this assumption.	103
2.4	Orientation via signposts [May+12]. Signposts connect a small, yet detailed focus region of a graph to the invisible “context”. They label outbound edges that connect invisible regions along their shortest path.	104
2.5	Model-driven guidance [Str+12]. (a) A domain-specific model is defined in a three-stage process. (b) The model is then utilized to support users during the data analysis.	106
3.1	Interface of the evaluation environment. In the upper right corner (A), a text-box shows the current task. A further text-box, gives indications to interpret the guidance suggestions. On the top left (B), some combo-boxes give the possibility to filter the dataset. On the bottom of the visualization (C), a text-box shows the step-by-step instructions to reach the desired results (in case prescribing guidance was provided). At the center of the visualization (D), the rectangular selection tool is shown.	123
3.2	General structure of the user’s study. We conducted two parallel evaluation sessions. After a short introduction, the participants performed two set of tasks. Each task was followed by a series of questions. Between the two task sessions, the participants had an active learning phase, where they were instructed either in interaction concepts or in domain concepts respectively. Subsequently, the participants completed the second set of tasks. A cross structure was chosen to minimize the learning effects on the participants.	125
3.3	The same domain task is supported with different degrees of guidance. (a) no guidance: a time-series line chart shows the variation of water salinity in different years. (b) directing guidance: possibly interesting data is highlighted to address the analysis (e.g., data representing high salinity values). These periods are signaled in red at the top of the visualization (1); (c) prescribing guidance: step-by-step instructions are presented to the user (2) together with the highlighting of interesting data points (3).	128

3.4	The same exploratory task supported with different guidance. (a) no guidance: a scatter plot shows values of salinity (x-axis) in relation to the change of water depth (y-axis). A widget allow users to filter the dataset in respect to the measuring station. (b) directing guidance: possibly interesting data points and filtering options are highlighted; (c) prescribing guidance: step-by-step instructions and highlighting of correct values as well as the filtering actions to be performed.	131
3.5	A chart representing aggregated values of nutrients (y-axis) according to the measurement station.	133
3.6	Box-plots for H1.1: We report time and correctness performance metrics for novice users (blue tones), according to different guidance degrees (x-axis). Dashed lines encode the total average (exploratory and domain tasks combined; in black) and individual average values for exploratory and domain tasks (in light and darker blue).	146
3.7	Box-plots for H1.1: We report distance metric and the number of steps for novice users (blue tones), according to different guidance degrees (x-axis). Dashed lines encode the total average (exploratory and domain tasks combined; in black) and individual average values for exploratory and domain tasks (in light and darker blue).	147
3.8	Box-plots for H1.2: completion time and total steps of knowledgeable users (green tones). Dashed lines encode the total average (exploratory and domain tasks combined; in black) and individual average values for exploratory and domain tasks (in light and darker green).	148
3.9	Box-plots for H2.1. Guidance influences positively the user's confidence. Confidence was measured on a five-point Likert scale, where 0 encodes no confidence. Novice users are represented with blue tones, and knowledgeable users with green tones. Dashed lines encode the total average (exploratory and domain tasks combined, in black) and individual average values for exploratory and domain tasks.	149
3.10	Box-plots for H2.2. Frustration of novice users (blue tones) and knowledgeable users (green tones), according to different guidance degrees. The level of frustration is ascending: values closer to 0 indicate less frustration. We also report, with dashed lines, the total average frustration (for exploratory and domain tasks combined, in black), as well as the average value for individual domain and exploratory tasks.	150
		201

4.1	Analysis of Soccer Matches [And+17]. We motivate the need of a proper guidance design: (a) Fragments of trajectories of the players and of the ball for a selected time interval of 5 minutes length. Finding coordinated movement behaviours by looking at trajectories is complicated due to the overcrowded visualization. (b) Time series representing ball attributes, for instance, ball status (in play or out of play), ball possession (which team), X-coordinate, and speed. Due to the high frequency of the data, it is difficult to spot patterns at first sights. (c) Time series of attributes of the players. The need to consider multiple attributes makes the discovery of behaviors and patterns complicated.	159
4.2	The guidance design framework. The framework aims to support the design of effective guidance (R0, see Section 4.5.1). We list a set of steps (Step 1–4) as well as quality criteria (R1–R5) that should guide designers during the design process. The arrows going back and forth between the steps illustrate the iterative nature of the design.	162
4.3	Guidance for exploring cyclical patterns [Cen+18b]. Analysts are supported in finding cycles. Suitable cycle length values are encoded in the sliders (see the gray bars on the left-hand side of the image) that control the visualization of patterns. By choosing the suggested values, cycles appear in the visualization.	175
4.4	The SignalLens approach indicates detected time series anomalies by level-of-detail and markers [Kin10]	179
4.5	Detecting anomalies by comparing sensor data correlation matrices of a reference test cycle (known normal) with an unknown test cycle. Larger differences in columns or rows hint at anomalous values, in comparison to the reference cycle. By means of visual pattern-driven exploration of cross patterns deviating sensors can be identified. Four examples of such cross patterns are highlighted in the figure by red circles at the cross’s intersection on the diagonal element.	180
4.6	Overview visualization for finding frauds in a money transactions network [LGM19]. The analytical exploration of the network data is supported by allowing analysts to formulate only semantically relevant queries and make the analysis effective. The money transaction graph is displayed in the middle. The left-hand side of the interface is where the guidance takes place. It shows the building blocks, i.e., predefined query components, that the analyst can use to formulate queries. As the exploration proceeds, the list of building blocks is updated, some of them are removed, so that just meaningful and non-empty queries can be formed.	183

4.7 Guidance-enhanced blind source separation (BSS). We enhance with guidance the task of separating a signal into its basic components. (a) The area above the sliders used to input parameters is used for showing the impact of different parameter choices to users, hence informing them about the possible outcome of a parameter selection. (b) Thanks to classification algorithms the output of the BSS is classified to enhance the exploration and interpretation of the results. 185



Die approbierte gedruckte Originalversion dieser Dissertation ist an der TU Wien Bibliothek verfügbar.
The approved original version of this doctoral thesis is available in print at TU Wien Bibliothek.

List of Tables

1.1	Table summarizing the classification of guidance papers. Columns represent the different aspects we took into consideration, while papers are listed as rows. The rows are sorted according to the guidance objective they support. We considered approaches providing guidance for different objectives : approaches supporting data Transformation , visual Mapping , Parameter setting, Model Visualization , Model Building , Exploration and Knowledge generation. We considered three Guidance degrees : Orienting , Directing , and Prescribing guidance. We also describe the user side of guidance, in particular the guidance Inference : Direct and Indirect actions. Finally, we provide details of the guidance Direction : whether it is Feedback or Feedforward . Empty black circles (o) show approaches that offer this functionality but to a minor extent. (Table taken from [CGM19a], ©2019 Wiley. Used with Permission)	40
1.2	Key findings of the first user study. In this table, <i>additional finding</i> refers to results that were not taken directly from the hypotheses, but inferred from them. Please refer to Section 3.7.4 for the details.	45
1.3	Summary of the questions designers should ask when designing guidance.	63
3.1	Performance Metrics. We recorded these metrics while the participants executed the tasks. The completion time was provided directly by the evaluation environment EvalBench [AHR13]. The others (correctness, distance, steps) were calculated from the interaction logs (see Section 3.7.1).	135
3.2	Indicators of users' mental state. We asked participants to answer some questions regarding their feelings after each task. Each variable was rated on a five-point Likert scale. We then related the users' feelings to the degree of guidance they received, and to their knowledge level.	136
3.3	Key findings. In this table we summarize the results obtained in our study. We provide references to the hypotheses where we discussed these results in more detail. <i>Additional finding</i> refers to results that were not taken directly from the hypotheses, but inferred from them. Please refer to Section 3.7.4 for the details.	143
		205

4.1 Questionnaire summarizing the design of guidance in three application scenarios. The questionnaire is based on the guidance design framework. (1) Guidance for Cyclical Patterns Exploration [Cen+18b]. The knowledge gap refers to the lack of knowledge regarding the length of the cycles in the univariate time series. The questionnaire shows that while many aspects are considered in this guidance design, question Q12 is not fully answered (n.a.). (2) Condition Monitoring and Failure Detection: The focus is on high dimensional multivariate time-series data. Guidance is needed to correctly set the parameters of the algorithms to detect anomalies and correlations across events. (3) Fraud Detection in Financial Systems [LGM19]. Guidance is needed to support the analyst in analysing a financial transaction graph and discerning whether such transactions are frauds or regular money movements. The knowledge gap refers to finding parameters and form non-empty and meaningful queries to the system. Also in this example, the designed guidance is quite comprehensive, as all the questions are answered.