2

Framing the Problem of Interactive Task Learning

Tom M. Mitchell, Simon Garrod, John E. Laird,
Stephen C. Levinson, and Kenneth R. Koedinger

Abstract

To support a precise discussion of interactive task learning, the problem setting in which teachers and learners interact in a shared world must be clearly defined and understood. This chapter provides a formalism to enable discussion of the different types of interactive learning: from teaching a robot to grasp a novel object, to instructing a mobile phone how to reach a friend in an emergency. It provides a way to speak precisely about notions such as shared knowledge between teachers and learners, presents working definitions of the internal structures of the agent, and describes the relationships between the task environment and the communication channel. It focuses on the *problem* of interactive task learning, not its solution, as a backdrop to further discourse in this volume.

Formulating the Problem

In interactive task learning (ITL), we assume that a set of agents jointly occupies a shared world. Each agent, A, has a set of sensors, $S = \{s_1, s_2, \ldots\}$, it uses to sense the world, and a set of effectors, $EF = \{ef_1, ef_2, \ldots\}$, it can use to alter the world through actions. For example, a robot might have sensors such as a camera and microphone and effectors such as wheels and arms. Similarly, a mobile phone might have both physical sensors and effectors, such as a microphone to sense the world and a speaker to affect it, as well as cyber-sensors, such as a web browser to observe the internet, and cyber-effectors, such as the ability to send text messages.

In general, we assume that these agents exist in a shared world, the state of which at time t is denoted by $w(t)$. The sensors and effectors of each agent operate on this world, producing a time series of its observations and actions. In general, the world may contain any number of agents with differing capabilities.

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.

Communication

To capture communication among agents (e.g., speech, nonverbal gestures), we assume there is a communication channel between each pair of agents. Given that all communications must go through the physical world in some form, we capture the communication channel in our model by designating a subset of effectors, EF_C (e.g., speaking, pointing), and sensors, S_C (e.g., listening and interpreting), that implement this communication channel. We think of these communication sensors and effectors as operating on a set of objects (e.g., words, symbols, fingers that point, underlines that highlight) that *refer* (among other things) to other objects in the world. However, communication differs from other forms of sensing and acting in at least two key ways:

1. Unlike the use of sensors to observe properties of nonhuman physical objects, the interpretation of sensor inputs in communication (e.g., input speech) is based on *shared conventions* between agents about the meanings conveyed by different communication actions and communication objects.
2. The communication channel can convey meanings that go beyond what can be perceived by other sensors of the physical world, including information about mental states such as "I believe Mary thinks the train is late."

Determining whether another agent is using its effectors in an attempt to communicate rather to perform some instrumental action in the world is part of any interactive task for a robot. This is discussed below under the section, "Task Environment and the Communications Channel."

Learning

Now let us consider some learning agent A. To define a learning problem precisely, we say that agent A learns to improve its performance, P, at task, T, through experience, E. In fact, we assume every learning problem can be defined in terms of some triple <P, T, E>. To illustrate, an agent might face a learning problem in which the task is to play the game Go (T), the training experience (E) consists of playing 1000 practice games against itself, and the goal of learning is to improve performance (P) measured by how frequently it can defeat a second agent in a 100-game tournament. A related learning problem might have the same task, T, and performance metric, P, but differ in the type of training experience available (e.g., learning from advice received by a teacher while playing games). As a further example, a robot might face the task of learning to set the dining table (T), from watching videos of a person performing that task (E), where performance (P) is measured by the precision of the final table setting minus the number of dishes dropped along the way.

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.

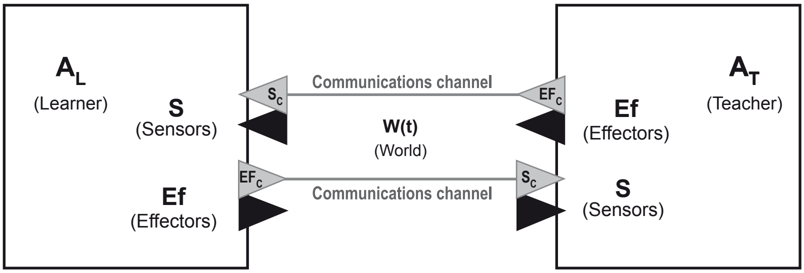

Figure 2.1 Simple world with learning agent A_L and teaching agent A_T. Black sensors and effectors sense and alter the agents' shared physical world. Gray sensors and effectors indicate the communication channel between agents.

It is important to keep in mind that the performance task (e.g., playing Go) is distinct from the learning problem itself (e.g., *learning* to play Go).

This definition of learning covers cases where the learning agent acts alone in the world to generate its own training experience (e.g., experimenting in the world via its sensors and effectors), as well as cases where the training experience involves other agents acting as teachers or collaborators to help shape the training experience. Figure 2.1 illustrates a prototypical situation of interactive learning between a learning agent, A_L, and a teaching agent, A_T.

Working Definition of ITL

We define ITL to be any process by which an agent (A) improves its performance (P) on some task (T) through experience (E), when E consists of a series of sensing, effecting, and communicating interactions between A, its world, and crucially other agents in the world.

There are many nuances to the ITL concept, as presented in subsequent chapters in this volume. These discussions introduce additional structure into the above problem formulation, including assumptions about how a teacher might impart different types of knowledge to a learner, the roles of demonstration and direct teaching by instruction, the internal architecture of the student and teacher, as well as the need for building common ground between the teacher and learner.

Internal Structures of an Agent

Having described the problem faced by a learning agent, we now turn to a discussion of the internal structures of the agent. The determination of what constitutes a proper internal structure for an agent to succeed at ITL is the focus of ongoing research, and although consensus has yet to emerge, many proposed approaches share certain assumptions; namely, that an agent possesses internal

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.

structures such as knowledge and goals. Here we introduce a simple vocabulary to support discussion of the knowledge and goals held by agents, as well as the mental states of others.

Knowledge

We assume that each Agent A possesses a set of beliefs which constitute its knowledge, K(A), about its world and agents. Knowledge can be about absolutely anything. It can include, for example, knowledge about the world (e.g., "there is a chair in front of Agent A"), other agents (e.g., "Agent A_2 is also in the room"), capabilities of other agents (e.g., "Agent A_2 is a good Go player"), as well as knowledge about the knowledge of other agents (e.g., "Agent A_2 knows that Agent A_1 knows there is a chair in the room").

Some of an agent's knowledge may be correct whereas other knowledge may be incorrect, and the agent might or might not know which is correct. Much of learning involves acquiring new knowledge and correcting or refining current knowledge.

Agents may possess (or act as though they possess) knowledge of which they are unaware, though their ability to communicate this knowledge depends on mental access to this knowledge in some declarative form. For example, a human agent may know how to recognize their mother yet be unable to communicate this knowledge to another agent in declarative form.

Goals

Agent A may have goals G(A) at any given time. These may include, for instance, goals applicable to the external world (e.g., set the dining table), to the agent's internal mental world (e.g., learn how to set dining tables), as well as to the mental world of another agent (e.g., help Agent A_L learn to set tables). Goals can have subgoals which contribute to an agent achieving the overall goal.

Common Ground

Shared knowledge and goals of agents are essential to communication, in general, and to ITL in particular. Communication and teaching can be viewed as processes used to establish and refine a set of shared goals and knowledge about the task at hand—the common ground between agents. When discussing common ground, it is helpful to define two related but distinct notions:

- The set of shared knowledge between agents A_1 and A_2 is defined as the intersection of knowledge between these agents: $IK(A_1, A_2)$.
- Shared goals is the set of goals held by both agents A_1 and A_2: $IG(A_1, A_2)$.

Again, it is important to note that agents A_1 and A_2 may themselves be unable to know perfectly which of their knowledge and goals are shared by the other

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.

agent. In fact, the ITL process is often driven by the attempts of these agents to better know what they already share, and the attempt to transfer knowledge between teacher A_T and learner A_L, thereby growing the shared knowledge, $IK(A_T, A_L)$.

The intersection of knowledge is, however, not sufficient to *ground* communication. For example, two people may both have the knowledge that one of two cars in a lot is electric, but they may not realize that the other person knows this as well; thus, when differentiating between the two cars in conversation, they will not refer to the car as "the electric car" but rather draw upon other criteria (e.g., color, model) to make the distinction. This demonstrates the need for a more specific notion than simply the intersection of knowledge (see Smith 1982). For successful communication, we need a set of assertions, S, such that A_1 knows S, A_2 knows S, A_1 knows that A_2 knows S, and A_2 knows that A_1 knows S. In this case, we will say that A_1 and A_2 *mutually know* S and that S is in the set of mutual knowledge $MK(A_1, A_2)$. This is the notion of *common ground* needed for establishing communication.

There is no foolproof way for agents to infer what is contained in common ground. So when human agents interact, they might try to discover a shared basis for common ground before treating anything as being in it. For instance, if two agents see a cup in front of them and both observe that the other also sees the cup, this may act as a shared basis for that cup to be considered in common ground. However, if it is not apparent that the cup is salient to both agents, one agent might attempt to *ground* the information by holding, for example, the cup up or pointing to it (Clark 1996:98). Interactions between human and nonhuman agents (e.g., robots) present challenging problems of finding the appropriate shared basis for inferring that something is in common ground.

In summary, common ground is always specific to two or more agents at a point in an interaction. It can be thought of as nested: all the mutual beliefs any human could be presumed to have, all the mutual beliefs members of a speech community or nation may be presumed to have, all the beliefs presumably shared by our special interest groups, or friendship network, partners, and so forth. Above all, common ground contains the history of previous interactions of not only what we have established as mutual beliefs, but also the particular referential tokens we agreed to use to refer to them (e.g., the blue car is the electric car) through a *referential pact* (Clark 1996).

Task Environment and the Communication Channel

Tasks are performed in the world, but that same world also provides the channel for communication between humans. Spoken language is the most obvious communication channel between humans but there are others as well (e.g.,

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.

gestures, documents). It is therefore worthwhile to distinguish the parts of the world used in task performance from those used in referential communication about task performance.

We define *task environment* as the set of objects in the world used or created in performing or attempting to perform tasks. When a learner agent has learned a task, it can manipulate these objects to achieve the task goals. A teacher agent may also manipulate these objects, for example, to demonstrate task steps.

Recall that above we defined a communication channel between agents in terms of sensors and effectors that operate on objects used to refer to the task environment, such as words or fingers that point. Depending on the nature of the task, objects can be in either the task environment or the communication channel, or both. For example, robot fingers may be used to point (communication) and to manipulate physical parts during a repair task (task environment). Similarly, words are typically part of the communication channel, but they can be situated in the task environment if the task domain is learning a second language. If an object is used or created while performing a task, then that object is part of the task environment. If an object is used to refer to an object in the task environment, then we say it is a "referring object" and is part of the communication channel. Because there is no clear *a priori* division between the sensors and effectors used for communication versus those used for accomplishing tasks, (e.g., a person may indicate agreement/disagreement with a word, a nod, eye gaze, a gesture, movement of an object, or even the length of pause before a response), one of the great challenges in social robotics is to read such personalized signals on the fly, as humans do. An important challenge for ITL is to understand how communicative signals are bootstrapped through other actions.

Some instructional strategies are executed primarily in the task environment, such as demonstration or corrective action (e.g., the teacher moves the fork to the left of the plate after the robot has placed it on the right). Other instructional strategies are executed primarily in the communication channel (e.g., the teacher says "put the fork to the left of the plate"). Some instructional strategies use both, whereby the teacher performs actions in the task environment while using the communication channel (see Chai et al., this volume).

In everyday terms, teaching which moves in the task environment includes "showing" what should be done or "correcting" what the learner has done incorrectly. Similarly, teaching that moves in the communication channel includes "telling" what should be done or providing "feedback" on what the learner has done incorrectly. When a teacher is showing and telling, performing corrective actions and giving feedback, that teacher is simultaneously using both the task environment and the communications channel. Whether "showing" or "telling," interaction is critical because of inherent ambiguities in communication through both the task environment and communications channel.

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.

Conclusion

This chapter has laid out a framing of the problem for ITL and proposed a vocabulary that will be used in the remaining chapters when discussing agents, their knowledge and goals, as well as their interactions.

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.