EuroVis Workshop on Visual Analytics (2024)
M. El-Assady and H.-J. Schulz (Editors)

Persistent Interaction: User-Generated Artefacts in Visual Analytics

Ignacio Pérez-Messina[1], Davide Ceneda[1], Victor Schetinger[2] and Silvia Miksch[1]

[1] TU Wien, Institute of Visual Computing & Human-Centered Technology, Austria
[2] FH St. Pölten, Department Medien und Digitale Technologien, Austria

Abstract

While traditional approaches in visual analytics (VA) prioritize insight generation and knowledge discovery, we argue that user-generated artefacts—annotations, model parameters, subset selections, spatializations, and other constructs—constitute a significant outcome of the analytical process. Drawing from theoretical models in VA literature, we introduce persistent interaction as techniques capturing user decisions. These interactions, called operations, provide a formalization of how users attach subjective judgments to datasets, condensing this input into artefacts serving specific purposes within broader workflows. We provide a description and classification of persistent interaction techniques and outcomes, demonstrating their practical implications in VA systems for system design, information transferability, and guidance capabilities.

CCS Concepts

• ***Human-centered computing*** → ***Visualization design and evaluation methods;*** *Visualization theory, concepts and paradigms;*

1. Introduction

Getting insight and generating knowledge from data is commonly cited as the main drive of visual analytics (VA) [SSS*14]. However, it is not the only one. Increasingly, we find approaches that focus on different kinds of outcomes: data labels, model parameter tunings, subset selections, etc., are all goals embedded in the design of VA solutions which do not necessarily have data-centered outcomes. Although there are surveys and theoretic models focusing on particular types of interactive analysis with data-centered outcomes (e.g., supervising model learning [BZSA18], optimization [LDT*20], parameter space exploration [SHB*]), there is yet no discussion of *persistent user-generated artefacts* as a characteristic aspect of VA. As these concepts become more predominant, particularly within mixed-initiative approaches, we believe it is important to draw a distinction between persistent and non-persistent interaction techniques and provide a theoretical common ground.

We frame this aspect of VA as *persistent interaction*, and decompose it into *operations*—the atomic actions—and *artefacts*—the persistent outcome of the process). Any kind of digital constructs created during the analytical process is a potential artefact, provided they are designed for persistence and involve the user's consent and decisions, and so can serve a definite purpose within a workflow that extends beyond the VA session. Persistent interaction allows users and analysts to introduce subjectivity, i.e., personal preferences and expert knowledge, into the VA workflow by capturing their decisions and attaching them to the subset of the data that they relate to, so that later this information can be fed back into future analyses. Unlike exploratory visual analysis, which yields insights for hypothesis generation and validation, artifacts possess a predefined structure embedded in the system's design. In this paper, we ground the concept of persistent interaction into different theoretical models from the VA literature to provide a formal definition of operations and artefacts based on *persistence*, *subjectivity* and *attachment* as their characteristic aspects (Sec. 2). From this, we differentiate five classes of persistent interaction: *selection*, *parameterization*, *spatialization*, *annotation*, *stand-alone construct*; and characterize them by their common dimensions as they appear in the application literature (Sec. 3).

The benefits of understanding persistent interaction and applying it to VA system design are several: (P1) *generalization*, as operations used to support the development of a class of artefacts can be shared across systems; (P2) *transferability*, as information stored through persistent interaction can be reused in following investigations; (P3) *guidance*, as artefacts can be fed back to mixed-initiative systems to improve their guidance capabilities. We elaborate on benefits P1-3 by describing multiple example scenarios where persistent interaction is or can be applied and potential workflows in which the outcome of analysis can be fed back into future analyses (Sec. 4). We elaborate on this and provide a discussion on interesting information theoretic aspects of artefacts, and how they open possible future research directions (Sec. 5).

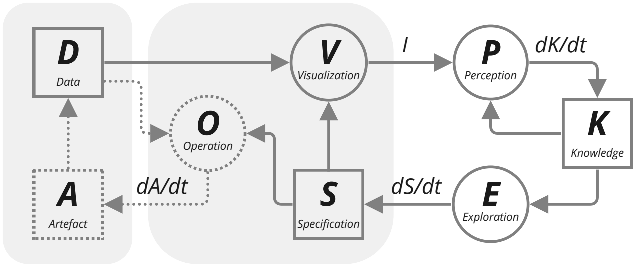

Figure 1: ***Persistent interaction within the simple visualization model.*** *The simple visualization model [vW05] illustrates the process of gaining knowledge from data through visualization V and user interaction E. The data D is processed into an output image I through V. S stands for any and all specification influencing the visualization process. The user perceives (P) the image I, which over time (dK/dt) contributes to the user's Knowledge K. This is a cyclic time-dependent process (through dt) where the user's exploration (considering the knowledge K) feeds back into the specification (dS/dt). We adapted the model to showcase the components of persistent interaction. The Artefact A is the result of the Operation process O attaching a subset of D to a subset of S that captures part of the user's subjectivity (dA/dt). Due to their persistent quality, A can then be considered to be part of D.*

2. Definitions of Persistent Interaction

Dimara and Perin [DP19] characterized interaction in visualization as having a data-related intent (i.e., a necessary but softer intention of the user towards data than in formal data analysis), and classified *data actions* into *input* and *processing* data actions, emphasizing the need to provide the user a range of allowances to act on the data. However, the difference has not been articulated yet, whether data actions are simple allowances to reach a certain visualization state or are functional to the generation of persistent data products from user decisions. In the paradigm of semantic interaction [EFN12], the difference has been conceptualized between "soft" and "hard" data, where the former stands for the stored user interaction data as interpreted by a model. This definition considers softness (i.e., subjectivity), storage (i.e., persistence), and model interpretation as aspects of user-generated artefacts, but it fails to acknowledge that soft data needs hard data to become information (i.e., data processed to be useful [CEH*09]).

In their multi-level typology of abstract visualization tasks, Brehmer and Munzner introduce the `produce` category "in reference to tasks in which the intent is to generate new artefacts" [BM13, p.2378], usually of persistent nature. The `produce` task has a special place in the user task typology, as it pertains to and stems directly from the *why* dimension. Dimara and Stasko proposed that a category, symmetrical to produce, should exist for `decision making` tasks [DS21] in this typology. We believe that certain decision making takes place within workflows that produce a special kind of artefacts, one that effectively captures the user's decisions, as the generation of these artefacts involves every stage of decision-making (intelligence, design, choice) [OCW*23]. However, the importance of `produce` tasks and their role in the generation of persistent user artefacts has, in our opinion, not been fully acknowledged yet as, there are no further distinctions for `produce` tasks as for the `consume` tasks. The typology only affords the description of `produce` tasks through a combination of `manipulate` (methods to interact with visual elements) and `introduce` (methods to add new data) tasks, and even so the terminology results in rather ambiguous constructs. The reason for this may be that the typology is tailored towards visualization and not VA, where interactions with persistent effects become more important.

Given these premises, to properly define persistent interaction, we have to introduce a critical distinction at the point where user intent and low-level tasks meet. We must distinguish between general interaction techniques (e.g., pan, zoom, encode, etc.) from *operations* that allow the user to *persistently attach* their *subjective input* to data. These types of operations possess three distinctive qualities that set them apart from other types of actions: **Persistence**, i.e., the effect of these actions is meant to survive a single analysis session; **Subjectivity**, i.e., the effect of these actions captures a judgement by the user which cannot be directly derived from existing data; **Attachment**, i.e., the effect of these operations relates to some part of the data under analysis. We expand on each of these aspects further and then propose a definition for the class of interactions that generate artefacts.

Persistence — As stated by Wall et al., "Interaction is data." [WDC*17]. Although we agree, not all interaction is *meant* to be data. Even if every interaction is stored in a history log for provenance, there is not necessarily an intention of the user for that data to be captured and have future use. We distinguish user artefacts from *provenance* data because the latter encompasses all kinds of data produced with or without a definite purpose during analysis [RESC15] while the former refers specifically to interaction data whose production is *by design* a main output of a system. An artefact is created, i.e., designed, by a user who is conscious and consents to their creation to survive beyond the end of the analysis. For instance, imagine the case of analysts having the task of selecting images that will be later analyzed (i.e., after the main task is concluded) or employed by a different group of experts for a subsequent task, or used to train a model for further classification and analysis. Persistence confronts us with a system dimension that is rarely discussed in the literature, i.e., how the use of user-produced artefacts allows for knowledge *transfer* which can enrich and inform future uses of systems in time.

Subjectivity — The role of subjective human judgment in visual analytics is often discussed in the context of decision-making and optimization, where multiple objectives compete, or objective measures fail to capture expert knowledge [SHB*]. Visual parameter space exploration serves as a typical example of such scenarios [PBM23]. In these contexts, user input takes various forms, ranging from adjusting variable values with slider movements to importing 3D models for physics simulations. However, not all user inputs in parameter space exploration constitute persistent interaction, as they may not necessarily imply subjective judgment or decision-making by the user. User-subjective inputs shape artefacts and drive the iterative refinement process, whereas non-subjective inputs pertain rather to the data domain to which artefacts attach.

Attachment — User-generated artefacts, within the framework of this work, are not merely an additional dimension introduced by the user (i.e., their subjectivity). Instead, we recognize these artifacts as a fusion of the user's subjective judgement and the original data to which it is imparted upon. This attachment is crucial for providing meaning to the information that the resulting artefact stores. The original data serves as the contextual anchor for the newly introduced subjective judgement, which would otherwise be meaningless. Therefore, persistent interactions consider not only what data is being *captured* but what is it being *attached* to.

Artefacts are usually created in an iterative and sometimes exploratory way, by trial-and-error, by interacting and receiving meaningful visual feedback. To further clarify this concept and differentiate persistent interaction from other interactions, we introduce a formalism to describe an *operation*.

Definition 1: Operation. *An operation function o is formally defined as a mapping from a data object D_m of dimension m, combined with user interaction and interpretation represented by S, to a new data object D_{m+n}, thus:*

$$o : (D_m, S) \mapsto D_{m+n} \quad (1)$$

where n signifies the dimensionality introduced by the user's subjective judgement. The resulting data object D_{m+n} encapsulates the extended dimensionality that includes both the original data and the user's subjective input.

In Fig. 1 the place of artefacts and operations within the simple visualization model [vW05] is presented. The model shows Artefacts A as the result of Operations O attaching a subset of Data D to a subset of Specification S that captures part of the user's Exploration E. Operations are thus a process that separates, captures, and processes the stream of change in specification dS/dt coming from the user's evolving knowledge K. Due to their persistent quality, artefacts can then be considered to be part of the data.

3. Persistent Interaction in VA Systems

Without attempting a taxonomic classification of persistent user-generated artefacts, we describe various classes found in existing visual analytics literature. There are several questions we can ask about an artefact, deriving from Def. 1: (1) what data objects does it attach to (what is the nature of D_m)?; (2) what are the interaction techniques and models involved in its production (what is the nature of the mapping $\mapsto$)?; and (3) what data objects is it made of (what is the nature of n in D_{m+n})? We observed that the answers to these questions are often correlated, leading us to identify five primary classes of persistent interactions: *selections*, *parameterizations*, *spatializations*, *annotations*, and *stand-alone constructs*. Although their naming may refer back to regular interaction techniques, we are naming an artefact-operation class.

The systems and approaches featured here as examples have been mainly sampled from surveys of VA systems for decision making [OCW*23], parameter space exploration [PBM23] and guidance-enhanced approaches [CGM19]. While systems featuring persistent artifacts were less common, those identified demonstrated properties of subjectivity and attachment in user input. In the following, we provide descriptions of each class.

Selection — A selection represents one of the simplest classes of artifacts within persistent interaction, with its datatype abstracted as an array of boolean values indicating selected elements from a dataset, effectively acting as a mask. When the dataset is of moderate size, a selection can be generated by the user by "direct manipulation", i.e., manually adding or removing each data point from the selection. Selections can also be generated on derived datasets, e.g., by selecting principal components after a principal component analysis. When datasets become too massive to handle directly, a model may be of help to "extend" a selection operation from a single element to many in a latent dimension neighbourhood, for example. Depending on the task, selection artefacts can be partially transferred between overlapping datasets, as the selection contains subjective information for each data point independently. Despite selection being a common task, persistent selections and thus selection artefacts are not found widespread in the literature. One notable example is found in Pérez-Messina et al. [PMCM23], a system for unexploded ordnance (UXO) detection by where domain experts select and refine a subset of images, which will be used in the next phase of analysis, while exploring the effects their decisions have on time and spatial coverage and overall quality of the selection.

Parameterization — Parameterization artefacts are usually numerical values (e.g., weights) attached to variables or dimensions of the data. This operation can be performed to assign subjective importance to data attributes to tune a visual or algorithmic model. The resulting artefacts are not specifically tied to a dataset, but instead to a task and a model, as the task defines which parameter value is the most suitable and the model which parameters are allowed. Parameter space analysis approaches are probably the most predominant artefact-producing systems in VA literature, as they by definition consider user input for exploration and generation of results, usually through a simulation model [SHB*]. In spite of that, that the user input is meant to persist is rare in the design of actual systems [PBM23][p.20]. There are several ways in which a parameterization can be operated, for example, by dragging boundaries between nodes in an icicle plot [SOL*15]. Some generative design systems, e.g., of building layouts [BYMW13], allow users to directly input parameter values, while others allow users to implicitly refine parameters while picking generated design options, e.g., of yacht hulls [KGS19].

Spatialization — Spatialization is an interaction technique introduced by Endert et al. [EFN12] where users directly manipulate data point positions in visual space to express their subjective perception of relations between elements, while a model tries to fit itself to these anchor points. The output artefact is twofold: one part attaches a numerical value to pre-existing data point coordinates while the other is the model's interpretation of those coordinates to create a parameterization. Thus, spatialization engenders two artefacts bound together, one "natural" and one "derived". The former artefact is not transferable between contexts but expresses the user's subjective model of a certain dataset, meaning the significance of the spatialization depends on the whole configuration of the dataset (i.e., on every other data point), so they are strictly associated to a particular instance of a dataset. The latter has, on the other hand, the property of being transferable as it is able to "generalize the user's intention" from one dataset to a set of model parameters. System's featuring spatialization are ForceSpire [EFN12] for topic model analysis, Dis-function [BLBC12] for scatterplots, and Podium [WDC*17] for rankings, and El-Assady et al.'s for trees [EASD*18]. In the latter, while points are placed sequentially as leaves in a progressively generated hierarchic structure by an iterative algorithm, the user is allowed to move the data points in the linear space between them, attaching them to other branches. This produces a new parameterization of the model and the difference in time can be visualized through speculative execution [SBS*18]. Speculative execution works in this case as the force-directed simulation in ForceSpire that gives feedback to the user after each operation about the model's change.

Annotation — An annotation artefact is text-based data that is attached to arbitrary subsets of a dataset. Annotations come in the form of labels, tags, notes, or comments introduced by the user to give context and record thought processes into data structures. Any data point or subset can have none, one, or more labels associated. Human-centered approaches frequently include annotation as a feature, as annotation artefacts are used to comment on and label datasets in a human-readable format. Annotations are important artefacts that can become integral parts of end products of a workflow, e.g., stories in investigative journalism [BISM14]. Mixed-initiative approaches go as far as considering an annotation attached to a document as part of it, adding new entities to it, or weighting them accordingly, e.g., ForceSpire [EFN*11]. Other mixed-initiative approaches feature semi-automatic annotation at different stages of the VA process [SGSR21].

Stand-Alone Constructs — There is a wide variety of VA approaches that allow users to create rather complex artefacts which are not attached to data (e.g., placing luminaries in a 3D model of a work environment [WSL*20]), or that attach, in the opposite direction, data *to* artefacts (e.g., in the design of idiosincratic timelines [OBCT23, FBM15]). In task-typological terms, their product is not the outcome of a *search* task [BM13]. We call these artefacts *stand-alone* (as their creation does not suppose the existence of external data) *constructs* (as they are made through "constructive operations" that can build upon each other), making this class a borderline case for our definition of artefacts. Outside VA, text editing, drawing and 3D modelling software fall into the category of systems that support stand-alone construct creation. Within VA, these artefacts are found in many domains such as flood simulation [WKS*14], industrial design [UIM15], and graphic design [DTSO20, KSG20], just to name a few.

4. Persistent Interaction in Practice

In this section, we show how to apply the concepts of persistent interaction for different purposes, giving an example of each through different VA approaches from the literature presented in Sec. 3.

(P1) Generalization — Artefacts can help to generalize interaction techniques. In particular, operations that produce artefacts of a higher number of dimensions (of n) can generalize to equal or lower dimensions, i.e., interaction techniques used to produce a certain type of artefact are mostly interchangeable and can also be used to produce "simpler" artefacts. For example, interaction techniques for spatialization are mostly shared across systems that deal with data of different dimensions: a ranking system [WDC*17] and a hierarchic topic-modeling system [EASD*18] allow the user to operate on data by visual transposition and derive a model state from it. A hypothetical system featuring persistent interaction for spatialization applied for generating selections instead, could allow a user to select a data point in a force-directed layout by dragging it to a "selection area", then fit its model to the differentiating attributes of the point and re-run the layout algorithm making other points get closer to the selection area. This interaction would output a spatialization artefact from which a selection is derived: by changing the operation "selecting" to "positioning freely over a continuous space of decision" we can enhance a simple selection artefact with a richer subjective valuation, producing implicitly a parameterization artefact (i.e., a model of user preferences). Such a generalization can be performed because the spatialization artefact has more dimensions than a selection artefact. If we start the exercise only with selection operations we cannot arrive at a spatialization type artefact as the former does not allow enough degrees of freedom.

(P2) Transferability — A property of artefacts is the ability to transfer information to analytical sessions taking place in the future. However, not all artefacts are equal in this sense. Different artefact classes we have identified have different dimensions in which they hold transferability. We distinguish three of them.

Dataset-transferability. Parameterizations can be transferred between datasets, as the parameters created for one dataset could be used to analyze other datasets with the same attributes. Also, parameters used in generative design can act as a seed for design variations or as a starting point for new explorations. In the ranking system Podium [WDC*17], for example, a parameterization artefact created for analyzing a dataset of US universities could be reused in a similar dataset of international universities or other kinds of institutions. Spatializations are not entirely transferable, as the user positioning of data points is attached to a whole dataset (as positions are always relative to other data points and encode their subjective relations), but the model parameterization implicitly derived from the spatialization holds the property of dataset-transferability.

User-transferability. Annotations can transfer between users as they can be read and understood by other users, and persistence of annotations is usually already assured in systems that feature them, otherwise it would defeat their very purpose. This does not hold useful only for different users, such as in collaborative VA settings [MHK*19], but also for a single user conducting an investigation spanning weeks or returning to an old project. Persistence of annotations is usually assured in systems that feature them, otherwise it would defeat their very purpose.

Task-transferability. Selections can transfer between tasks as they can provide useful information about the data in different scenarios. In image selection for UXO detection [PMCM23], for instance, there are several post-session implications of a new selection: (1) a selected image needs to be acquired if it is not yet, meaning its status in its metadata must be updated, which will change its visual encoding for future analyses; (2) an image that has been selected has been visually validated by a human expert, meaning it is at least of reasonably good quality and it is not damaged or cloud-obstructed; (3) a group of images are selected to work together in time and space, implying they have an affinity for each other. It is possible that a different task performed over the same (or an overlapping) dataset will have a use for stored selection artefacts.

(P3) Guidance — Persistent interactions can significantly impact guidance-enhanced systems, yet many existing systems do not leverage the potential of persistent user artifacts for enhancing their guidance models. These artifacts contain valuable information for future analyses, which can be visually encoded or automatically provided as guidance. For instance, in historical image selection for UXO detection [PMCM23], a guidance model supports users in their task by highlighting possibly useful images and recommending targets for selection. The heuristic optimization model which enables the guidance has been carefully crafted from expert knowledge; however, it misses out on maybe the most predictive dimensions of the data that can only be obtained through the use of persistent artefacts: the information about which images have been previously selected and which pairs have been selected together. This information could be leveraged to inform the interest function and provide automatic suggestions for image pairing. However, it is important to note that the concern that guidance may have a disruptive effect on analysis by constraining user action [PMCEA*22] attains directly to artefacts, as mixed-initiative systems may suggest and enact operations. This can reduce the subjective content put forth in analysis, diminishing the new information that artefacts contribute to the improvement of the guidance model, but ultimately the users holds responsibility over the artefacts they create, even with guidance support, and so subjective judgement is still an aspect of mixed-initiative persistent interaction.

5. Discussion and Future Research Directions

The concepts behind persistent interaction are not entirely novel. This type of architecture (which captures traces of user subjectivity through their interactions and records them) has been fundamentally related to knowledge-based VA [CEH*09]. However, recent work has tended to treat this kind of data as secondary metadata, or essentially analytics that could be used to improve the users' performance in their tasks, hopefully in a passive, automated, and non-invasive way. We believe artefacts should be a core aspect of VA system design, and that the role of artefacts needs to be further explored and researched.

Artefacts add a historical dimension to a system by capturing (Fig.1, arrow from O to A) traces of user subjectivity in time (dS/dt) in relation to the data (a subset of D relevant for the operation O). We can also glimpse their potential value by serving us of information theoretic arguments. In their seminal work, Chen and Golan [CG15] use the Data Processing Inequality (DPI) from Information Theory (IT) to show that subsequent transformations on any data source can only reduce the total "useful" information (the fuel for the dK/dt engine, the evolution of the user's knowledge), which is the mutual information (MU) with the phenomena that generated the data [MJ10]. In information theory, MU is a formal metric that, in broad terms, captures how much of their "essence" two distributions share, and is one of the best approximations VA has for measuring "insight": a good visualization will maximize the mutual information of an user's decision with the original data, i.e. making it as informed as possible. It is also sometimes defined as the reduction in uncertainty about a variable X that can be gained from observing another variable Y. However, the DPI says the information content of any data can only diminish over sequential applications of algorithms and visualisation techniques. To break these (information) bounds a human in the loop is needed to plug in their mutual information (K) to the system [MJ10]. From an information theoretic perspective, then, because artefacts share mutual information with both the data and the human in the loop (Fig.1 subsets of D and S being joined to A through O), they can effectively increase the information bounds (the capacity for being informative, in IT terms) of the system. That is, going against the DPI by enriching the original data (with its fixed information) with subjective data.

An important assumption is that this subjectivity is going to be useful, i.e., effectively making future decisions better informed. This cannot be guaranteed for all cases, and this is why we stress the importance of the design dimension of persistent interaction within VA systems. Picture, for example, a system that gets flooded by artefacts from novice or malicious users, adding persistence to mistakes. The more causal power persistent interaction has on future decisions, the more their life-cycle should be considered. An IT optimization model for a VA system with persistent interaction (as in Chen and Golan [CG15]) needs to account for this historical dimension too, posing several questions: (1) *What is being maximized or learned over time in such kind of system?* (2) *Can the design of a system be robust to future changes of tasks, of scope?* (3) *How to model the weight of different evidence into future decision?* We believe these questions to be interesting venues of investigation within VA, and they are connected to a larger scope of challenges such as data reuse, longevity of systems, and most importantly modeling uncertainty.

6. Conclusion

In this paper, we have introduced a model of persistent interaction within the framework of VA, defining the resulting artefacts as products of operations that assure three key characteristics: persistence, capturing user subjectivity, and data attachment. Through our exposition of the literature, we have aimed to distinguish persistent interaction from a mere recording of user interaction, putting forward a perspective that, we hope, proves useful for VA theory.

Acknowledgements. This work was funded by Vienna Science and Technology Fund (WWTF) under grant [10.47379/ICT19047], by the Austrian Science Fund (FWF) with the grant P31419-N31 (KnoVA) and by DoRIAH (FFG Grant #880883).

References

[BISM14] BREHMER M., INGRAM S., STRAY J., MUNZNER T.: Overview: The design, adoption, and analysis of a visual document mining tool for investigative journalists. *IEEE transactions on visualization and computer graphics 20*, 12 (2014), 2271–2280. 4

[BLBC12] BROWN E. T., LIU J., BRODLEY C. E., CHANG R.: Dis-function: Learning distance functions interactively. In *2012 IEEE conference on visual analytics science and technology (VAST)* (2012), IEEE, pp. 83–92. 4

[BM13] BREHMER M., MUNZNER T.: A multi-level typology of abstract visualization tasks. *Transactions in Visualization and Computer Graphics 19*, 12 (2013), 2376–2385. 2, 4

[BYMW13] Bao F., Yan D.-M., Mitra N. J., Wonka P.: Generating and exploring good building layouts. *ACM Trans. Graph. 32*, 4 (2013). URL: https://doi.org/10.1145/2461912.2461977, doi:10.1145/2461912.2461977. 3

[BZSA18] Bernard J., Zeppelzauer M., Sedlmair M., Aigner W.: Vial: a unified process for visual interactive labeling. *The Visual Computer 34* (2018), 1189–1207. 1

[CEH*09] Chen M., Ebert D., Hagen H., Laramee R. S., Van Liere R., Ma K.-L., Ribarsky W., Scheuermann G., Silver D.: Data, Information, and Knowledge in Visualization. *IEEE Computer Graphics and Applications 29*, 1 (Jan. 2009), 12–19. URL: http://ieeexplore.ieee.org/document/4736452/, doi:10.1109/MCG.2009.6. 2, 5

[CG15] Chen M., Golan A.: What may visualization processes optimize? *IEEE transactions on visualization and computer graphics 22*, 12 (2015), 2619–2632. 5

[CGM19] Ceneda D., Gschwandtner T., Miksch S.: A review of guidance approaches in visual data analysis: A multifocal perspective. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 861–879. 3

[DP19] Dimara E., Perin C.: What is interaction for data visualization? *Transactions in Visualization and Computer Graphics 26*, 1 (2019), 119–129. 2

[DS21] Dimara E., Stasko J.: A critical reflection on visualization research: Where do decision making tasks hide? *Transactions in Visualization and Computer Graphics* (2021). 2

[DTSO20] Dayama N. R., Todi K., Saarelainen T., Oulasvirta A.: Grids: Interactive layout design with integer programming. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (New York, 2020), CHI '20, ACM, p. 1–13. URL: https://doi.org/10.1145/3313831.3376553, doi:10.1145/3313831.3376553. 4

[EASD*18] El-Assady M., Sperrle F., Deussen O., Keim D., Collins C.: Visual analytics for topic model optimization based on user-steerable speculative execution. *Transactions in Visualization and Computer Graphics 25*, 1 (2018), 374–384. 4

[EFN*11] Endert A., Fiaux P., North C., et al.: Unifying the sensemaking loop with semantic interaction. In *IEEE Workshop on Interactive Visual Text Analytics for Decision Making at VisWeek 2011* (2011). 4

[EFN12] Endert A., Fiaux P., North C.: Semantic interaction for visual text analytics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2012), pp. 473–482. 2, 3, 4

[FBM15] Fulda J., Brehmer M., Munzner T.: Timelinecurator: Interactive authoring of visual timelines from unstructured text. *IEEE transactions on visualization and computer graphics 22*, 1 (2015), 300–309. 4

[KGS19] Khan S., Gunpinar E., Sener B.: Genyacht: An interactive generative design system for computer-aided yacht hull design. *Ocean Engineering 191* (2019), 106462. doi:10.1016/j.oceaneng.2019.106462. 3

[KSG20] Koyama Y., Sato I., Goto M.: Sequential gallery for interactive visual design optimization. *ACM Transactions on Graphics (TOG) 39*, 4 (2020), 88–1. doi:10.1145/3386569.3392444. 4

[LDT*20] Liu J., Dwyer T., Tack G., Gratzl S., Marriott K.: Supporting the problem-solving loop: Designing highly interactive optimisation systems. *IEEE Transactions on Visualization and Computer Graphics 27*, 2 (2020), 1764–1774. 1

[MHK*19] Mathisen A., Horak T., Klokmose C. N., Grønbæk K., Elmqvist N.: Insideinsights: Integrating data-driven reporting in collaborative visual analytics. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 649–661. 4

[MJ10] Min Chen, Jäenicke H.: An Information-theoretic Framework for Visualization. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (Nov. 2010), 1206–1215. URL: http://ieeexplore.ieee.org/document/5613460/, doi:10.1109/TVCG.2010.132. 5

[OBCT23] Offenwanger A., Brehmer M., Chevalier F., Tsandilas T.: Timesplines: Sketch-based authoring of flexible and idiosyncratic timelines. *IEEE Transactions on Visualization and Computer Graphics* (2023). 4

[OCW*23] Oral E., Chawla R., Wijkstra M., Mahyar N., Dimara E.: From information to choice: A critical inquiry into visualization tools for decision making. *IEEE Transactions on Visualization and Computer Graphics* (2023). 2, 3

[PBM23] Piccolotto N., Bögl M., Miksch S.: Visual parameter space exploration in time and space. In *Computer Graphics Forum* (2023), Wiley Online Library. 2, 3

[PMCEA*22] Pérez-Messina I., Ceneda D., El-Assady M., Miksch S., Sperrle F.: A typology of guidance tasks in mixed-initiative visual analytics environments. In *Computer Graphics Forum* (2022), vol. 41, Wiley Online Library, pp. 465–476. 5

[PMCM23] Pérez-Messina I., Ceneda D., Miksch S.: Guided visual analytics for image selection in time and space. *IEEE Transactions on Visualization and Computer Graphics* (2023). 3, 4, 5

[RESC15] Ragan E. D., Endert A., Sanyal J., Chen J.: Characterizing provenance in visualization and data analysis: an organizational framework of provenance types and purposes. *IEEE transactions on visualization and computer graphics 22*, 1 (2015), 31–40. 2

[SBS*18] Sperrle F., Bernard J., Sedlmair M., Keim D. A., El-Assady M.: Speculative Execution for Guided Visual Analytics. In *Workshop for Machine Learning from User Interaction for Visualization and Analytics at IEEE VIS* (2018), pp. 1–8. 4

[SGSR21] Schmidt C., Grundel B., Schumann H., Rosenthal P.: Annotations in different steps of visual analytics. In *VISIGRAPP (3: IVAPP)* (2021), pp. 155–163. 4

[SHB*] Sedlmair M., Heinzl C., Bruckner S., Piringer H., Möller T.: Visual Parameter Space Analysis: A Conceptual Framework. *IEEE Transactions on Visualization and Computer Graphics 20*, 12, 2161–2170. doi:10.1109/TVCG.2014.2346321. 1, 2, 3

[SOL*15] Sorger J., Ortner T., Luksch C., Schwärzler M., Gröller E., Piringer H.: Litevis: integrated visualization for simulation-based decision support in lighting design. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (2015), 290–299. 3

[SSS*14] Sacha D., Stoffel A., Stoffel F., Kwon B. C., Ellis G., Keim D. A.: Knowledge generation model for visual analytics. *Transactions in Visualization and Computer Graphics 20*, 12 (2014), 1604–1613. 1

[UIM15] Umentani N., Igarashi T., Mitra N. J.: Guided exploration of physically valid shapes for furniture design. *Communications of the ACM 58*, 9 (2015), 116–124. doi:10.1145/2185520.2185582. 4

[vW05] van Wijk J.: The value of visualization. In *VIS 05. IEEE Visualization, 2005.* (2005), pp. 79–86. doi:10.1109/VISUAL.2005.1532781. 2, 3

[WDC*17] Wall E., Das S., Chawla R., Kalidindi B., Brown E. T., Endert A.: Podium: Ranking data using mixed-initiative visual analytics. *IEEE transactions on visualization and computer graphics 24*, 1 (2017), 288–297. 2, 4

[WKS*14] Waser J., Konev A., Sadransky B., Horváth Z., Ribičić H., Carnecky R., Kluding P., Schindler B.: Many plans: Multidimensional ensembles for visual decision support in flood management. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 281–290. 4

[WSL*20] Walch A., Schwärzler M., Luksch C., Eisemann E., Gschwandtner T.: Lightguider: Guiding interactive lighting design using suggestions, provenance, and quality visualization. *IEEE Transactions on Visualization and Computer Graphics 26*, 1 (2020), 569–578. doi:10.1109/TVCG.2019.2934658. 4