4

What People Learn from Instruction

Christian Lebiere

Abstract

Computational models offer a precise, quantitative way to represent the cognitive processes and representations involved when an agent interacts with another agent: from the receiving of instructions, to their interpretation, to the processes involved in learning to perform a task. This chapter discusses various forms of knowledge and skills involved in interactive task learning (ITL). It describes the components and processes in cognitive architectures relevant to ITL, organized around dichotomies of declarative knowledge and procedural skills, symbolic representations and subsymbolic statistics, as well as cognitive, perceptual, and motor processes. One specific cognitive architecture, ACT-R, serves to focus discussion. Using a model of interactive learning in decision making, it demonstrates how these components and processes interact. Representation, learning, and processing issues are discussed both in isolation as well as in the context of this integrated task learning model.

Introduction

A substantial part of human knowledge and skills comes from instruction, formal or otherwise. Children in modern societies spend at least a dozen years in school, sometimes significantly more well into adulthood. After school, people enter the workforce where a significant part of their knowledge and skill acquisition comes in the form of training sessions, followed by practical experience in the field. Outside of school or work, our environment in recent years has become increasingly characterized by the ubiquitous availability of information at the touch of a button (or touch screen). This information stems from an exponentially increasing array of facts about topics common and obscure, whether organized in formal repositories such as Wikipedia or marshaled in by search engines from the most remote corners of the Internet. It includes tutorials on how to perform any number of practical tasks, from cooking recipes to replacing a cracked screen, as well as entire classes on a range of topics from

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.

websites such as Khan Academy to university online courses. The idea that humans derive new knowledge and skills predominantly through painstaking reasoning and experimentation is becoming outdated, if it ever was accurate in the first place. Information is there for the taking, but the central question remains: How do people assimilate and integrate new facts with existing knowledge to understand, instantiate, and apply task instructions in both familiar and challenging settings?

Human-centered fields such as education and psychology contain a wealth of information on these questions, yet much of it is qualitative in nature, limited in scope, and thus cannot be easily integrated into actionable form. For instance, psychology experiments in the laboratory or naturalistic interventions in educational settings can measure the effectiveness of various forms of instructions but cannot, in general, explain how and why they work. What we need to guide interactive task learning (ITL) is a formalized, computational representation of the human learning process. It must be emphasized that ITL is fundamentally different from the typical forms of learning exemplified by machine learning techniques. Those approaches are typically based on algorithms that crunch through very large, uniform sets of context-free training instances. For humans, the task of learning from instructions and interaction with the world, by contrast, requires making sense of and organizing a quite limited and sparse set of information, applying it, and generalizing it on the fly. Requirements for a computational representation of human ITL must thus include a number of cognitive capabilities, including memory, action selection, decision making as well as perceptual and attentional processes and their interactions. The best match for these requirements is the concept of cognitive architectures (Newell 1973b), which originated in response to ad hoc domain-specific task models. A cognitive architecture is a computational instantiation of unified theories of cognition (Newell 1990) that integrates the various cognitive, perceptual, and motor mechanisms into one general, domain-independent framework. An increasing number of cognitive architectures have been proposed over the years. Here, discussion is grounded in one specific architecture, ACT-R (Anderson and Lebiere 1998; Anderson et al. 2004), although increasing evidence suggests that the cognitive architecture program is converging on an implementation-independent consensus (Laird et al. 2017b).

There is reasonable concern that cognitive architectures provide a too detailed, mechanistic account of human learning from instruction to be relevant in this setting. Newell (1990) formalized the idea of levels of description in his concept of bands of cognition, distinguishing between subsecond mechanisms in the neural and cognitive bands and phenomena in the rational and social bands, which can take place over days, weeks, months, or years. Anderson (2002), however, argued that while not every aspect of the fine-grained mechanisms will be relevant at the larger timescales, key characteristics will still be relevant and even determinative. A demonstration of such detailed relevance can be seen in the application of cognitive models

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.

to intelligent tutoring systems (e.g., Anderson and Gluck 2001). In that application context, cognitive models are formal representations of the student learning process; they are comparable to actual student behavior and can be used to infer their current state of learning and guide interventions (e.g., selection of problem set, hints, review of material) to improve the learning process. In particular, cognitive models have been shown to provide the right abstractions to enable proper understanding of performance curves as a function of the underlying knowledge and skills at the grain scale provided by cognitive architecture mechanisms and representations (e.g., Corbett and Anderson 1995). Thus, it is not much of a leap to suggest that cognitive architectures also provide the best tool to guide the ITL process, both in terms of design as well as run time.

Below, I describe components and processes in cognitive architectures that are relevant to ITL, organized around dichotomies of declarative knowledge and procedural skills, symbolic representations and subsymbolic statistics, as well as cognitive, perceptual, and motor processes. The interaction between these components and processes is exemplified in a model of interactive learning in decision making developed in my lab (Lebiere et al. 2013b). I conclude with a general discussion regarding potential avenues of research in developing cognitively based approaches capable of supporting ITL.

Learned Knowledge and Skills

Our discussion begins with a computational description of the knowledge and skills that are learned in the context of ITL (for further discussion, see also Wray III et al., this volume). Following the distinction made in many cognitive architectures, declarative knowledge is separated from procedural skill and symbolic structures from statistical parameters. For each category of knowledge, I describe the characteristics and instances of how it can be learned and supported through ITL:

	Symbolic	Subsymbolic
Declarative	Chunks: instructions, general knowledge, situation knowledge, beliefs	Activations: environmental statistics, quantitative uncertainty, distributional semantics
Procedural	Productions: heuristic strategies, general procedures, compiled skills	Utilities: learned rewards, probabilistic selection, adaptive generalization

The terminology used follows that of the ACT-R cognitive architecture, although similar constructs exist in most other popular cognitive architectures (Laird et al. 2017b). ITL systems have been developed using a number of different cognitive architectures, including ACT-R (e.g., Taatgen et al. 2006) and Soar (e.g., Kirk and Laird 2014; Mohan and Laird 2014; Mininger and Laird 2016).

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.

Symbolic Declarative Knowledge

Declarative knowledge represents explicit, conscious knowledge as symbolic structures, equivalent to propositional representations. These structures, called "chunks," bind together a set of values (which can themselves be chunks) into specific roles. Chunks are created and stored automatically in long-term memory by recording the state of working memory at specific points in time. Chunks can represent any type of explicit knowledge related to instructional task learning: knowledge of specific situations, general factual knowledge, knowledge of procedures and instruction steps, as well as knowledge of other people's beliefs.

Cognitive architectures make strong mechanistic commitments as to how knowledge is learned, represented, and accessed but typically do not make ontological commitments as to what the knowledge types are and how they are organized. This largely reflects the historical split between mechanistic and ontological approaches to cognitive science and artificial intelligence. A number of attempts have been made to import knowledge ontologies into cognitive architectures and, more generally, to adopt stronger and more systematic knowledge commitments (e.g., Ball et al. 2004; Wray et al. 2004; Best et al. 2010; Oltramari and Lebiere 2012; Salvucci 2014). Recently, Lieto et al. (2018) examined issues systematically with knowledge representation in cognitive architectures and compared the approaches as well as potential solutions to those problems in various architectures.

Subsymbolic Declarative Knowledge

Symbolic representations provide structure but not all (or even most) of the knowledge. Much real-world knowledge is not conscious but rather implicit, resulting from automatic learning of statistical patterns. Anderson (1990) argued that human cognition adapts to the statistics of the environment to optimize its performance given limitations (e.g., working memory capacity, attentional limitations, and various other architectural bottlenecks). This adaptivity is fundamentally heuristic in nature and can lead to errors and cognitive biases (e.g., Lebiere et al. 2013b) when its assumptions do not match the actual nature of the environment. Although "soft" aspects of human cognition (e.g., adaptivity, generalization, and stochasticity) can be described at higher (e.g., Bayesian) or lower (i.e., neural) levels of abstraction, in cognitive architectures such as ACT-R and Soar they are represented in terms of subsymbolic mechanisms that control access to symbolic knowledge structures. Specifically, they are formalized in terms of activation processes that determine which chunk is selected from long-term declarative memory for a given retrieval request. Activation is represented by a base-level term that captures pervasive regularities, such as the power law of practice and the power law of forgetting. Strengths of associations from context elements to memory structures enable

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.

task-sensitive knowledge access. Generalization from similar situations is achieved through a partial matching mechanism which leverages semantic similarities between memory chunks. Stochasticity, useful in exploring solution spaces, results from a noise term that makes retrieval probabilistic.

While these soft characteristics of cognition seem to have little to do with learning to perform tasks (i.e., taking instructions and executing them), purely symbolic information processing is not enough in naturalistic settings. Intelligent behavior, by definition, is ill-defined; otherwise traditional algorithmic or optimization techniques could be applied, as has happened in many domains, from chess to program trading. Performing tasks effectively in the real world requires the very characteristics that human cognition has evolved in that environment. For instance, adversarial behavior requires sensitivity to the statistics of the opponent's actions (Lebiere et al. 2003). Moreover, learning complex patterns of events requires associative links to capture context sensitivity (Lebiere and Wallach 2001). Generalization to similar situations is essential for making decisions (Sanner et al. 2000) and controlling dynamical systems (Wallach and Lebiere 2003) in complex environments, where exhaustive or even extensive experience is impractical to achieve. Stochasticity is essential not only to avoid being exploited in adversarial situations (e.g., West and Lebiere 2001) but to explore complex spaces and plan solutions in combinatorial environments (e.g., Lerch et al. 1999). Even tasks that can be defined in a purely symbolic manner (e.g., learning arithmetic tables in school) leverage cognitive characteristics such as stochasticity, adaptivity, and generalization to overcome our cognitive limitations (Lebiere 1999). Indeed, ontologies can be most effectively integrated in cognitive architectures by representing not only the symbolic knowledge itself but by expressing the underlying statistics and regularities of the knowledge. Those regularities can be represented using the subsymbolic parameters that control retrieval and application of that knowledge (Oltramari and Lebiere 2013). Next I will discuss in more detail how these various characteristics can be integrated into ITL.

Symbolic Procedural Skill

While knowledge contributes crucially to performance in many tasks, skill is even more essential. Skill refers to the procedural control of behavior involving both external actions (e.g., manipulating objects in the environment) and internal operations, which include a broad set of internal cognitive capabilities. The ability to manage attention, given well-known limitations in scope and time, requires paying attention to the right things at the right time. The maintenance of appropriate context involves strategies for allocating limited working capacity to preserve the relevant parts of the task representation at all times. The control of knowledge requests from long-term memory involves constantly deciding what type of knowledge is relevant, how to access it most efficiently, and how to make use of it to perform required inferences, such as

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.

anticipating future developments. These three components, and related abilities, are often referred together under the term *situation awareness* (Endsley 1995). Those basic functionalities support key cognitive skills that implement complex task functions. Decision making involves selecting among competing courses of action, based on estimations of their outcomes. Planning results in a potentially complex sequencing of actions that alter the state of the environment, by oneself and potentially in concert with other agents. Managing communication with other agents requires skills that involve not only the generation and understanding of natural language, including resolving references to objects and events relevant to the task, but also social skills such as taking another's perspective (theory of mind), reasoning about their intent, or engaging in persuasive arguments.

Since the origin of cognitive architectures, the most popular form for representing procedural skills has been *production rules* (Newell 1973a). Formulated as condition-action pairs, production rules are not logical rules but rather tiny pieces of skills that have to be acquired in each step of the task learning process and then assembled at run time into coherent behavior. The flexibility and modularity of the framework is its major source of power in generating robust real-world behavior. Acquiring and assembling those skills can be a challenging and time-consuming process. Sometimes skills originate by random trial and error. Other times they originate through a more directed search process. The most direct process of skill acquisition is learning from step-by-step instructions that specify precisely how to accomplish goals (e.g., Taatgen et al. 2006). While other approaches to learning are possible, they tend to be complementary to the instructional process. For instance, if only the goal is specified in the instructions, a problem-solving process can be added to generate the instruction steps. Similarly, learning by demonstration can be conceptualized as requiring an additional process to generalize observations into instructional steps.

The path for this direct instruction process that is supported by the cognitive architecture is to learn the instructions declaratively (e.g., from interacting with a teacher), then retrieve each step from memory one by one and interpret it. The cognitive architecture, through a process called proceduralization (Anderson 1987), then speeds up the process of instruction interpretation and execution in at least two ways. The first, which applies primarily to the interpretation aspect, is to compile away the retrieval of each instruction step into a new unit of skill (a production rule) which directly applies that step without the need for explicit retrieval from memory. For instance, with enough practice, we can type in a password to one of our devices without having to recall it consciously: we just let our fingers "do the walking." The second step applies primarily to the execution aspect: consecutive steps are assembled into a single production rule, thereby making the process increasingly efficient (within the limit of cognitive resource conflict and external task structure, of course) by constructing a set of increasingly complex hierarchical task skills.

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.

Subsymbolic Procedural Skill

The challenging part of skill acquisition, therefore, is not the specific process of creating the skill itself, since the proceduralization process is automatic, but rather the creation of the mental environment in which that process takes place. This involves either knowledge-rich methods (e.g., reasoning and inference) as well as weak methods (e.g., trial-and-error or means-ends analysis) or the much more efficient method of receiving and retrieving the proper task instructions through interaction with an expert teacher. Once new skills have been acquired, applying them properly is not a trivial matter. The new production rules have the same conditions of applicability as the previous, less efficient rules and will have to demonstrate their effectiveness. As for declarative knowledge, production rules are not selected purely symbolically but utilize subsymbolic processes similar to memory activation. The basis for selecting competing rules is that of utility, acquired for each rule according to reinforcement learning processes (e.g., Fu and Anderson 2006) that reflect the effectiveness of each rule in attaining positive rewards. As with memory activations, production utilities include a stochastic component which makes rule selection probabilistic and allows for an exploration-exploitation trade-off that allows for the most efficient skill to emerge gradually and be selected reliably.

A final challenge is that the context of the current problem might not necessarily perfectly match production rule conditions, for instance, because of minor variations in the task environment. As for memory retrieval, a partial matching process has been integrated in the production utility calculus to enable approximate matching of rule condition to the situational context. Analogous to how a matching penalty reduces chunk activation in the retrieval process, imperfect matching to a production condition reduces its utility by the degree of mismatch (Best and Lebiere 2006). What results is an adaptive process of skill application: production rules expand their range of generalization as long as they are successful; then they contract it once they reach situations for which they are poorly suited. Therefore, the process of task learning is not characterized by an all-or-none development of perfect skills; it is a gradual process in which independent skills are acquired and composed to find their range of applicability. Below I will describe two instances of how this approach can be applied to an experimental task.

Learning Complex Decision Making

Depending on the nature of both the task and the interaction, ITL can take on many distinct forms (see Wray III et al., this volume). Different tasks—from decision making in abstract spaces, such as games, to the manipulation of physical objects in 3D space—display fundamentally different characteristics, which will be reflected in the knowledge and skills that are acquired.

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.

Similarly, modes of interaction—from passive instruction interpretation to bidirectional building of common ground through interaction in shared physical spaces—involve distinct processes and representations (Chao et al. 2011). To illustrate the concepts introduced above and sketch out the space of mechanisms involved, I will describe two ways in which these tasks involve substantially different domains and processes to demonstrate the generality of the approach.

The focus is on procedures for general decision making, independent of domain, and was initially defined in the context of a human–robot interaction scenario (Lebiere et al. 2013a). The task to be learned is framed as making arbitrarily complex decisions by decomposing them into a complex process that involves simpler decisions. The scenario is represented by a set of quantitative variables, each describing the value of a particular feature of a specific object or entity. The perceptual and attentional processes that lead to the encoding of each scenario are not explicitly represented. Instead, each situation variable for a given scenario is encoded as a separate chunk in memory, representing the outcome of the encoding process. Chunk activation values are set high enough to assume that access to the situation chunks is fast and reliable. This approach abstracts away from potential perceptual issues, such as grounding problems, that would require distinct solutions, such as perceptual learning.

The decision process is represented by a treelike chart that involves flowing through a set of basic decisions which lead to a number of actions. For each basic decision in the process decomposition, the situation variables entering into that decision are explicitly specified as part of the instructions. Again, the entire decision structure is represented as memory chunks, abstracting over an explicit interactive instruction process. While specific issues might arise as part of the interaction itself, there is nothing inherent in the approach that would present a limitation, as the sequential instruction retrieval process could easily be replaced by an interactive dialogue with an instructor.

Each instruction chunk in memory represents one particular step (e.g., encoding a situation variable, making the decision, or deciding on the next decision to make based on its outcome) of a basic decision process. Again, activation values of the instruction chunks are set sufficiently high as to ensure that they can be reliably retrieved.

The overall decision process (see Figure 4.1) follows the general pattern laid out in the previous section: each individual instruction step must be retrieved before a corresponding action can be taken. Each basic decision entails the following steps:

1. The value of each relevant situation variable is retrieved from memory.
2. A decision chunk is gradually composed.
3. Once that chunk is complete, a decision is made.
4. The outcome of that decision serves as the basis to move on to the next decision.

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.

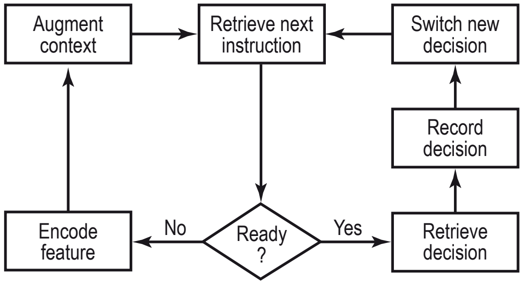

Figure 4.1 Decision-making process for instruction execution.

Each decision is made not according to a specific rule, but by leveraging memory retrieval processes to generalize from a set of previous decision instances held in memory. Those decisions could have resulted from demonstration by experts as part of the instruction process or experienced by the learner as part of a trial-and-error process. This process of instance-based learning (Gonzalez et al. 2003) leverages both the symbolic representation of previous experience in memory and their subsymbolic activation parameters. Thus, the task learning process is a combination of instructions and experience, reflecting a pervasive pattern in both formal education and field training. The result of that overall decision process generates a set of outcomes that closely reflects how humans gain expertise: the model follows the same decision procedure as experts, as specified by the instructions, and has also been trained from their decision instances.

Nothing about this instructional decision-learning process is specific to the domain or scenario. We applied the same instructional learning engine to an entirely different task of making probability judgments among hypotheses in a geospatial sense-making domain and obtained excellent results that match human performance, in particular, by exhibiting similar types and degrees of cognitive biases (Lebiere et al. 2013b; Thomson et al. 2015). The structure of the task consisted of a sequential presentation of spatial layers of information: each layer resulted in a probability judgment about a discrete set of hypotheses. Subjects were then instructed on how to combine judgments from the various layers to refine their probability distribution over the set of hypotheses. Those instructions were converted into the format described above. Only two minor generalizations to the approach used in the first domain were necessary. The first extended the test on the result of an intermediate decision to quantitative nonbinary outcomes. The second extended the set of values used in decisions to include the outcome of previous decisions in addition to the external environment variables. These changes were not domain specific in any way but rather increased the generalization of the instructional process. Interestingly, while the decision process in the original domain was specified in tree form, it could be generalized in this domain to a graphlike iterative procedure without

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.

requiring any change because of the nature of the representation of instructions as independent memory chunks. This illustrates the power of cognitive techniques to provide humanlike flexibility.

Conclusion

Computational models offer a precise, quantitative way to represent the cognitive processes and representations involved when an agent interacts with another agent: from the receiving of instructions to their interpretation and the processes involved in learning to perform the task. Computational models can be used in a number of ways to predict human performance in interactive learning tasks, to design ITL systems that work naturally with humans, to implement intelligent assistants that can learn as a human assistant would, as well as to implement intelligent agents and robots that behave as a human instructor.

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.