Robot Learning from Humans in Everyday Life Scenarios

Matthias Hirschmanner, Markus Vincze

Abstract

Robots need to be able to learn about novel environments and acquire new capabilities during deployment. Robot learning from humans is a paradigm to enable the human user to teach robots certain information and skills without programming knowledge. In this chapter, we provide an overview of this domain and present some of our work as concrete examples. First, we address grounded language learning with the goal to create connections between words and references (e.g., objects, locations) in social environments. We present our incremental word learning systems using the Pepper robot. Following that, we introduce to learning low-level actions from demonstrations. We present our systems with an industrial robotic arm and a dexterous robotic hand. Then, we address the role of the teacher in the learning process. We investigate the human factors that are important for facilitating the learning process and present the results of our user studies. We conclude with open challenges and opportunities for further research.

Keywords

Robot Learning, Human-Robot Interaction, Learning from Demonstration, Grounded Language Learning

1 Introduction

Robots are increasingly being placed in unconstrained environments, such as homes, where they must adapt to new situations. They cannot be preprogrammed to perform every task with every object in every environment. They need to be able to learn about new tasks with unseen objects in novel environments. Learning from users' input is one way to acquire this knowledge. Examples of information provided by the user could include demonstrating a task or providing language feedback via speech. Robotic learning from humans enables novice users to teach new tasks to a robot without extensive programming knowledge. Therefore, the topic of learning from human teachers has received increased attention in recent years [Ravichandar et al. 2020].

Chernova and Thomaz [2014] motivate learning from humans using robots in the household. Vacuum cleaning robots have become ubiquitous in recent years. They can be placed in an unknown environment and start operating immediately. They can even create a map of the environment to navigate from room to room autonomously. This works well as long as certain constraints are met, such as a flat floor without stairs, cables, or other obstacles.

A general-purpose household robot must complete a much wider and more complex set of tasks. A user would expect it to empty the dishwasher, clean the bathtub, or store objects in their designated storage location. These tasks are not only more complex in terms of manipulation and perception but also need to be performed in less constrained environments. Each household is unique and

https://doi.org/10.34727/2022/isbn.978-3-85448-052-5_8

different from other households. There could be similarities that can be exploited, such as the same type of existing object (e.g., cupboards, drawers, or fridges) or the same type of room (e.g., kitchen, bathroom). However, the storage location of certain objects, such as plates, mugs, or cookie jars, can be unique and arbitrary for each home. These conditions cannot be preprogrammed into the robotic knowledge base in the factory but must be learned by a robot, once it arrives in a new household, similar to a new person moving in. There has to be the possibility for the user to extend the robot's knowledge and modify its behavior. Learning from demonstrations (LfD) methods attempt to learn information, and action policies (i.e., how to perform a task) from examples provided by humans.

Additionally, household robots must be controlled by users directly. A popular and intuitive approach is to use voice commands such as *"Put the strawberry jam into the food storage cabinet."* Modern speech recognition algorithms perform well and can convert a spoken language to text even from a distance, as demonstrated by stationary voice assistants integrated into speakers at users' homes [Berdasco et al. 2019]. A more challenging task is to make sense of what has been said. A robot might not know which object is meant by *"strawberry jam"*, what location by *"food storage cabinet"* and maybe not even how to perform the action *"put"*. "Grounded Language Learning" [Matuszek 2018] is the process of assigning words to references in physical and social spaces. It is a subfield of robotic learning from humans but is often not mentioned in the context of LfD.

Human factors are an important consideration when learning from human teachers. Many papers focus on algorithms for learning policies from demonstrations. The role of human teachers is often overlooked. Especially, novice users cannot be treated as infallible oracles who always provide perfect demonstrations to the robot. Instead, users are part of the learning loop and influence the final performance of the robot immensely. A learning system must consider the human in the loop and accommodate their needs.

The field of "Robotic learning from humans" is very broad, with many different application fields. However, we focus on two domains as an example to provide a starting point for discussing the human factors connected to the learning system. The main contributions of this chapter are:

- We give an introduction to the field of grounded language learning and present our framework with the Pepper robot. It is focused on iterative language learning and being transparent towards the human teacher.

- We discuss the topic of learning low-level actions from human demonstrations and give an overview of recent approaches. We present our setups with an industrial robotic arm and a dexterous robotic hand in simulation.

- We investigate different human factors involved in the learning process such as the teacher's workload, self-efficacy, transparency and trust. We present the results of our experiments with a robot teleoperation setup, a language-learning setup and an interaction scenario.

In Section 2, we provide a brief overview of the field of "grounded language learning," highlighting our two approaches with the Pepper robot. Section 3 discusses "learning of low-level actions" with examples using industrial robotic arms and simulated robotic hands for dexterous manipulation, with a special emphasis on input methods. In Section 4, we discuss the teacher side of the learning loop to identify human factors that must be considered when building learning systems, such as workload, self-efficacy, and trust. Section 5 concludes the paper and mentions opportunities for further research in this field.

2 Grounded Language Learning

Robots are increasingly being used in environments where they must be controlled by untrained nonexpert users. Using one's voice to give commands or communicate intent is a very natural approach in everyday life. Therefore, speech is a very popular modality for giving instructions to robots and has been extensively studied [Matuszek 2018].

Grounded language (also known as situated language) connects the natural language to references in physical and social spaces [Tellex et al. 2020]. For example, the word *"mug"* can be connected to a class of objects, the word *"fridge"* to a storage location different for each home, or the word *"put"* with a series of motor controls dependent on the specific object. The purpose of grounded language learning is to create these word-reference connections.

Many datasets have been introduced because of the various scenarios to which grounded language learning can be applied. An early example was the MARCO dataset [MacMahon et al. 2006], which addresses the problem of navigation instructions. It consists of navigation instructions for a simulated robot (e.g., *"With the wall on your left, walk forward."*). The goal of the system is to understand and follow these instructions with a simulated robot. Other examples of datasets that can be used as starting points for language learning are object detection datasets. They provide natural language class labels for the images. Imagenet has many class labels (e.g., *snail*, *broccoli*, *teapot*) [Deng et al. 2009]. It uses the WordNet [Fellbaum 1998] hierarchy of sets of synonyms that describe meaningful concepts by adding images to each set. Other datasets extend image labels to describe entire images, such as *“a kid sitting on the side walk eating a slice of pizza.”* in the COCO dataset [Lin et al. 2014].

Robots have a various sensors that enable them to use different modalities for grounded language learning such as detected objects, human movements, and recognized actions. Multimodal datasets are used to cover more modalities of real-life scenarios than the above-mentioned. Gaspers et al. [2014] present a dataset where human participants show object manipulation actions to a robot and explain what they are doing. It includes video, audio and human posture data.

We introduced the action verb corpus dataset geared toward object manipulations [Gross et al. 2018], consisting of 390 simple actions (i.e., *take*, *put*, and *push*) of 12 humans following pictured instructions of tasks and describing what they are doing. It includes audio, video and motion data of hand joints and objects. The dataset is annotated with utterance transcriptions, part-of-speech tags, which object is currently moved, and whether a hand touches an object, or an object touches the ground/table. This type of cross-modal and cross-situational data can be used to create systems that learn from humans demonstrating actions while explaining what they are doing. A robot could infer the object name of a manipulated object and the name of the action. The action could be defined by its outcome or by its trajectory. The data can also be used by the robot to replicate the presented action.

Cooccurrence statistics of words and references are often used in computational models that learn from this type of cross-situational data [Krenn et al. 2020]. Taniguchi et al. [2017] provided an overview of different approaches. However, these methods often require large datasets or batches of examples for learning, which is often disadvantageous when deploying a robot in a new environment to learn about new concepts from a human teacher. Additionally, noisy real-world data collected by a robot usually differ from those provided on datasets. Consider a situation where to teach a new concept to a robot, the user must first gather a dataset, which is of course cumbersome and not feasible for a robot at home. However, an incremental learning system uses each new sample to update the probability of a word-reference pair.

We introduced a word-learning system for the Pepper[1] robot, as a concrete example, in Hirschmanner et al. [2018a]. The goal is to learn word-object and word-action mappings in a human-robot interaction scenario. The setup and system architecture are shown in Figure 1. The human teacher demonstrates actions (i.e., *take*, *put*, *push*) to the robot and explains what they are doing. The system infers the type of action from the movement of an object obtained from an object detector processing visual data. The output of the speech recognition module

1 https://www.softbankrobotics.com/emea/en/pepper

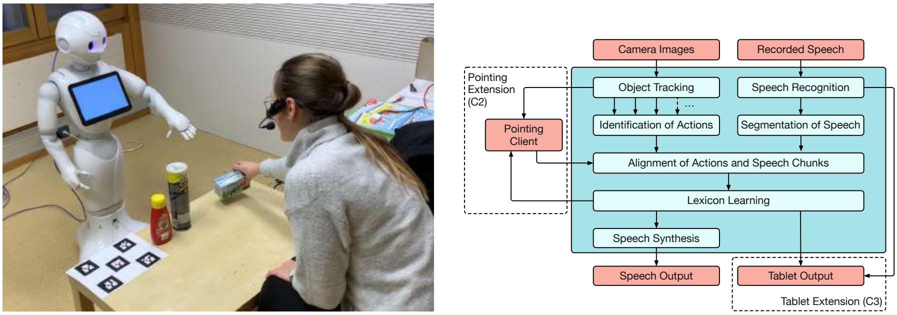

Figure 1 A user performs an object manipulation action (left). Overview of the system architecture (right) used for the language learning system. Tracked objects are used to identify actions and then aligned with the utterances. Normalized pointwise mutual information is used to estimate object/action-word cooccurrence. From Hirschmanner et al. [2021].

is aligned with the action to create utterance-situation pairs. An example of an episode with two utterance-situation pairs would be <*I take the box - ACTION1 OBJECT1*>, <*and put it next to the can. - ACTION2 OBJECT1 OBJECT2*>. We use the normalized pointwise-mutual information (*npmi*), which is a measure of the likelihood of an object/action-word cooccurrence. The *npmi* value is updated after each detected situation-utterance pair. We propose two extensions to this system to increase transparency for the human teacher, in Hirschmanner et al. [2021]. These extensions will be addressed in more detail in Section 4.

The approaches described above usually treat the robot as a passive observer. However, unlike a computer program, a robot is an embodied agent, which can actively request new information by directing the attention of the human teachers toward some unknown references through pointing, gaze, or verbal utterances. This can also be motivated by findings in the developmental psychology of children during language acquisition. They actively request the names of objects using deictic gestures, such as pointing or gaze [Krenn et al. 2019]. A robot can formulate full sentences to acquire knowledge of its surrounding. At public events, we experimented with a Pepper robot that points at objects and formulates questions about the objects pointed at [Hirschmanner et al. 2018b]. The questions did not only refer to the name of an object (i.e., *"How do you call this object?"*) but also to its function (i.e., *"What do you use it for?"*) and the users' preferences (i.e., *"How do you like it?" "What does it mean to you?"*). We used a relatively simple approach that uses part-of-speech tagging to identify nouns, verbs, and adjectives in users' responses. The number of occurrences of each word in these categories is summed up for each object, providing the robot information about the objects modeled using the cultural space model [Schürer et al. 2018]. In this preliminary study, we looked at how human teachers respond to questions from robots.

This section provides a brief introduction to grounded language learning. We want to motivate further development of incremental and active word learning systems for robots, similar to Bisk et al. [2020]. For a general introduction to robots that use language, we refer to Tellex et al. [2020].

3 Learning Low-Level Actions

When deploying a service robot at home, it can already perform certain actions, such as grasping objects and placing them somewhere. In our example of the household robot, the user might give the voice command *"Put the salad bowl into the dishwasher."* Assume that it has already learned which object is meant by *"salad bowl"* and which location by *"dishwasher"* through grounded language learning. There could be a problem in which the robot puts the bowl into an unsatisfactory position or is unable to place the bowl at all. The user will probably know a good strategy for positioning the bowl in the dishwasher. The user can teach the robot the low-level action of placing this specific bowl into the dishwasher using an LfD algorithm.

When creating an LfD system, the following numerous design decisions must be addressed. Which input method is used by the teacher to demonstrate the action? How is the demonstration represented (i.e., which state space is used)? Which algorithm is used to learn the presented demonstration? We give a short overview of the different possibilities to address these design decisions. At the end of the section, we present some concrete projects where we implement learning action policies from human demonstration. We direct the interested readers to Billard et al. [2016] and Chernova and Thomaz [2014] for a general introduction to the topic. A detailed view of the algorithms used in LfD can be found in Osa et al. [2018]. Recent advances are summarized in Ravichandar et al. [2020].

A human teacher can provide demonstrations to a robot in several different ways. Teleoperation is a popular method. The human teacher controls the robot via some device, such as a keyboard, mouse, or joystick to make the robot directly perform the action that is to be learned, which is often cumbersome and difficult to do for novice users [Whitney et al. 2020]. To overcome these limitations, researchers investigated using methods, such as motion tracking to replicate the human motion on the robot [Chernova and Thomaz 2014]. Kinesthetic teaching is an alternative to teleoperation, in which a human manually guides the end-effector of the robot to perform the task [Ravichandar et al. 2020]. For teleoperation or kinesthetic teaching, the sensor data of the robot (e.g., joint angles, end-effector positions, and torques) can be recorded directly and used as input for the machine learning algorithm. We compare kinesthetic teaching to teleoperation on a Pepper robot concerning to the workload on the human teacher in Hirschmanner et al. [2019], which is summarized in Section 4.

Alternatively, some approaches exist that learn directly from observing a human performing an action, making teaching much easier and more natural to the human teacher. The drawback the machine learning problem becomes more difficult because the human movements must be encoded or mapped to the robot's movement [Ravichandar et al. 2020]. Other technical problems may occur if the human performs the task in a way that the robot cannot properly perceive (e.g., fast movements, occlusions, or leaving the field of view).

The next design decision is how to store and process the demonstrations. In this chapter, we will mainly discuss deriving a policy $\pi : \mathcal{S} \to \mathcal{A}$ that maps from a state vector $\boldsymbol{s} \in \mathcal{S}$ to a low-level action $\boldsymbol{a} \in \mathcal{A}$. Other approaches learn policies that output complete trajectories instead of low-level actions. Instead of policies, alternative learning outcomes in LfD can be plans or a reward function for reinforcement learning (i.e., Inverse Reinforcement Learning) [Ravichandar et al. 2020]. The choice of state space $\mathcal{S}$ and action space $\mathcal{A}$ depends on the concrete problem statement. A very simple state space $\mathcal{S}$ may represent the current time, resulting in an open-loop control, where no feedback on the robot or its environment is provided to the policy. Additionally, the robot's sensor data, such as end-effector positions, joint angles, joint velocities, and torques can be used. Sensor data from the environment of the robot can also be included, which can be high-level, such as the pose of an object received by an object pose estimator or low-level, such as a light detection and ranging sensor (LIDAR) or raw camera images.

Similarly, the action space $\mathcal{A}$ can be defined in different ways. Low-level policies could output torques applied to each robot joint. Motion controllers can be used to output actions as end-effector poses or velocities in Cartesian or joint space. Actions can also be defined as trajectories or even sub-tasks as a high-level representation. The choice of granularity of the state and action space depends on the concrete problem, as previously stated. Naturally, the state depends on the available sensors and the teaching approach. For example, when using kinesthetic teaching, using raw camera images as the input might be problematic because the human teacher moving the robot is only present during the demonstration phase. Thus, the image would necessarily be different when the robot executes the action without a teacher. Additionally, a balance should be found between providing enough information to accurately represent the demonstration and not introducing too many dimensions to make the machine learning problem too difficult (“curse of dimensionality” [Bellman 1957]). Similarly, for the action space, a simple representation that can still perform the required task is preferable. For example, for a task involving pushing an object on a table, the two-dimensional (2D) position of the end-effector at a fixed distance to the table might be sufficient. If a device is used to teleoperate the robot, the obvious choice for the action space would be the same domain that is used by the demonstrator, such as the steering angle and acceleration for a remote-controlled car.

A recorded trajectory $\boldsymbol{\tau}$ consists of a state vector $\boldsymbol{s}$ and an action vector $\boldsymbol{a}$ per timestep. The complete demonstration $\mathcal{D}$ can then be defined as

$$\boldsymbol{\tau} = [\boldsymbol{s}_0, \boldsymbol{a}_0, \boldsymbol{s}_1, \boldsymbol{a}_1, \dots \boldsymbol{s}_{T-1}, \boldsymbol{a}_{T-1}, \boldsymbol{s}_T, \boldsymbol{a}_T], \quad \mathcal{D} = \{\boldsymbol{\tau}_i\}_{i=1}^{N}.$$

Training a policy $\pi(\boldsymbol{s}) = \boldsymbol{a}$ from these demonstrations can be seen as a supervised learning problem. Over the years, many different supervised learning approaches have been applied to LfD. Popular approaches include support vector machines (SVM) [Chernova and Veloso 2009], Gaussian mixture models (GMMs) [Khansari-Zadeh and Billard 2011], and Gaussian processes [Choi et al. 2016]. In recent years, artificial neural networks (ANNs) have gained popularity (e.g., Rahmatizadeh et al. [2018]; Zhang et al. [2018]; Young et al. [2021]). There have also been many approaches that address specific problems occurring in LfD. For example, the DAGGER algorithm reduces the number of demonstrations required and, therefore, the load on the human teacher by generating additional demonstrations [Ross et al. 2011].

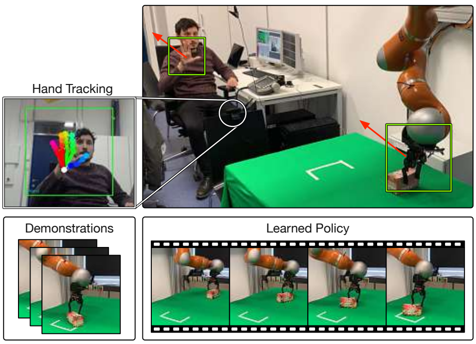

Figure 2 A user teleoperating a Kuka robotic arm using hand tracking to perform a task. The demonstrations are used to learn a policy represented as a neural network. From Hirschmanner et al. [2020].

We present an LfD approach in Hirschmanner et al. [2020], as a concrete example of how to address the different design decisions. We trained a policy on the Kuka LWR IV+ [Bischoff et al. 2010] robot to push a box to a certain position on the

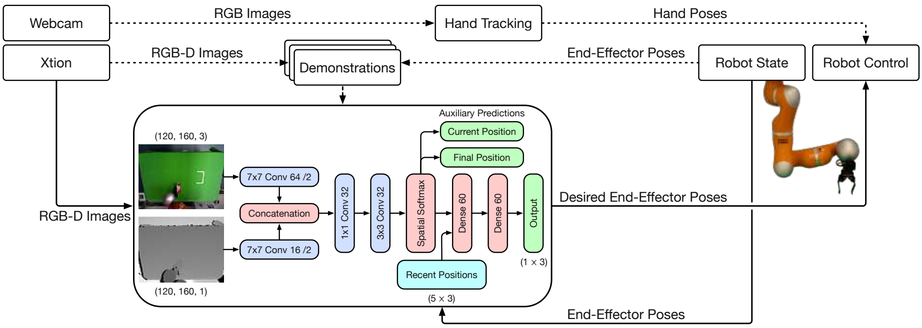

Figure 3 Overview of the system. The dashed lines represent the procedure for collecting demonstrations for training. The continuous lines represent the information flow during policy execution. From Hirschmanner et al. [2020].

table. The demonstrations were recorded using a teleoperation setup based on hand tracking from an RGB webcam. The setup is shown in Figure 2. The state of the robot and its environment are represented as an RGB-D image and the end-effector position of the robot in Cartesian space at the five previous timesteps. For the actions, we use the relative end-effector position $\Delta \boldsymbol{p} \in \mathbb{R}^3$ in Cartesian space. These representations were chosen to capture the entire scene without requiring a separate method to obtain the object pose. The policy is represented as a convolutional neural network (CNN) based on the architecture of Zhang et al. [2018]. It includes two auxiliary tasks during the policy training to predict the current and final end-effector position from the input images. The architecture is shown in Figure 3. We recorded 98 demonstrations at a 10 Hz sampling rate. For the evaluation, we placed the box in different positions on the table, which were unseen during the demonstrations. The robot started to push the box in 86.1% of the trials and reached the goal in 58.3%.

These results indicate some problems with pure supervised learning methods. Demonstrations will not cover each possible configuration in the problem space. During the policy execution, the agent encounters situations unseen during the demonstrations. The situation when the source domain distribution differs from the target domain distribution is referred to as a "covariate shift" [Osa et al. 2018]. Several data-efficient trajectory-learning methods addressed this generalization problem recently. Task-parametrized models of movement [Calinon 2016] use GMMs and represent demonstrations in different frames of reference to improve generalization. Probabilistic movement primitives [Paraschos et al. 2018] represent movement policies in the form of a distribution of trajectories that can be conditioned on desired via-points to adapt to new situations. Kernelized movement primitives [Huang et al. 2019] extend this idea to a nonparametric approach

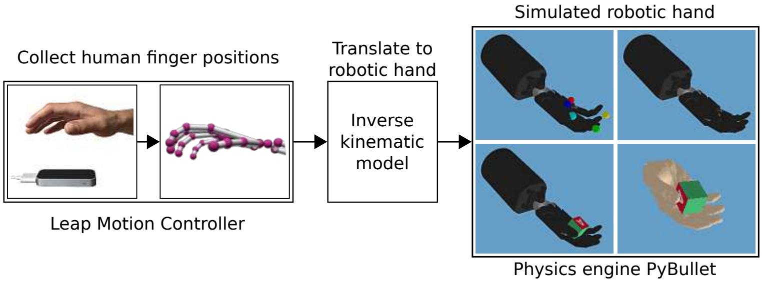

Figure 4 Teleoperation system used to collect the dexterous manipulation tasks. The Leap Motion hand tracker is used to control a simulated robotic or human hand in simulation. The two tasks are shown in the left column of the image on the right. From Zahlner et al. [2020].

geared toward high-dimensional inputs and extrapolation of demonstrated trajectories. One limitation of these trajectory-learning approaches is that the task parameters, such as object poses or obstacles, must be provided to the system when executing the policy, for example, by computer vision algorithms [Pervez et al. 2017]. Additionally, they require a motion planner that converts the trajectory to low-level actions.

One problem with supervised learning is that a learned policy will not outperform the teacher. Researchers have worked on using expert demonstrations in reinforcement learning, as an alternative. In this alternative learning paradigm, the agent can discover new policies through exploration. A reward function is required, which returns a value depending on how beneficial a certain step is to achieve the goal of the task. The machine learning algorithm attempts to maximize the sum of rewards over all timesteps. In the previous box pushing example, this reward function could be the negative distance of the box from the goal. Reinforcement learning is usually very time-intensive because actions that solve a certain task must be discovered through exploration. When expert demonstrations that solve the task are available, this process can be speed-up (e.g., Nair et al. [2020]).

Similarly, we used demonstrations to accelerate the learning process for two dexterous manipulation tasks in Zahlner et al. [2020]. The setup consists of the Shadow Dexterous Hand[2] in the PyBullet[3] simulator. The tasks involved reaching a target position for each fingertip and manipulation of a block to rotate it to a certain orientation. Demonstrations are provided using a teleoperation system that tracks the human hand using a Leap Motion Controller[4] to replicate the current

2 https://www.shadowrobot.com/dexterous-hand-series/

3 https://pybullet.org

hand poses on the simulated hand. The teleoperation system and the different tasks are shown in Figure 4. The state space consists of the absolute angle and velocity of all 20 joints and additional task-specific data. For the reaching task, the current and target Cartesian positions of the fingertips are added to the state. For the object manipulation task, the cube's current and target Cartesian poses, as well as its linear and angular velocities, are provided. The action space of both tasks consists of the 20-dimensional noncoupled hand joints. Both tasks were designed to be similar to the ones presented by Plappert et al. [2018]. We trained the policy with deep deterministic policy gradient (DDPG) [Lillicrap et al. 2016] and hindsight experience replay (HER) [Andrychowicz et al. 2017]. The policy was represented as a neural network. We used demonstrations for pre-training the policy using supervised learning. We saw a speed-up compared to reinforcement learning without pre-training from $2.2 \cdot 10^6$ to $1.2 \cdot 10^6$ timesteps for the reaching task. No comparable speed-up was observed for the cube manipulation task. We hypothesize that this is because the goal in the manipulation task is often reached randomly during exploration and thus does not profit from demonstrations. Additionally, the quality of the demonstrations was low because of the difficulty in manipulating a cube in the simulation without haptic feedback.

Learning the reward function from expert demonstrations is another approach to combining demonstrations and reinforcement learning. This domain is known as inverse reinforcement learning (IRL). The main idea is that the teacher performs demonstrations that optimize an unknown reward function. IRL approaches try to find this reward function. This problem is ill-posed since the expert's behavior could be explained using multiple functions. The retrieved reward function is then used to train a motion policy using standard reinforcement learning algorithms in a subsequent step. Because of to the limited scope of this chapter, we refer to Osa et al. [2018] and Arora and Doshi [2021] for an extensive overview of IRL.

4 Human Factors

In the previous sections, we have addressed how a robot can use the information provided by a human teacher to acquire new skills and knowledge. We did not discuss the influence of the learning process on the user and vice-versa. The human teacher is a part of the learning loop and significantly affects the final performance of the robot. A learning system must consider the humans in the loop and accommodate their needs. However, few studies have been conducted to evaluate the role of the teacher and how the teaching behavior influences a learn-

4 https://www.leapmotion.com

ing system [Sena and Howard 2020]. See Vollmer and Schillingmann [2018] for a recent review.

When designing a system that learns from humans, several factors must be considered. The teaching process must be designed in a way to keep the workload of the user minimal. Low mental and physical workloads lead to higher quality and quantity of training data by keeping the human teacher motivated and engaged [Cui et al. 2021]. The quality of the training data directly affects the learning outcome.

In the context of learning low-level actions, we compared the workload of human demonstrators using a virtual-reality teleoperation setup and kinesthetic guidance in Hirschmanner et al. [2019]. The human teacher wears a virtual reality headset with an attached Leap Motion Controller to teleoperate the Pepper robot shown in Figure 5. The camera stream is displayed in the headset. The current head orientation of the user is imitated by the robot. The hand pose is tracked using the Leap Motion Controller, which is also transferred to the robot. Thus, the robot imitates the upper-body movements of the user. The robot's physical dimensions and constraints are different from those of humans. However, humans can still complete the task successfully because they receive immediate feedback and can adapt to the situation. We compared this setup to kinesthetic guidance, in which users moved the arms of the robot manually. In a user experiment (n=21), participants performed an object grasping task and a pouring task that required controlling both of the robot's arms. Most of the users preferred the teleoperation system for both tasks stating because it was easier to learn. The workload was measured using the NASA-TLX questionnaire [Hart and Staveland 1988]. Compared to kinesthetic guidance, the workload of the users was lower when using teleoperation for the pouring task. We also observed a reduction in task duration for the pouring task when using the teleoperation setup, as an objective measure. Contrary to these results, previous research demonstrated that users preferred kinesthetic guidance to teleoperation [Fischer et al. 2016; Praveena et al. 2019]. This is not contradictory; rather, it emphasizes the importance of tailoring the teaching method to the concrete scenario.

Another important factor that contributes to the workload of the teacher is the number of demonstrations required to train an algorithm. Approaches based on deep learning often require many demonstrations to reach satisfactory performance. Mandlekar et al. [2021] report that 40 demonstrations from a proficient teacher were sufficient to train simple actions, such as lifting an object. For a more complex task, such as transporting a hammer from the workspace of one robot arm to the workspace of another robot arm with a handover operation, the success rate dropped from 72% when using 200 demonstrations to 30.7% when using 40 demonstrations. To overcome the sample inefficiency of deep learn-

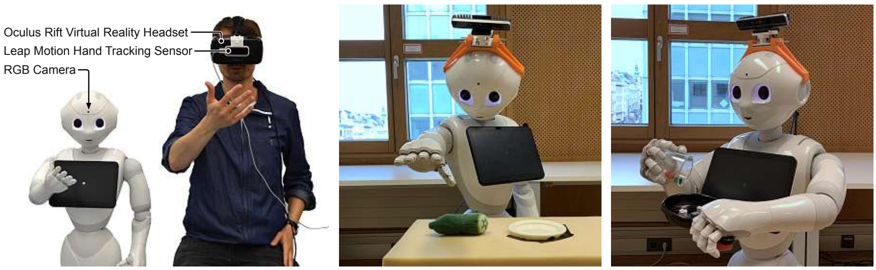

(a) Teleoperation Setup **(b)** Grasping Task **(c)** Pouring Task

Figure 5 The human demonstrator uses the virtual reality teleoperation setup to control the Pepper robot to perform two different tasks. The human head and hand poses are then transferred to the robot. From Hirschmanner et al. [2019].

ing approaches, one can start with a pretrained policy and only ask for human demonstrations if the robot fails (e.g., DelPreto et al. [2020]) or a policy trained for a different task and apply meta-learning with a low number of demonstrations to transfer it to a new task (e.g., Finn et al. [2017]). Algorithms that learn trajectories instead of low-level actions using GMMs or movement primitives (e.g., Calinon [2016]; Paraschos et al. [2018]; Huang et al. [2019]) are designed to require few demonstrations (<10) but require task-parameters, such as object poses.

Additionally, the teacher's mental model of the learning system should align with the actual model to facilitate good teaching behavior of the user [Cakmak and Thomaz 2014]. A robot needs to be able to communicate the current state of the learning system and how the teacher can improve teaching examples to the teacher. These topics are also investigated in the context of transparency in human-robot interaction and explainable artificial intelligence (AI) to increase trust in robots [Papagni and Koeszegi 2021].

Robots, as embodied agents, can expose the current state of the user through various means, such as visualization, movements, text, speech, lights, and imagery [Wallkötter et al. 2021]. A combination of these different modalities is often used. In Hirschmanner et al. [2021] we investigated the efficacy of different modalities. We integrated transparency mechanisms using visualization and deictic gestures in our word-learning system described in Section 2. As a visualization, the Pepper robot displays its current lexicon and the output of the speech recognition system on its screen. The robot uses deictic gestures, such as looking and pointing at objects to either request additional information or to announce the learned word of the object. This behavior is motivated by early-childhood language learning in humans [Krenn et al. 2019]. We did not observe any significant performance difference between the base, visualization, and deictic gestures conditions in a user experiment (n=32). However, the users' knowledge of the system's state positively correlates with the self-reported perception of control and perceived learning success. Users exhibited more interactive behavior when the robot used deictic gestures which might help keep the user engaged, but it also increases noise in the training data. These results encourage further investigation of the transparency mechanism in LfD systems.

Additionally, a learning system should consider factors that influence the user's self-efficacy and perceived control when teaching the system. Self-efficacy is the confidence of a user in being able to perform actions to accomplish a certain task [Bandura 1982]. In the context of a teaching system, self-efficacy is the confidence in being able to teach a new task or concept to a robot. High self-efficacy is important to increase the user's willingness to engage with a robot and to keep them motivated to interact with a learning system in a long-term deployment [Pütten and Bock 2018; Robinson et al. 2020].

The way the robot interacts with the user can influence these factors. We conducted a user experiment (n=29) in Zafari et al. [2019] to study the effect of the interaction style. The task of the user was to build a house of cards. The Pepper robot observed the user and interacted with them using natural language output, such as *"Very nice, keep up the good work."* The speech output was controlled by a researcher following a script. In the person-oriented condition, the robot used motivational sentences to support the user. In the task-oriented condition, the interaction was focused on the task progress and pushed the participant to improve their performance. In the neutral control condition, the robot was only a game instructor and commented on the task progress. We did not tell the participants that they were demonstrating how to build a house of cards to the robot, but the scenario could be used for an LfD system. We found that users in the person-oriented condition reported higher self-efficacy and that they experienced the interaction as less frustrating than in the task-oriented condition. Additionally, participants performed the task significantly longer and thus stayed engaged for a longer time in the person-oriented conditions than in the neutral condition. These results indicate that the interaction style of a robot can also be used to positively influence the human demonstrator and as a consequence, they might be willing to provide more training data in learning from a human setting.

Another important factor to investigate is how trust is influenced in learning from humans setting. A low trust may cause the human teacher to abandon the system. Over-trusting the system may lead to the user ending the teaching process before the system has learned a task reliably and failing to monitor the trained agent, which may result in unwanted behavior or even accidents [Lewis et al. 2018]. DelPreto et al. [2020] found that low accuracy in an LfD task reduces trust and increases the users' workload. They also found tendencies that users overestimate the robot's skills. Hedlund et al. [2021] found that when robots fail to perform the learned tasks, participants' trust in the robot and themselves as teachers decreases.

5 Conclusion and Open Challenges

In this chapter, we presented our work and set it into the context of the field of robotic learning from humans. First, we motivated the need for grounded language learning of social robots, i.e., connecting words with references such as objects. Learning these connections is required for a robot to follow voice commands. We presented two word-learning systems using the Pepper robot. Following that, we addressed the field of learning low-level actions from demonstrations. We covered the main design choices that must be made when developing a learning system. We presented two systems with different robotic setups to demonstrate different design decisions. Furthermore, we discussed the role of the human teacher in the learning system. We emphasized the importance of considering factors, such as workload, self-efficacy, and trust during the teaching process to obtain good training examples and keep the user motivated. We presented three user studies for different robotic setups that investigated workload, transparency, and self-efficacy.

The field of learning from human users is emerging, as more robots move into living spaces. There are still many open problems to be tackled. Robots need to be able to acquire information from spoken language to make interactions with humans more natural. Grounded language learning methods that can incrementally process the high-dimensional multimodal data that robots will encounter in everyday situations must be developed [Bisk et al. 2020]. Additionally to spoken language, they need to be able to understand nonverbal communication to better interact with humans.

Learning action policies from demonstrations has accelerated in recent years [Ravichandar et al. 2020]. Many algorithms have been developed to address the special conditions and constraints associated with learning from human teachers. However, it is often difficult to compare the approaches because of the limited number of available benchmarks using real demonstrations provided by humans that have advanced other fields such as computer vision or reinforcement learning. Two of the few examples are Mandlekar et al. [2021] and [Sharma et al. 2018]. New standardized benchmarking methods on real robotic systems will be required to advance the field of learning motion policies from human demonstrations.

Demonstrations are usually task-specific and do not cover the entire problem space. A promising direction for further research is to develop algorithms that generalize better across tasks, domains, and robots. A combination of demonstrations with reinforcement learning could be useful in this regard and should be examined further. Demonstrations can be used to shorten the long training times of reinforcement learning algorithms. Additionally, these systems often require tedious hyperparameter tuning, which is not feasible for novice users. Further research is required to develop methods that require few hyperparameters and are easy to tune automatically.

The role of the teacher and teaching behavior have been under-represented in the robotic learning from humans pipeline [Vollmer and Schillingmann 2018]. High-quality training data from the human teacher facilitates the learning process. To ensure this, the teacher must be considered as a part of the learning loop when designing a system. Further research should aim to create non-intrusive and intuitive teaching systems to minimize the workload of the user and keep them motivated and engaged.

If we want to deploy robots that learn from humans in users' homes, the effect of the learning system on the users must be studied further. Users will only accept these systems, if they see an added value in them and if they enjoy using them [de Graaf et al. 2017]. We believe that self-efficacy is an important concept in that regard. We must investigate which factors influence the trust of the user in the system to find a balance between not overtrusting the system and trusting it enough to use it continuously.

Bibliography

Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. 2017. Hindsight experience replay. In *Advances in neural information processing systems*. 5048–5058.

Saurabh Arora and Prashant Doshi. 2021. A survey of inverse reinforcement learning: Challenges, methods and progress. *Artificial Intelligence* 297 (2021), 103500. https://doi.org/10.1016/j.artint.2021.103500

Albert Bandura. 1982. Self-efficacy mechanism in human agency. *American psychologist* 37, 2 (1982), 122. Publisher: American Psychological Association.

Richard Bellman. 1957. *Dynamic Programming* (1 ed.). Princeton University Press, Princeton, NJ, USA.

Ana Berdasco, Gustavo López, Ignacio Diaz, Luis Quesada, and Luis A Guerrero. 2019. User experience comparison of intelligent personal assistants: Alexa, Google Assistant, Siri and Cortana. *Multidisciplinary Digital Publishing Institute Proceedings* 31, 1 (2019), 51.

Aude G Billard, Sylvain Calinon, and Rüdiger Dillmann. 2016. Learning from Humans. In *Springer Handbook of Robotics*, Bruno Siciliano and Oussama Khatib (Eds.). Springer International Publishing, Cham, Chapter 74, 1995–2014.

Rainer Bischoff, Johannes Kurth, Günter Schreiber, Ralf Koeppe, Alin Albu-Schäffer, Alexander Beyer, Oliver Eiberger, Sami Haddadin, Andreas Stemmer, Gerhard Grunwald and Gerhard Hirzinger. 2010. The KUKA-DLR Lightweight Robot arm-a new reference platform for robotics research and manufacturing. In *ISR 2010 (41st international symposium on robotics) and ROBOTIK 2010 (6th German conference on robotics)*. VDE, 1–8.

Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020. Experience Grounds Language. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 8718–8735. https://doi.org/10.18653/v1/2020.emnlp-main.703

Maya Cakmak and Andrea L Thomaz. 2014. Eliciting good teaching from humans for machine learners. *Artificial Intelligence* 217 (2014), 198–215.

Sylvain Calinon. 2016. A tutorial on task-parameterized movement learning and retrieval. *Intelligent Service Robotics* 9, 1 (Jan. 2016), 1–29. https://doi.org/10.1007/s11370-015-0187-9

Sonia Chernova and Andrea L Thomaz. 2014. *Robot learning from human teachers.* Morgan & Claypool Publishers.

Sonia Chernova and Manuela Veloso. 2009. Interactive policy learning through confidence-based autonomy. *Journal of Artificial Intelligence Research* 34 (2009), 1–25.

Sungjoon Choi, Kyungjae Lee, and Songhwai Oh. 2016. Robust learning from demonstration using leveraged Gaussian processes and sparse-constrained optimization. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 470–475.

Yuchen Cui, Pallavi Koppol, Henny Admoni, Scott Niekum, Reid Simmons, Aaron Steinfeld, and Tesca Fitzgerald. 2021. Understanding the Relationship between Interactions and Outcomes in Human-in-the-Loop Machine Learning. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence.* International Joint Conferences on Artificial Intelligence Organization, Montreal, Canada, 4382–4391. https://doi.org/10.24963/ijcai.2021/599

Maartje de Graaf, Somaya Ben Allouch, and Jan van Dijk. 2017. Why Do They Refuse to Use My Robot?: Reasons for Non-Use Derived from a Long-Term Home Study. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction. ACM, Vienna Austria*, 224–233. https://doi.org/10.1145/2909824.3020236

Joseph DelPreto, Jeffrey I Lipton, Lindsay Sanneman, Aidan J Fay, Christopher Fourie, Changhyun Choi, and Daniela Rus. 2020. Helping Robots Learn: A Human-Robot Master-Apprentice Model Using Demonstrations via Virtual Reality Teleoperation. In *2020 IEEE International Conference on Robotics and Automation (ICRA).* IEEE, 10226–10233.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. IEEE, 248–255.

Christiane Fellbaum (Ed.). 1998. *WordNet: An Electronic Lexical Database.* MIT Press, Cambridge, MA.

Chelsea Finn, Tianhe Yu, Tianhao Zhang, Pieter Abbeel, and Sergey Levine. 2017. One-shot visual imitation learning via meta-learning. In *Conference on robot learning.* PMLR, 357–368.

Kerstin Fischer, Franziska Kirstein, Lars Christian Jensen, Norbert Krüger, Kamil Kukliński, Maria Vanessa aus der Wieschen, and Thiusius Rajeeth Savarimuthu. 2016. A Comparison of Types of Robot Control for Programming by Demonstration. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction (HRI '16)*. IEEE Press, 213–220. http://dl.acm.org/citation.cfm?id=2906831.2906868

Judith Gaspers, Maximilian Panzner, Andre Lemme, Philipp Cimiano, Katharina J Rohlfing, and Sebastian Wrede. 2014. A multimodal corpus for the evaluation of computational models for (grounded) language acquisition. In *Proceedings of the 5th Workshop on Cognitive Aspects of Computational Language Learning (CogACLL)*. 30–37.

Stephanie Gross, Matthias Hirschmanner, Brigitte Krenn, Friedrich Neubarth, and Michael Zillich. 2018. Action Verb Corpus. In *Proc. 2018 Language Recognition and Evaluation Conference.* Miyazaki, Japan.

Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. *Advances in psychology* 52 (1988), 139–183.

Erin Hedlund, Michael Johnson, and Matthew Gombolay. 2021. The Effects of a Robot's Performance on Human Teachers for Learning from Demonstration Tasks. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction.* 207–215.

Matthias Hirschmanner, Stephanie Gross, Brigitte Krenn, Friedrich Neubarth, Martin Trapp, and Markus Vincze. 2018a. Grounded Word Learning on a Pepper Robot. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents*. ACM, 351–352.

Matthias Hirschmanner, Stephanie Gross, Setareh Zafari, Brigitte Krenn, Friedrich Neubarth, and Markus Vincze. 2021. Investigating Transparency Methods in a Robot Word-Learning System and Their Effects on Human Teaching Behaviors. In *2021 30th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN).* Vancouver, BC, CA (Virtual Conference), 175–182.

Matthias Hirschmanner, Ali Jamadi, Bernhard Neuberger, Tim Patten, and Markus Vincze. 2020. Learning Manipulation Tasks from Vision-based Teleoperation. In *Proceedings of Joint Austrian Computer Vision and Robotics Workshop (ACVRW2020).* 42–47. https://doi.org/DOI:10.3217/ 978-3-85125-752-6-11

Matthias Hirschmanner, Oliver Schürer, Brigitte Krenn, Christoph Müller, Friedrich Neubarth, and Markus Vincze. 2018b. Improving the Quality of Dialogues with Robots for Learning of Object Meaning. In *Workshop on Language and Robotics at IROS2018*.

Matthias Hirschmanner, Christiana Tsiourti, Tim Patten, and Markus Vincze. 2019. Virtual Reality Teleoperation of a Humanoid Robot Using Markerless Human Upper Body Pose Imitation. In *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids).*

Yanlong Huang, Leonel Rozo, João Silvério, and Darwin G Caldwell. 2019. Kernelized movement primitives. *The International Journal of Robotics Research* 38, 7 (June 2019), 833–852. https://doi.org/10.1177/0278364919846363

Seyed Mohammad Khansari-Zadeh and Aude Billard. 2011. Learning stable nonlinear dynamical systems with gaussian mixture models. *IEEE Transactions on Robotics* 27, 5 (2011), 943–957.

Brigitte Krenn, Sepideh Sadeghi, Friedrich Neubarth, Stephanie Gross, Martin Trapp, and Matthias Scheutz. 2020. Models of Cross-Situational and Crossmodal Word Learning in Task-Oriented Scenarios. *IEEE Transactions on Cognitive and Developmental Systems* 12, 3 (2020), 658–668. https://doi.org/10.1109/TCDS.2020.2995045

Brigitte Krenn, Christiana Tsiourti, Friedrich Neubarth, Stephanie Gross, and Matthias Hirschmanner. 2019. Active Language Learning Inspired from Early Childhood Information Seeking Strategies. In *Workshop on Cognitive Architectures for Human-Robot Interaction at AAMAS 2019*.

Michael Lewis, Katia Sycara, and Phillip Walker. 2018. The role of trust in human-robot interaction. In *Foundations of trusted autonomy.* Springer, Cham, 135–159.

Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2016. Continuous control with deep reinforcement learning. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings.*

Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*. Springer, 740–755.

Matt MacMahon, Brian Stankiewicz, and Benjamin Kuipers. 2006. Walk the Talk: Connecting Language, Knowledge, and Action in Route Instructions. In *Proceedings of the 21st National Conference on Artificial Intelligence - Volume 2 (Boston, Massachusetts) (AAAI'06)*. AAAI Press, 1475–1482.

Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. 2021. What Matters in Learning from Offline Human Demonstrations for Robot Manipulation. In *5th Annual Conference on Robot Learning (CoRL)*. https://openreview.net/forum?id=JrsfBJtDFdI

Cynthia Matuszek. 2018. Grounded Language Learning: Where Robotics and NLP Meet.. In *Proceedings of the International Joint Conference on Artificial Intelligence.* 5687–5691.

Suraj Nair, Silvio Savarese, and Chelsea Finn. 2020. Goal-Aware Prediction: Learning to Model What Matters. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119),* Hal Daumé III and Aarti Singh (Eds.). PMLR, Virtual, 7207–7219. http://proceedings.mlr.press/v119/nair20a.html

Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J Andrew Bagnell, Pieter Abbeel, Jan Peters, and others. 2018. An algorithmic perspective on imitation learning. *Foundations and Trends in Robotics* 7, 1-2 (2018), 1–179. Publisher: Now publishers.

Guglielmo Papagni and Sabine Koeszegi. 2021. Understandable and trustworthy explainable robots: A sensemaking perspective. *Paladyn, Journal of Behavioral Robotics* 12, 1 (2021), 13–30. https://doi.org/doi:10.1515/pjbr-2021-0002

Alexandros Paraschos, Christian Daniel, Jan Peters, and Gerhard Neumann. 2018. Using probabilistic movement primitives in robotics. *Autonomous Robots* 42, 3 (March 2018), 529–551. https://doi.org/10.1007/s10514-017-9648-7

Affan Pervez, Yuecheng Mao, and Dongheui Lee. 2017. Learning deep movement primitives using convolutional neural networks. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*. IEEE, Birmingham, 191–197. http://ieeexplore.ieee.org/document/ 8246874/

Matthias Plappert, Marcin Andrychowicz, Alex Ray, Bob McGrew, Bowen Baker, Glenn Powell, Jonas Schneider, Josh Tobin, Maciek Chociej, Peter Welinder, et al. 2018. Multi-goal reinforcement learning: Challenging robotics environments and request for research. *arXiv preprint arXiv:1802.09464* (2018).

Pragathi Praveena, Guru Subramani, Bilge Mutlu, and Michael Gleicher. 2019. Characterizing Input Methods for Human-to-robot Demonstrations. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI).* IEEE, 344–353.

Astrid Rosenthal-Von Der Pütten and Nikolai Bock. 2018. Development and Validation of the Self-Efficacy in Human-Robot-Interaction Scale (SE-HRI). *J. Hum.-Robot Interact.* 7, 3 (Dec. 2018). Place: New York, NY, USA Publisher: Association for Computing Machinery.

Rouhollah Rahmatizadeh, Pooya Abolghasemi, Ladislau Boloni, and Sergey Levine. 2018. Vision-Based Multi-Task Manipulation for Inexpensive Robots Using End-to-End Learning from Demonstration. In *2018 IEEE International Conference on Robotics and Automation (ICRA).* IEEE, Brisbane, QLD, 3758–3765. https://doi.org/10.1109/ICRA.2018.8461076

Harish Ravichandar, Athanasios S. Polydoros, Sonia Chernova, and Aude Billard. 2020. Recent Advances in Robot Learning from Demonstration. *Annual Review of Control, Robotics, and Autonomous Systems* 3, 1 (May 2020), 297–330. https://doi.org/10.1146/annurev-control-100819-063206

Nicole L Robinson, Teah-Neal Hicks, Gavin Suddrey, and David J Kavanagh. 2020. The Robot Self-Efficacy Scale: Robot Self-Efficacy, Likability and Willingness to Interact Increases After a Robot- Delivered Tutorial. In 2*020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, Naples, Italy, 272–277. https://doi.org/10.1109/RO-MAN47096. 2020.9223535

Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics.* 627–635.

Oliver Schürer, Benjamin Stangl, Christoph Hubatschke, and Christoph Müller. 2018. Experiments with a First Prototype of a Spatial Model of Cultural Meaning through Natural-Language Human-Robot Interaction. *Technologies* 6, 1 (Jan. 2018), 6. https://doi.org/10.3390/technologies6010006

Aran Sena and Matthew Howard. 2020. Quantifying teaching behavior in robot learning from demonstration. *The International Journal of Robotics Research* 39, 1 (2020), 54–72. Publisher: SAGE Publications Sage UK: London, England.

Pratyusha Sharma, Lekha Mohan, Lerrel Pinto, and Abhinav Gupta. 2018. Multiple interactions made easy (mime): Large scale demonstrations data for imitation. In *Conference on robot learning.* PMLR, 906–915.

Akira Taniguchi, Tadahiro Taniguchi, and Angelo Cangelosi. 2017. Cross-Situational Learning with Bayesian Generative Models for Multimodal Category and Word Learning in Robots. *Frontiers in Neurorobotics* 11 (Dec. 2017), 66. https://doi.org/10.3389/fnbot.2017.00066

Stefanie Tellex, Nakul Gopalan, Hadas Kress-Gazit, and Cynthia Matuszek. 2020. Robots That Use Language. *Annual Review of Control, Robotics, and Autonomous Systems* 3, 1 (May 2020), 25–55. https://doi.org/10.1146/annurev-control-101119-071628

Anna-Lisa Vollmer and Lars Schillingmann. 2018. On Studying Human Teaching Behavior with Robots: a Review. *Review of Philosophy and Psychology* 9, 4 (Dec. 2018), 863–903. https://doi.org/10.1007/s13164-017-0353-4

Sebastian Wallkötter, Silvia Tulli, Ginevra Castellano, Ana Paiva, and Mohamed Chetouani. 2021. Explainable Embodied Agents Through Social Cues: A Review. *J. Hum.-Robot Interact.* 10, 3, Article 27 (July 2021), 24 pages. https://doi.org/10.1145/3457188

David Whitney, Eric Rosen, Elizabeth Phillips, George Konidaris, and Stefanie Tellex. 2020. Comparing robot grasping teleoperation across desktop and virtual reality with ROS reality. In *Robotics Research*. Springer, 335–350.

Sarah Young, Dhiraj Gandhi, Shubham Tulsiani, Abhinav Gupta, Pieter Abbeel, and Lerrel Pinto. 2021. Visual Imitation Made Easy. In Proceedings of the 2020 Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 155), Jens Kober, Fabio Ramos, and Claire Tomlin (Eds.). PMLR, 1992–2005. https://proceedings.mlr.press/v155/young21a.html 6

Setareh Zafari, Isabel Schwaninger, Matthias Hirschmanner, Christina Schmidbauer, Astrid Weiss, and Sabine T. Koeszegi. 2019. ”You Are Doing so Great!” - The Effect of a Robot’s Interaction Style on Self-Efficacy in HRI. In *2019 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE.

Stefan Zahlner, Matthias Hirschmanner, Timothy Patten, and Markus Vincze. 2020. Teleoperation System for Teaching Dexterous Manipulation. In *Workshop on Hand-Object Interaction: From human demonstrations to robot manipulation at RO-MAN2020.*

Tianhao Zhang, Zoe McCarthy, Owen Jow, Dennis Lee, Ken Goldberg, and Pieter Abbeel. 2018. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In *IEEE International Conference on Robotics and Automation (ICRA).*