From “Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,”
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

15

Learning Task Knowledge

Dario D. Salvucci, John E. Laird, Franklin Chang, Kenneth D. Forbus, Parisa Kordjamshidi, Tom M. Mitchell, Shiwali Mohan, Michael Spranger, Suzanne Stevenson, Andrea Stocco, and J. Gregory Trafton

Abstract

How does an agent acquire (i.e., learn) knowledge and information about a specific task by interacting with a teacher, so that ultimately the agent is able to execute the task successfully? This chapter reviews critical aspects of the learning process in interactive task learning (ITL). It discusses learning task knowledge through interaction, capabilities that facilitate learning, aspects of interaction that relate closely to learning, and evaluation dimensions and metrics for ITL systems. Given the interconnected nature of ITL, it also explores relationships between learning, knowledge, interaction, and tasks: how tasks influence learning, how knowledge should be represented, and what types of information and communication are needed to facilitate learning.

Introduction

Research into interactive task learning (ITL) focuses on how agents (biological and artificial) acquire new tasks through natural interactions with one another within a shared environment. A core component of ITL involves (a) learning new knowledge, (b) integrating it with existing knowledge, and (c) operationalizing that knowledge to perform novel tasks. This process, in turn, is heavily influenced by the very nature of ITL: a learner continually acquires information from both the teacher and the environment, and such interactions have the potential to fundamentally affect what is learned as well as when and how learning takes place.

Group photos (top left to bottom right) Dario Salvucci, John Laird, Parisa Kordjamshidi, Michael Spranger, Suzanne Stevenson, Tom Mitchell, Greg Trafton, Andrea Stocco, Shiwali Mohan, Ken Forbus, Dario Salvucci, Suzanne Stevenson, Franklin Chang, Shiwali Mohan, Greg Trafton, Tom Mitchell, Andrea Stocco, Franklin Chang, Michael Spranger, John Laird, Parisa Kordjamshidi

From “Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,” edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

We begin by reviewing multiple aspects of learning in the context of ITL and explore two questions in depth: What are the characteristics of learning mechanisms most appropriate for ITL? What overall capabilities of the agent are desired (or even essential) to enable and facilitate interactive learning? We then examine the relationship between interaction and the task itself and propose a set of metrics that may be used to understand and evaluate ITL systems with respect to the dimensions of learning, task, and interaction.

Background to Learning Task Knowledge

Because ITL can be applied in a wide range of domains (e.g., real-world robots, agents in virtual environments, video game agents, and virtual personal assistants), learning mechanisms must be able to respond to different levels of complexity and diverse characteristics inherent to each domain. Such variance means that there is no single process of learning to be followed in ITL. The unique characteristics of a domain will determine what needs to be learned, the degree of difficulty of the learning problem, and the appropriateness of a learning algorithm. Consideration must be given to whether the environment is partially or fully observable, whether the environment is discrete or continuous, whether actions are deterministic or stochastic, and whether the environment has complex dynamics. In addition, the extent of uncertainty in sensing and acting in the environment must be evaluated: the more complex and uncertain a domain is, the more difficult it will be for an agent to learn new tasks quickly.

Currently, ITL agents focus primarily on tasks within a single domain, and designers tailor their learning algorithms to that domain. As we look to expand our research into ITL, we need to assess whether there are fundamental representations and learning algorithms that would work across all domains. Do certain domains require specialized learning approaches and, if so, to what degree? Here, neuroscientific studies of the human brain may provide inspiration for general learning algorithms, as the human brain has demonstrated an unrivaled ability to learn across the lifespan and multiple domains (Cole et al. 2013).

We begin our discussion with a review of characteristics of the task being learned and the knowledge needed to perform the task. This is especially important in the context of the knowledge and knowledge representations that it elicits in the learner.

Tasks and Task Knowledge

The space of ITL tasks centers around goal-oriented tasks, and thus one of the first aspects of the task to be learned is the goal itself—what is being achieved. For some tasks, the goal is straightforward to describe, such as the end state in checkers or chess; for other tasks, the goal might be harder to state formally or even informally, such as writing an essay or giving an eloquent but humorous

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

speech. The goal might also be dependent on the current environment—for instance, when setting a table, the appropriate setting might depend on the types of food and drinks being served.

The actions that can and cannot be performed in the context of a task offer additional constraints. On one hand, there may be actions that are tempting to perform but not available to the learner agent (e.g., going through a locked door); on the other hand, other entirely possible actions might not be realized at first by the learner (e.g., insight problems that require "thinking outside the box"). The former constraints embody limitations of the task and/or external world themselves, whereas the latter constraints may embody limitations (perhaps temporary ones) associated with the learner. There is a more general question of whether we need an explicit metric to assess how well the learner has learned the basic properties of a task and distinguish it from how well the learner performs the task successfully. Though this distinction seems useful and indicative for evaluating the learner in domains such as games, in general, the line between what are "legal" actions and what is simply good performance is not clear cut. Moreover, there is not always a need to learn all the rules to be able to be successful—some situations never occur, and sometimes even a learner who knows a few basic rules can perform well. These challenges are even more difficult in ill-constrained problems, such as writing and telling jokes.

Besides the general performance metric described above, learners must often form their own notion of this metric, and its formation becomes part of the learning process. This metric is important for deciding and learning the sequence of effector actions while pursuing a specific task, as it is useful in evaluating various actions available in a state. The metric may be explicitly defined using general criteria: for example, for some tasks, the time needed to execute actions may be more important than the number of actions taken; this performance metric will cause the agent to prefer a longer sequence of actions that takes less time over a shorter sequence of actions that is more time consuming. While a performance metric might be explicitly communicated, it is arguably more commonly implicitly communicated: for example, the Argentine tango has no explicit goal state, but instead is danced in a counter-clockwise movement around the outside of the dance floor.

In broad terms, then, learning can be thought of as an *iterative refinement* of a number of aspects of the learner's knowledge: the overall knowledge base being extended and refined, including not only the typical notions of skill and factual knowledge, but also iterative refinement of the learner's understanding of the goal, performance evaluation metric, and problem state.

Characteristics of Learning in ITL

There are a number of characteristics of learning in the context of ITL that, all together, distinguish this challenge from others in the broader space of problems in machine learning, cognitive science, and artificial intelligence.

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

One of the most significant ones is the necessity of learning from a small number of training instances. In contrast to approaches that rely on thousands if not millions of training examples, a human teacher typically provides only a handful of examples from which to learn. While these examples are typically of much higher quality than in a large-scale data set—carefully chosen to illustrate specific points—the paucity of data often necessitates a radically different approach to the learning problem. This idea is related to the work done on "one-shot learning," although in the ITL case, the potential for continual interaction that influences the learner (not to mention the teacher) provides for a richer set of data and experiences.

Another distinguishing feature of ITL relates to the interaction of the learner and teacher—in particular, the possibility that the learner can directly ask the teacher for feedback and guidance. While most learning algorithms operate in batch mode on large data sets, an ITL agent may direct its own learning by prompting the teacher for specific information—for instance, more instructions, clarifications on ambiguities, or additional specially chosen examples. In turn, the learner's ability to ask such questions implies the existence of metacognitive abilities on the part of the learner, namely an ability to inspect its own knowledge and to identify gaps or potential sources of errors. This ability is needed to provide the teacher with specific feedback. A simple "dump" of the agent processes would not be enough, or not be informative for a human teacher; instead, the agent must identify specific problems, or explain its own difficulties in a targeted way.

The learner's interaction with the teacher also creates conditions for a "push-pull" alternation between two major types of learning, generalization and specialization. For example, typically the examples provided by the teacher to the learner would be specific in nature (e.g., "grab this cup"), and it would be up to the learner to generalize these procedures to related tasks. In fact, children are typically conservative in their use of new artifacts (Casler and Kelemen 2005) and tend to imitate even causally irrelevant actions in learning new sequences (over-imitation; Lyons et al. 2007). But, if and when any overgeneralization takes place, specialization can be invoked to make instances of the more general rules, and/or augment the general rule with exceptions to the rule.

A teacher can also help the learner to construct hierarchical structures of understanding around their behaviors. Again, although the teacher may start with specific instances, they can also point out higher-level general properties of their examples (e.g., "notice that we always fill the pot with water before turning on the stove"). Recent cognitive modeling frameworks have some of these properties and would allow the learner agent to combine smaller behaviors into larger ones in an incremental manner. Learning about multiple concepts and the way they relate to each other, even in the simplest case, takes us from classical statistical learning to a structured learning paradigm, but this structure is usually expressed in a high-level abstraction and in a relational form (e.g., Kordjamshidi et al. 2015). The type of learning techniques that are used must

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

consider the relational structure of the domain and background knowledge, and to learn from examples in the context of these challenging issues. A natural consequence of these characteristics is that the learning examples may become structured as well—for example, expressed in terms of demonstrations, explanations, and interpretations of the world.

Last but certainly not least, one of the most important elements of ITL agents is a rich endowment of knowledge. Learning almost never happens "from scratch," but rather combines existing pieces of knowledge, large or small, in new ways to achieve a new task. An agent's background knowledge will necessarily be incomplete (if they knew everything already, they would be operating in a static world and no training would be needed). To date, existing large-scale open-source ontologies, such as CYC (Lenat 1995) and YAGO (Suchanek et al. 2007) can be used to provide a backbone for such domain knowledge. The advantage of starting with a large-scale ontology is that it reduces the time to build new systems, since reasonable knowledge representations for predicates and concepts have already been created. In ontologies such as OpenCyc, there is also a context structure, called *microtheories*, that is useful for focusing reasoning and also enables irrelevant subsets of the knowledge base to be removed if desired, to reduce an agent's memory footprint.

Capabilities to Support Interactive Task Learning

People are, of course, currently the best interactive task learners in existence. A deeper understanding of how people learn tasks interactively may provide valuable insights in how to create artificial systems that can learn as flexibly. At the functional level, we will catalog some capabilities that interactive task learners need.

There are a number of capabilities that can be identified for a learner agent in the context of ITL—capabilities that are not necessary per se, but certainly desirable as a part of the learner, with a degree that depends on the task domain. Some of the most important cognitive capabilities are:

- *Theory of mind*: In establishing common ground in communication (see Levinson, this volume), humans draw on "theory of mind" capabilities: knowledge or assumptions about the state of mind of others (e.g., their beliefs, desires, and intentions) and the ability to reason about these (e.g., Wimmer and Perner 1983; Leslie et al. 2004). Similarly, establishing a shared understanding and common knowledge in ITL—necessary to support communication of relevant information between the student and teacher—will require an agent to have some elements of theory of mind, such as detecting goals and intentions, and having beliefs about their human teachers and other agents.

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

- *Causality*: Understanding causality is fundamental to human rationality (Gopnik and Schulz 2007). This ability will not only improve an agent's ability to reason about the world as it learns and performs its tasks, but it is crucial to the process of determining its human users' intentions and goals as part of interacting and establishing common ground.
- *Self-monitoring*: Another key capability of ITL agents is the ability to monitor their own performance and learning. When being taught a new task, potential collisions with prior learned knowledge (perhaps trained by a different teacher) should be noticed by the agent and resolved in some way (e.g., by inferring the best solution, or by explicitly asking the teacher). Due to the fact that the agent is responsible for its own internals, it needs to monitor its own performance and use that to either ask for additional instruction or formulate and tackle its own learning goals.
- *Retrospection*: An ITL learner will be continually engaged in interactions with an instructor as well as with its environment. These interactions occur in real time and therefore, the agent may not have an opportunity or computational resources to recognize general patterns in the data it is observing or to judge how similar the current situation is to something in the past. In such cases, the capability of retrospection—to be able to recover, reason about, and learn from data at a later time may become crucial. Retrospection has been shown to be useful in category learning (Kuehne et al. 2000), where the agent maintains a set of exemplars from the past and previous generalizations. Each new observation is compared to these sets to discover new patterns in experiential data. Another example is explanation-based learning, which requires an analysis of a complete trace of task performance for generalization, but the complete trace may not be available during incremental interactions with a tutor (Mohan and Laird 2014). While using an explanation-based strategy, the ITL learner can store away its interactive experience and at a later point reason about why the instructed trace was useful in achievement of the goal. This retrospective analysis of its own behavior could be critical to produce generalizations and further learning.
- *Explainability* (or transparency): The ability of the agent to explain elements of its learning can promote generalizability and improve its potential to communicate effectively with the human user. Although this ability may be especially evident in cases where the learner agent then becomes a teacher of others (human or agent), this type of transparency can play a key role in self-monitoring and learning from retrospection, regardless of the end task.
- *Directed attention*: When learning any task, but especially in a busy real-world environment, the ability to focus on task-relevant details—and to ignore task-irrelevant details—is critical to the success of learning and performance. Given the ambiguous nature of most

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

communicative acts, directed attention is also very important to successful interaction by narrowing the search space over what is being referred to or what is intended.

- *Operationalization of knowledge*: One of the most challenging aspects of interactive task learning is the problem of converting the externally specified knowledge—whether from language or examples—into an internally executable form. It is somewhat reminiscent of a run-time compiler, but for task knowledge instead of a computer program. There are several possibilities, including creating a declarative representation of task knowledge that is interpreted by an agent, and compiling declarative task knowledge into more efficient rules or even code. One option is to create a declarative representation of the task knowledge that is interpreted by the agent; for example, in the Northwestern Tic-Tac-Toe program (Hinrichs and Forbus 2013b), task information extracted from natural language and sketching is translated into a simple representation and then interpreted by an interpreter. Another option is exemplified by Rosie (Kirk and Laird 2016), in which declarative task descriptions are automatically converted into a native rule representation, which execute 80 times faster than if they were interpreted. Both of these approaches seem to find support in the neuroscience of how the human brain rapidly learns new tasks (Stocco et al. 2012; Cole et al. 2013). Existing data suggest that tasks are represented in a declarative and unified format (Cole et al. 2013) and that these representations are interpreted "on the fly" or "just in time" by subcortical brain structures, like the basal ganglia (Frank et al. 2001; Stocco et al. 2012). With time, however, tasks undergo significant reconfigurations, with the original declarative knowledge being recoded into more efficient, procedural terms (Chein and Schneider 2005).
- *Handling large amounts of knowledge*: Given the need to learn multiple tasks over time, through multiple means (instructions or examples) and possibly in multiple modalities, an ITL agent should be able to handle and maintain a large amount of knowledge in a scalable and efficient format. In particular, it should be able to use representations that can integrate different task descriptions (from instructions or examples) and different actions (for different tasks), avoiding conflicts between representation and building up a knowledge base (world knowledge, teacher preferences) that can be used to facilitate learning future tasks.

Interaction in Task Learning

Although interaction has arisen as an issue throughout the earlier chapters, it is worth highlighting the critical role of interaction in task learning in its own

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

right. One way to provide context to this problem involves asking a series of questions about the role of interaction in ITL.

First, why is interaction so critical to ITL? Interactive learning can lead to agents that are much more adaptable to the contexts in which they are deployed. However, relying on a human user to train an artificial agent on a range of tasks imposes some degree of burden on the user, and interaction for its own sake is not recommended. Incorporating interaction into a task-learning scenario should be motivated by restrictions imposed by the task goals and environment, such as the variety of tasks that might be important to learn, not all of which (or all variations/parameters of which) are known *a priori*; the lack of sufficient training data; and the difficulty of hand-coding sufficient knowledge into the agent.

We define the goal of interaction in task learning as communicating any and all aspects of the task to be learned, and an important requirement for effective communication is *common ground*. Common ground refers to the knowledge in each agent of which aspects of the situation are shared between the two agents (Clark and Brennan 1991; Clark 1996). For example, a teacher and the robot can both know how to do a task (e.g., making coffee knowledge is in their situational awareness), but they might not know that the other knows how to do this task and that would mean that it is not in their common ground. Hence in their interaction, the teacher might think that she has to teach the robot how to make coffee. Interactive task learning can take place without common ground (e.g., the robot simply memorizes the commands of the teacher), but when common ground is also created between teacher and robot, this interaction becomes more natural because there is more shared information to use in the interaction.

How can interaction best facilitate establishing common ground between learner and teacher in the context of achieving effective learning? One of the main challenges that a participant in an interaction faces in establishing common ground is determining the intentions of the other participant(s) in the interaction. Natural communicative acts between humans (as opposed to artificial symbolic systems) are generally highly ambiguous and/or underspecified (Keysar et al. 1998). Humanlike interaction between a human user and artificial agent will face these same issues. For example, if the human tells her domestic robot "put the fish in the refrigerator," background knowledge will help resolve the ambiguity between "fish" as a pet and "fish" as a food. If she says "it's cold in here," only experience with that person (or further interaction) will enable the robot to decide whether the person prefers a blanket or having the thermostat adjusted. At the same time, humans have expectations about communication, as in the Gricean maxims of quantity, quality, relation, and manner (Grice 1975), and pedagogy (Csibra and Gergely 2009) which guide how they interpret utterances and these kinds of expectations could be incorporated into ITL robots.

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

Given the current state of the art, we have a long way to go before artificial agents can fully interact naturally with a person, the way people so effectively and efficiently do with each other. It is useful to specify various dimensions of interaction and to consider the level of sophistication of the mechanisms—and the resulting level of naturalness of the communicative capabilities—of the agent along each of those dimensions, as indicated in Table 15.1.

How much should interaction be used in ITL systems? System designers will need to determine an appropriate balance of the burden of effort between the human user and the computational agent—or perhaps the learner agent can determine its own need for interaction and behave accordingly. For the human teacher, issues to consider include: how much can they be expected to effectively elaborate task/background knowledge, especially for implicit knowledge; and how much can they be expected to (learn to) interact in a restricted way. The design of the agent will also need to examine to what degree user needs can be anticipated, how feasible it is to program in the needed knowledge, and what the limits are of the state-of-the-art in the mechanisms and algorithms of interactive learning.

To better understand interaction as part of ITL, we need to consider both the types of knowledge being communicated between agents, and separately, the types of communication and information that package this knowledge for transmission between agents. First, there are a number of types of knowledge that could be communicated as part of the ITL learning process, including:

- Perceptual and attending knowledge (e.g., object identification in the world, perceptual chunking)
- Goals: final state(s), space of possible goals

Table 15.1 Dimensions of interaction: general questions about communicative interaction abilities are listed in column one, followed by the corresponding dimensions along which agents can be assessed and the scale endpoints of those dimensions (columns two and three, respectively).

	Dimensions of an Agent's Communicative Abilities:	
Who can communicate, and who can take initiative in communicating?	Initiative	Passive ↔ Active
What can be communicated, and what can be understood?	Expressivity	Simple task knowledge ↔ Sophisticated task knowledge/strategy
	Theory of mind	Task/world knowledge ↔ Knowledge of other agent's goals/ intentions/beliefs
How is communication carried out?	Multimodality	Single mode ↔ Multiple coordinated modes
	Interactivity	Single turns ↔ Multiple coordinated/ nested turns with repair

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

- Possible actions, given the effectors and world constraints
- Action sequences, at multiple levels of abstraction, with action hierarchies, and possibly full and partial plans
- Action-world model that describes preconditions and effects of given actions
- Measure of performance for intermediate and final states
- Metacognitive strategies (e.g., learning strategies)

For each of these types of knowledge, there are multiple ways for the learner to acquire this information from the teacher as discussed by Thomaz et al. (this volume). Below is a list of different methods for communicating this information.

- Illustrative examples
- Demonstration
- Questions
- Feedback
- Explanations
- Sketches
- Gestures
- Facial expressions
- Language in the general case

The multiple types of knowledge and ways of communicating information are far from mutually exclusive, but instead often go hand-in-hand. For example, one way of communicating task knowledge is to communicate how the task can be learned—a metacognitive process—which can be done by communicating a learning strategy. A teacher might point out where task knowledge can be learned in the environment (e.g., a textbook to read), or might provide a self-training strategy that allows the learner to acquire a particular task. For instance, a search-and-rescue training instructor might advise on structuring the learning problem in terms of strategies for identifying signs of movement, potential hazards, and so on (each of which could be further decomposed). This type of metacognitive knowledge often arises in the context of other types of knowledge, such as knowledge about the goal(s) to perform and/or the performance measure to be used in self-evaluation.

In other situations, the learner might guide itself through the process as a form of learner-led interaction. For example, in the field of developmental robotics (Cangelosi et al. 2015), the goal is to develop robots that learn in humanlike ways and often these robots have exploratory behaviors in which the learner leads the interaction. While ITL is typically focused on a teacher guiding the agent to learn a given task, some learner-driven interaction could enhance ITL capabilities. For example, an ITL robot may have the task to sort a set of novel objects into two boxes. One approach is for the teacher to lead

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

the interaction, by telling the robot which object to select and then directing them to the target box (e.g., "pick up the green squarish thing and put it in the red box"). The alternative learner-led approach is that the robot picks up a random object and moves it toward a box (a child might put objects in boxes just for fun as part of their exploratory program). The teacher then only has to say "yes" or "no, the other box" to teach the task. In both approaches, the robot learns the task, but the second approach may be faster and more natural. Combining both learner and teacher-led approaches could lead to systems that learn more quickly with more natural interactions from teachers.

Evaluating ITL Systems

Any discussion of learning in the context of ITL—whether it relates to tasks, interaction, or learning itself—naturally leads to discussion of evaluation, namely how to evaluate whether an ITL system is performing well along some dimensions. Because an ITL agent is an interactive system in itself, many general evaluation measures from the fields of human–computer interaction, user experience, and human factors apply to ITL systems as well. However, ITL also raises a number of issues and challenges for evaluation that are either unique to, or especially important for, ITL in particular. This section reviews what we consider to be the most significant evaluation metrics and criteria with regard to the three dimensions of interaction, tasks, and learning.

Evaluating Interaction

As mentioned, there are many ways of evaluating interactive user systems in the general sense, from empirical user studies to heuristic techniques to conceptual models to computational cognitive models. For ITL in particular, some of the most significant factors for evaluating interaction include:

- *Usability*: An interaction system must be both easy to use and learnable. For example, if a human teacher must learn a new, complicated language to communicate with the agent, it is typically not considered usable. There are many common and acceptable ways to evaluate usability, including formal modeling methods (e.g., GOMS; Card et al. 1983), task analysis, and empirical methods.
- *Naturalness of interaction*: A criteria for evaluating ITL systems is the "naturalness" of interaction. Are there important aspects of natural interaction that are missing from today's systems? Is it really enough to add natural language to ITL systems, especially as a replacement for programming or scripting? Considerations of the degree of natural interaction would normally require other dimensions. For example, interactive turn-taking is an important problem that might be more

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

necessary for the system to interact naturally than vast language capabilities. Naturalness does not come simply from adding natural language but may be a function of the knowledge of the teacher versus the knowledge of the system to be trained. For example, the teaching of complex business processes could come naturally in a formalism known to experts (e.g., a business process language) if the teacher of the ITL system is an expert in that domain.

- *Expressivity*: Both the agent learner and the human teacher benefit from being expressive; if one highlights a point through nonverbal means, the interaction contains a richer set of data that can be used by the other agent for understanding and clarification. For example, a human saying "put the BLUE ball in the pail" suggests that the color is especially important in the current environment. It may be difficult to define, however, how exactly to evaluate expressivity in an ITL context.
- *Customizability*: Another dimension of an ITL agent's interaction is the degree to which it can adapt to the needs of a specific teacher. An ITL softbot that resides on a phone, for example, might be able to learn, over time, the specific needs and capabilities of the phone's user. For example, it might learn a user's preference for certain ways of organizing his or her calendar (e.g., no meeting before 9 a.m.) or add user-specific concepts to its ontology (e.g., the concept of "ultra-marathon running"), which would aid in the learning of future user-specific tasks. Extendibility can be empirically measured as the improvement in the quality of interaction (e.g., fewer examples, less questions, less feedback, or shorter instructions needed) for a repeated teacher, when compared to a novel teacher instructing the same task.
- *Flexibility*: A good ITL system should be flexible in how it interacts with a teacher; that is, it should allow multiple modalities (instructions, examples, and demonstrations) and be able to learn from all of these to the maximum extent. This is particularly important because different teachers (and certainly different human teachers) might have different preferences on how to describe a task, or might have different abilities. Consider, for example, the case of individuals with disabilities; a mute teacher might prefer to demonstrate a task, while a bedridden one might prefer to use linguistic descriptions. Flexibility can be empirically evaluated by preventing teachers, in an experimental setting, from using one or more modalities; for example, by preventing them from speaking, or by preventing them to use gestures, or to perform the actions. The extent to which an agent's performance suffers from such restrictions will be inversely correlated to it flexibility.
- *Expertise required for effective use*: One of the most promising aspects of ITL is that it could allow nonexperts to directly modify the knowledge and behavior of AI agents without having to learn programming.

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

For example, instead of having to use a programming language such as Java or C++, or even a scripting language, an untrained human could extend the capabilities of a home robot or an intelligent assistant on their phone. There are two metrics that would be relevant to this. The first is how long it would take a person with professional level skills in programming to develop an agent with the same behavior, possibly using different paradigms (e.g., using Java vs. a scripting language vs. an ITL agent). This could be compared to two measures: how long it takes a nontechnical person to learn how to interact with an ITL agent (e.g., the idiosyncratic nature of its language capabilities and ways of presenting examples) to some level of proficiency (e.g., 95% correctness); and how long it takes them to teach an ITL agent so it has the appropriate skills. Note that it would be useful to measure this over multiple, different types of tasks.

- *Diversity of modalities*: People communicate with each other via multiple modalities. Often this occurs simultaneously, as when someone talks while sketching or circles something on a photo to illustrate a point of reference in an email. One dimension for evaluating the naturalness of interaction for ITL systems is both the number of modalities and how sophisticated they are. For example, if we consider textual languages as one dimension, we can view it as being anchored by programming languages on one end, and unrestricted natural language dialogue on the other, with more restricted subsets of natural language understanding (e.g., controlled vocabularies, simplified syntax) in the middle. Similarly, visual communication as a modality can range from selection on a fixed map or diagram on one end to open-ended motion on the other end, with sketch understanding and image understanding as intermediate. Sounds and audio are other useful dimensions for communication, whether they be fixed signals (e.g., audible alarms or indications), speech, music, or other forms of audible communication. The complexity of dialogue supported is yet another dimension, ranging from unidirectional on one end, to fully mixed-initiative open dialogue on the other, with clarification questions and subdialogues as possibilities in between.
- *Natural interaction versus radical transparency*: Work so far on interactive task learning has focused on using natural modalities because that is the way that people normally train other people. This is challenging for ITL, of course, because it means that ITL agents must have communication skills that are much closer to those of people compared to the state of the art. We would be remiss if we did not mention another possible alternative: radical transparency. In a radically transparent system, the internals of the system are visible in some easily understandable way by the people who train them and work with them, with the

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

critical parts of their internal operations made visible by some powerful high-level language (perhaps visual in nature). Thus, we could program them graphically, in a manner reminiscent of RoboFlow (Alexandrova et al. 2015), and beyond that, monitor a high-level summary of their internal state and browse what they know about the world and about tasks. It seems unlikely to us that the complex operations of agent cognition can be neatly summarized in ways that would make training a teacher or user no more burdensome than training them to work with a dog or horse, but it is a possibility that might be worth considering for simpler forms of agents.

- *Context sensitivity*: Context sensitivity defines an agent's ability to adapt to different human users in their communication and performance of tasks. In the terminology of common ground, the agent is able to maintain knowledge of common ground specific to each human, including their shared referential background and context for effective communication, and any user preferences for performing the tasks. This could be assessed by demonstrating that an agent adapts appropriately, both in its communication and task performance, to each of several human users.
- *Tolerance to errors*: When working with an interactive task learning system, it is important to be robust to errors that either the human trainer or the agent may make. For example, if the human trainer provides an incomplete (or even incorrect) description, the agent must be able to learn the correct knowledge so that it will not keep a persistent "bug." A reasonable method of evaluating tolerance and robustness to errors is to demonstrate how a system might recover from errors and perhaps even learn to avoid them in future instances.
- *Safety*: The purpose of this evaluation would be to determine if ITL is able to provide some guarantees of safe operation in the context of a malicious, or possibly mischievous instructor, where the agent is taught undesirable behaviors. The challenge is to develop capabilities in the agent that enable it to recognize behaviors that violate some set of norms, and to distinguish actual facts from "alternative facts." Measuring and evaluating an agent along this dimension will be challenging, but could possibly involve corpora of example instructions, examples, and demonstrations that include both valid and invalid (unsavory?) knowledge for the agent to learn.

Evaluating Tasks

A similar set of evaluation metrics can be derived for evaluating tasks in the context of ITL—that is, evaluating an ITL agent with respect to the tasks that can be handled by the agent. These include:

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

1. *Diversity and coverage of tasks* (agent generality): One of the promises of ITL is that an agent will be able to learn more than just one or two tasks, and more than just tasks from a limited repertoire (although ITL agents that are targeted for domain clusters with limited diversity could be very valuable in certain contexts). Diversity can vary along many dimensions and often depends on the overall domain(s) in which the agent will be used.
   - Different task formulations (problem space, procedure, and optimization) that determine what needs to be learned to perform a task, as described by Laird et al. (this volume).
   - Different terminology needed to specify the task: Are many of the terms shared between tasks, or is there a small set shared by the tasks?
   - Different levels of abstraction of the task and preexisting knowledge can be required for the task: Does the task primarily involve perceptual-motor interactions (e.g., picking up different types of bottles) or more abstract actions that are decomposed into such primitives (e.g., cleaning a room or delivering a box)?
   - Different types of environments, such as whether the task takes place internal to the agent, involve interaction with other software or a physical embodiment (e.g., a robot).
   - Different level of complexity of tasks can arise in many forms across tasks. Some possible dimensions include the complexity of the task specification (how hard it is to understand the task) and the complexity of the task problem space (how hard it is to perform a task once it is learned).
2. *Explainability*: Another important aspect of tasks is their explainability, or ease of describing from one person to another. For example, learning the sensorimotor skills of soccer is more complex than learning the legal moves of chess, because verbalizing the physical movements in chess is straightforward but verbalizing the physical actions of soccer is not. Other tasks such as "setting tables" might also suffer in terms of explainability from implicit versus explicit knowledge. In general, explainability varies as a function of the task and the teacher, as what is introspectively accessible to one teacher might not be to another. However, at least when the teacher is predefined (e.g., it is predetermined to be a human), explainability can be measured for a specific task.
3. *Noise and randomness*: While some tasks have little to no "noise" or randomness associated with the world environment, others are highly variable and subject to such noise. For example, chess is completely determined by the players' moves (except for the possibility of an

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

extreme event, like someone upending the board); a dice-based board game is somewhat determined by the players but somewhat by the randomness of the dice; and driving down the highway is a relatively straightforward task in terms of the goal, but could potentially be a very difficult one because of the variability of the environment (e.g., darkness, rain, snow, potholes).

4. *Risk*: As related to a learning agent, risk refers to how dangerous the agent may be to itself or to others as it is learning to perform, or ultimately performing, a new task. For example, if the agent is a vehicle learning to drive itself, there is a potential risk that it may cause an accident while learning.
5. *Observability*: One final aspect is whether the agent directly observes all of the relevant aspects of its environment (e.g., tasks performed by a tabletop robot), or whether there are aspects of the environment that the agent cannot sense without explicit actions (e.g., tasks performed by a mobile robot that must move between rooms and interact with objects that are initially unobservable).

Evaluating Learning

Evaluating the learning of an ITL agent is an interesting challenge in comparison with previous work on machine learning and evaluation measures. For many machine learning systems like classifiers, there is a well-known trade-off between the amount of training effort, the accuracy of the final system, and the complexity of things that are trainable. However, a particular classifier is typically targeted for a specific domain, often with a large set of training examples. In contrast, ITL would typically involve a small number of higher-quality, targeted training examples, and potentially could improve learning by receiving help (explicitly or implicitly) from a teacher guiding the learner through the process. Thus, we can identify several dimensions along which evaluating learning is especially important:

- *Speed*, *efficienc*y: One of the most significant evaluation criteria for ITL involves how quickly the agent can learn from its teacher. Speed can be measured in real time, although for some mechanisms or algorithms, other measures may be more appropriate—such as the number of experiences or the number of iteration cycles required to achieve adequate performance. In most cases, we prefer ITL systems that learn as quickly as possible, since this minimizes the user time needed with the system. There may be some cases, however, in which it is desirable for an agent to learn at a rate roughly comparable to a human—for example, when using an ITL learner as a stand-in for a human learner when evaluating the learnability of concepts (e.g., a simulated student

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

learning new mathematical concepts to be integrated into an intelligent tutoring system).

- *Performance*: As a complement to speed and efficiency, another criterion is how well the agent performs the task that it has been taught. There are many possible ways to quantify and measure performance, including resources required to perform the task (e.g., time or energy), quality of the solution (how well it achieves all aspects of the goal), percentage goal completion, and time to completion.
- *Transfer*: An ITL learner is expected to learn multiple tasks. In the case where some of the given tasks are similar, ideally the agent will exhibit transfer of knowledge from one task to another. An ITL learner should exploit similarity of structure, action control hierarchies, and so on to generalize previously learned knowledge to new tasks. For example, after learning how to wipe down a table, learning to wipe down the windows should be easier (from certain perspectives) as the action control is similar. This can be measured as number of interactions (e.g., demonstrations, instructions) taken to completely learn a task, number of trials (independent exploration or time) taken to learn a task, and so on. We would expect that for similar tasks (as measured, e.g., by overlap of features, action hierarchy), a better ITL learner requires fewer interactions or trials to learn a task at a specific level of performance.
- *Interference tolerance*: As an ITL learner is learning to perform multiple tasks, knowledge acquisition has the potential to suffer from interference—that is, if learning a new task reduces performance in an already-learned task. This might happen if the algorithms are overzealous in generalization, or if the task representations do not sufficiently discriminate between different tasks. An ITL learner should be able to recover from such interference effects—for example, by asking for more instruction that allows it to produce separable representations. This capability can be measured by the number of generalization errors made (less for a better ITL learner), complexity of instruction needed for recovery, or amount of training needed for correction. The transfer-interference metrics will likely be in terms of a trade-off, and a good ITL learner will maximize transfer and minimize interference.
- *Robustness with respect to teachers*: The quality of an agent's learning capabilities should be assessed with respect to its robustness in the face of varying quality in the information it is given by the teacher. While it is comparatively easy to design a system that learns from carefully crafted instructions and perfectly selected examples, real-life human teachers vary in the quality of their directions. For instance, less capable teachers might give incomplete task descriptions, use ambiguous language, or provide misleading examples. This form of learning robustness can be empirically measured by experimentally degrading

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

the quality of the directions given to the learner; for example, verbal instructions might be shortened by deleting an increasing number of words, and the number and coverage of examples given to the agent might be reduced.

- *Incremental learning*: One overarching desired property is the allowance for incremental learning, to provide a trade-off space between training time and task performance. Ideally, if the user only wishes to spend a very short time training an ITL agent, that agent could perform at least reasonably well; but if the user wishes to spend more time (then or later) in doing further training, the agent would likewise improve in its performance.

Evaluation Discussion

With all of the evaluation criteria above, it should be noted that the criteria are most useful when evaluated within the context of an appropriately diverse set of tasks. In particular, a broad evaluation across criteria would ensure that task-specific approaches have not been incorporated into the agent to simply improve its behavior relative to a metric. The actual set of criteria used depends on the overall objective of the ITL agent; whether the agent is being used as a personal assistant, a good game player, or an assistive robot would presumably skew the priorities of the evaluation criteria in favor of those most important for the respective domain.

Because ITL is an emerging area, research may need to blossom further before these evaluation criteria are used across the research community in more prescribed ways, as in the development of benchmarks. At this stage, exploration of different task types, means of interaction, and learning mechanisms should be relatively unconstrained to allow for maximum creativity, and allow the research to push the boundaries of the science in all areas of this multidisciplinary effort. Nevertheless, the evaluation metrics are still critical even in these early stages for researchers and system designers to identify the most important aspects of ITL for evaluation—those that are particularly salient to the ITL context and that distinguish it from other machine learning and artificial agent approaches.

That being said, there could be benefits from developing a standardized set of tasks for evaluating ITL agents in simple yet realistic domains. One advantage of this approach is that it would provide a constrained domain in which experimental psychologists can test theories and collect data on ITL in humans. Currently, ITL behavior in humans is rarely studied because of the experimental challenges it poses as a very high-level domain. Providing simple tasks (e.g., simple card games) and a well-defined, constrained set of interaction models (e.g., instructions in restricted English, set of positive and negative examples of legal moves) would set the stage to examine how ITL

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

occurs in humans while carefully controlling and measuring the experimental variables. Furthermore, ITL tasks formulated within a formal framework provide objective metrics to control the experimental stimuli and design, which is what is needed to conduct statistical analysis on experimental variables—and to make valid statistical inferences. This would open up the door to more research of ITL behavior in humans, which could in turn provide important clues as to the fundamental mechanisms that should or could be implemented in artificial ITL agents.

There is at least one other important reason why a standardized set of tasks would facilitate experimental research on humans: the fact that the same task could be used in humans and artificial agents, and especially that artificial agents could be used as a model to analyze human neuroimaging data. Consider, for instance, a realistic example: teaching a human lying in an fMRI scanner how to play a simple card game that has been taught to the Soar-based agent Rosie. Rosie's internal states provide a formal model of learning, whose internal states that can be traced and used to generate a regressor for human brain activity. The estimation of various parameters for the regressor could be done using a Bayesian framework to maximize its correspondence to behavior and imaging data. When compared to brain activity, the Soar-based model could be used to locate regions of the brain whose activity more likely correspond to the model traces. In turn, this could be used to interpret the specific computations of different brain circuits. Finally, brain data can be compared to the internal processes of different agents, and insights from human data can be used to revise and improve on existing agents.

General Discussion

Interactive task learning establishes a problem focus that is AI-complete—drawing on knowledge representation and reasoning, computational linguistics, machine learning, robotics, computational perception, and so on. As a research area, ITL emphasizes integrative approaches that bring together all or many of these sub-areas into sophisticated systems with multiple interacting functions that must work in harmony with each other to solve overarching problems. This kind of research agenda is valuable in ensuring that research in the various subfields is grounded and contextualized in real-world problems whose solutions must coordinate the multiple intelligences required. At the same time, research that reduces aspects of any one of these areas into isolated components to be more thoroughly explored continues to be necessary. On one hand, the benefit of the overarching integrative framework is to inform the narrower focus and to ensure that component solutions will have relevance when brought into contact with other elements of the broader framework (e.g., Anderson 2007). On the other hand, the deep-dive research can find potential solutions to problems that may be too difficult to explore within the full

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions," edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26, J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

complexity of the integrated system. Ideally, each approach serves to provide motivation and guidance to the other, and work to their mutual benefit.

In this chapter we have largely focused on the relationship of ITL to other areas, especially noting the subsets of these areas that we believe are most relevant and important to the ITL effort. At the same time, our hope is that ITL also helps to push the boundaries of innovation in these areas themselves, since any improvement in any of the components of ITL will ultimately benefit ITL as well. One of the clearest areas of benefit comes in the field of machine learning. ITL holds the potential to have a major impact on machine learning, which is defined by the question: How can we create machines that automatically improve their competence through experience? To date, machine learning has focused on only a small fraction of the types of learning that humans exhibit—that is, mostly on statistical learning from large amounts of passively acquired data. For humans, however, interactive learning from conversations, demonstrations, and experimentation is central. This may well become a major thrust of machine learning over the coming decade.

ITL could also make a significant impact in the domain of cognitive neuroscience. Cognitive neuroscience has the seemingly impossible goal of understanding the human brain at a functional level—that is, how its activity gives rise to human cognition. The dominant approach in this field has consisted of dividing human cognition in putatively isolated functions, such as object recognition or decision making, and examining them with small-scale, controlled paradigms. However, many authors have suggested that this approach is limited, pointing out that perfect decomposition into basic functions might be impossible, or that any reasonably complex task ultimately requires the coordination of interconnected systems and functions. ITL research can benefit this area in two ways. First, it brings *formal models* of task learning, which could be specified mathematically, implemented computationally, and used to analyze neural data, as in the previously described example of using Rosie to interpret neuroimaging data from a well-defined card-game task. The approach of using models to understand and explain brain activity has been perhaps one of the greatest advances of modern neuroscience, with the application of reinforcement learning to the interpretation of the functional role of basal ganglia in procedural learning being an excellent example (Schultz et al. 1997). Second, it brings a general and integrative approach to learning in which multiple components exist at the same time. This is particularly important because it allows for study of very complex cognitive phenomena without losing track of the different role played by different components.

Yet another potential benefit comes in the area of linguistics and psycholinguistics. The rich models of human interaction from these areas will need to be formalized within a computational framework in which they interface to varied aspects of agent systems involving representation, learning, and processing of task, background, and interaction knowledge. Research into dialogue systems will need to consider embodiment in robots and the contexts of varied software

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

agents. Fundamental issues in interaction have great insights to contribute to ITL, and in turn ITL will identify interesting problems of interaction in a real-world context that highlight the connections to other cognitive capacities such as visual interpretation, various forms of reasoning, and physical interactions with the world and interlocutors.

From "Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,"
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.

From “Interactive Task Learning: Humans, Robots, and Agents Acquiring New Tasks through Natural Interactions,”
edited by K. A. Gluck and J. E. Laird. Strüngmann Forum Reports, vol. 26,
J. R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-03882-9.