

---

# Nonverbal Task Learning

---

**Maithilee Kunda**

MKUNDA@VANDERBILT.EDU

Department of Electrical Engineering and Computer Science, Vanderbilt University  
PMB 351679, 2301 Vanderbilt Place, Nashville, TN 37235-1679 USA

## Abstract

Nonverbal task learning is defined here as a variant of interactive task learning in which an agent learns the definition of a new task without any verbal information such as task instructions. Instead, the agent must 1) learn the task definition using only a single solved example problem as its training input, and then 2) generalize this definition in order to successfully parse new problems. In this paper, we present a conceptual framework for nonverbal task learning, and we compare and contrast this type of learning with existing learning paradigms in AI. We also discuss nonverbal task learning in the context of nonverbal human intelligence tests, which are standardized tests designed to be given without any verbal instructions so that they can be used by people with language difficulties.

## 1. An Anecdote

Several years ago, while visiting researchers at Zoo Atlanta, I was able to observe one of the resident orangutans performing a cognitive task on a touch screen, as in Figure 1. The task used a delayed match-to-sample design: the orangutan first saw a single image, and then after a short delay, he had to select the matching image from a set of choices, as illustrated in Figure 2. (As I recall, the actual task used photographs of other orangutans, though this figure just uses shapes for clarity.)



Figure 1: Orangutan working on cognitive task (Zoo Atlanta, 2014).

On the first trial, the orangutan touched the correct matching item in the top-right quadrant of the screen. He then received, and happily consumed, a food reward. On the second trial, however, he did not select the correct item, instead choosing the distracter in the top-right quadrant.

*Poor guy*, I thought. *No food for you*. But then, he got another chance; because he was still in the “training phase” of this particular experiment, the same trial was presented again, except with two of the incorrect distracters removed. So now the orangutan had a 50/50 chance...or so I thought. However, instead of selecting one of the two visible images, he again poked at the top-right quadrant, even though it was completely black. He was very insistent, jabbing repeatedly at that spot several times. Unfortunately, as they say, there was no “there” there.

Two things struck me about this incident. First, the orangutan apparently did not fully understand the *format* of the task. It is easy for us to look at the final item presented in this task and say, “Now there are two available responses to choose from,” but of course, that is not actually true. While there are only two actions that we consider to be valid responses, there are, in fact, an infinite




Figure 2: Simplified illustration of orangutan behavior once observed during cognitive testing.

number of other actions the orangutan could take. Even though he did seem to understand part of what was needed (i.e., to touch a single point on the screen with one finger), he apparently was not grasping the idea that the contiguous chunks of non-black pixels each represented a valid option, and that the black regions should be disregarded.

Second, even after correctly completing the first item of the task, the task *goal* was still under-determined by what had taken place so far. While we cannot say for sure, it certainly seems like the orangutan may have thought the goal of the task was to touch that particular x-y location on the screen. And why not? After the first item alone, and without additional priors, there was really no reason to suppose that the task goal was about visual similarity and not about spatial location.

## 2. Key Concepts of Interactive Task Learning

How can intelligent agents learn the definition of a task, i.e., the task format and goal, without being verbally instructed with that information, and using only a single observed example problem as the training input? We call this problem *nonverbal task learning* (NTL).

NTL is a variant of interactive task learning, an area of AI research that investigates how “an agent actively tries to learn the actual definition of a task through natural interaction with a human instructor, not just how to perform a task better” (Laird et al., 2017, p. 2). A lot of research in interactive task learning involves designing AI systems or robots that learn from both verbal and nonverbal information, i.e., instructions along with examples or situated experiences (e.g. Hinrichs & Forbus, 2014; Kirk et al., 2016). Such multi-modal inputs are used all the time in human learning; for instance, we might teach a friend how to play a new board game by explaining the rules while demonstrating example moves with various pieces.

However, people can also learn new tasks without verbal instructions. While everyday instances are hard to find due to the ubiquity of language, we can imagine teaching a new task to someone who speaks a different language, or to someone in a very noisy environment that precludes speech. In these language-free situations, we can still teach a new task by providing an example to the learner, though without language, the learner needs to have more sophisticated learning capabilities to infer the intended task format and goal in a way that can be generalized effectively to new problems.

Before discussing the problem of nonverbal task learning (NTL) in further detail, we first define some basic terminology related to tasks and problems and present concepts related to task learning in general.

## 2.1 Terminology

**Definition 2.1.** A **task**  $T$  is a set of problems  $p_i$  that share a task definition, which consists of:

- A **task format**: the form of valid inputs (a subset of possible states of the world) and outputs (a subset of possible actions available to an agent).
- A **task goal**: the characteristics of outputs, or of resulting world states, that represent desired solutions to problems in  $T$ .

For instance, consider the visual delayed-match-to-sample task described in Section 1. Valid task inputs consist of a single image, a time delay, and then a set of multiple images, one of which matches the initial image. Valid task outputs consist of an agent taking some action (e.g., touching a location on a screen) to select one of the final set of multiple images. The task goal is defined as the condition in which the selected image visually matches the initial image. Table 1 lists this task along with additional example tasks broken down according to this definition.

As observed by Laird et al. (2017), the vast majority of work in AI to date has assumed the task definition as a given, focusing instead on how to select the correct, goal-fulfilling outputs from among a set of valid output options. However, for humans and other animals, or for artificial agents that aim to learn through interaction with the world and not through direct programming, the challenge of learning, and successfully applying, a task definition is a significant one.

## 2.2 Task format: Valid inputs and outputs

A conventional AI system designed to perform a task  $T$  will only receive valid inputs, or might throw an exception for invalid inputs. Humans, in contrast, have to actively maintain a representation of the valid input format for a task, and we use this representation to parse the continuous stream of sensory inputs we receive into a structured instance over which we can reason. We too can generate an “invalid input exception,” but we can also modify our representation of task inputs as needed, which is a potent aspect of our task learning abilities.

For instance, imagine if you sat down to play tic tac toe with someone, and they presented you with the image shown in Figure 3a. Most people would probably decide this is not a valid input for

Table 1: Various tasks described using the terminology in Definition 2.1.

<b>Task</b>	<b>Task format: Inputs</b>	<b>Task format: Outputs</b>	<b>Task goal</b>
<i>Delayed match-to-sample</i>	Single image + delay + multiple images	Select one image from final image set	Selected image should visually match initial image
<i>Shortest path</i>	Connected graph with real-valued edge values + two nodes	Select subset of edges that connects the two nodes and is non-branching	Selected subset of edges has minimal total edge cost
<i>Tic tac toe</i>	3-by-3 grid + opponent + assigned mark	Take turns writing one assigned mark in empty grid square	Grid has contiguous row or column of own marks before opponent’s marks
<i>Object recognition</i>	Image depicting an object + set of labels	Select one label from set of labels	Selected label describes object in image




Figure 3: Input variants for tic tac toe that many people would probably consider: (a) invalid; (b) valid, and easily generalized from the original game; (c) valid, though not as easily generalized.<sup>1</sup>

the task of tic tac toe (and probably a tic-tac-toe-playing AI system would return an invalid input exception as well). However, even though tic tac toe, by definition, involves a three-by-three grid, most people would probably have no trouble accepting a four-by-four grid as a valid input problem, as shown in Figure 3b. Going further, we can imagine other, less similar variants, like that shown in Figure 3c, that at least some people would likely bend their task input representation to fit.

### 2.3 Task format: Valid outputs

Conventional AI systems generally have task output formats directly programmed, but this too is something that humans must actively represent. Consider the example of the orangutan given in Section 1. He did not have an effective representation of task output formats, i.e., actions to consider taking. It seemed that he considered touching any part of the screen to be a valid output, whereas most of us would define the task output format as selecting any one of the visible images.

Of course, he could have done much worse in many ways; he could have touched multiple parts of the screen, or swiped the screen, or vocalized at it, or tried to break it. He clearly had gotten the idea of a single-finger touch as being the right type of action to take. It is likely that learning task output formats is closely tied to our knowledge of the affordances of various objects in our environment. Representing task output formats in terms of affordances likely enables us to generalize our task definitions to very different settings.

One task output format that shows up again and again in tasks for humans, animals, and AI systems is the concept of multiple choice, i.e., selecting from a set of the same type of action applied to each of a group of alternatives. Like other aspects of task definitions, representing the concept of multiple choice response is likely much more complicated for humans and other animals than it is for AI systems that are explicitly programmed with this information.

### 2.4 Task goals

There are many different types of task goals. In some cases, task goals are represented by a concrete, fully specified task state. For example, the Tower of Hanoi task involves moving disks on pegs until

<sup>1</sup>Incredibly, I am not the only person to dream up this variant...though not a very interesting game, as it turns out: <<http://redfrontdoor.org/blog/?p=1497>>




Figure 4: Illustration of an example problem and solution (left), from which learning a somewhat abstracted goal could enable you to solve new problems (right).

a known, well-defined configuration is reached. In other cases, task goals are represented by a partially specified task state such as in tic tac toe or chess, in which a concrete goal condition holds over part of the task state (e.g., the king is dead) but the rest of the state is unspecified.

In still other cases, the task goal involves more abstract concepts that cannot be defined purely in terms of concrete conditions over the task state, and that instead require considering relationships over task elements. For instance, consider the example problem and solution shown on the left of Figure 4. From this, you can likely induce a task goal that would enable you to imagine a solution to the unsolved problem shown on the right. Your goal is likely represented in terms of certain abstract visuospatial relationships among the elements in each image having to do with the concept of “fit.” Representing such goals in a way that is both learnable and generalizable is an interesting challenge for AI systems, which we return to later in this paper.

### 3. Nonverbal Task Learning

Using the terminology discussed in the prior section, we now return to our original problem of nonverbal task learning. We begin with the interactive example shown in Figure 5.

Suppose someone taught you a new task by showing you the example problem and solution illustrated on the left of this figure. Now, take what you have learned from this example, and solve the five new problems given on the right of the figure. What do your solutions look like?




Figure 5: Interactive example of nonverbal task learning.

Here are some surmised solutions:

1. Most people would likely draw a dot in the middle of the blue circle.
2. Most people would likely draw a dot into the middle of each blue square. (The author's spouse said he would either do that or, "calculate the centroid of all the shapes and put a dot there.")
3. Many people would likely draw a dot into the middle of each blue shape. Some might put a dot into the middle of each of the four shapes.
4. Most people would likely decide this is not a valid input.
5. Many people would likely decide this is not a valid input. Some might "invert" the task and draw squares around each dot.

This example illustrates many interesting properties of nonverbal task learning (NTL) that we discuss below.

**NTL exists.** Humans can do this type of learning, as you have just demonstrated. With no verbal instructions about the task, you inferred a task format and goal, and used this representation of the task definition to parse new problems that you were then able to solve. (The figure did include some text labels, it is true; please imagine doing this task in an interactive setting, in which you watched the demonstration and then were presented with the new problems one at a time.)

**Learning continues within a task.** As discussed previously, humans can readily modify their task definitions to accommodate new problems. In the series shown in Figure 5, your initial task definition learned from the solved example on the left might have been something like, "Put a dot in the blue square." When the first new problem arrives, there is no blue square, and so perhaps you modify your representation of the task input format to include any blue shape. When the second problem arrives, you modify your task input format again to accommodate the presence of multiple shapes, and you also modify your task output format to allow for taking multiple actions instead of just one. And so on. In other words, NTL is a clear manifestation of within-task transfer, in which what is transferred are aspects of the learned task definition.

**Ambiguity often remains.** For some of the problems on the right of Figure 5, you could argue for several different versions of a task definition, all of which are justifiable. For example, in problem 3, do you put a dot in the orange square or not, i.e., is the task goal to put dots in all shapes, or in all blue shapes? From what has preceded, there is really no "correct" answer in this situation.

**Prior knowledge is necessary.** Because the learning input is so sparse in terms of its information content, there is no way to perform NTL without significant prior knowledge—otherwise the ambiguities are overwhelming. Taking Figure 5, a perfectly good set of solutions could have been to place a dot in the center of each outer square, regardless of the visual content within. (Likewise, for the orangutan example in Section 1, the orangutan's solution of touching the same x-y location on the screen that was previously rewarded was a reasonable approach.) Our prior experiences with shapes, games, visual reasoning test problems, and many other cultural and developmental situations likely feed into the inductive biases that we bring to bear during NTL.

**Learning continues across tasks.** Just as task definitions are updated across problems within a given task, we undoubtedly take parts of task definitions and reuse them across tasks. General task concepts such as turn-taking and multiple choice are likely reused and reinforced as a person's task repertoire grows. And so in addition to more general forms of prior domain knowledge, the knowledge about task definitions that we bring from old NTL scenarios into new ones is also key.

In other words, in addition to within-task transfer, NTL also involves across-task transfer of aspects of previously learned task definitions.

**Social reasoning plays a deep role.** Why would a person do NTL in the first place? I.e., having seen someone put a dot in a blue square with a pencil, why would you even bother to stick around and do anything, let alone meticulously put dots in shapes? The answer is that is a lot of social context involved in NTL, including our propensity to imitate, inferring the intentions of others, our perception of affordances and functional artifacts in our environment, etc. These are deep issues for cognitive systems that we do not pursue further in this paper, though they are important to keep in mind as a critical ingredient for NTL.

**NTL is interactive, though asynchronously so.** Following the prior point about social reasoning, it is worth noting that NTL is indeed interactive, in that the example problem and sequence of additional problems has been designed purposefully by an intelligent agent separate from the NTL agent. For example, the problems in Figure 5 were designed by me, the author, and later attempted by you, the reader. Even though the interaction is not taking place synchronously and in real time, the NTL context for you is still interactive, though our interaction is asynchronous. This is why we consider NTL to be a variant of interactive task learning.

## 4. Applicability

If NTL only involved contrived example tasks and problems like that shown in Figure 5, it would not be a very interesting research challenge for AI. However, there are at least three real settings in which NTL plays an important role.

### 4.1 NTL as a component of integrated interactive task learning

As mentioned in Section 2, humans usually learn new tasks using both verbal and nonverbal inputs. It is useful to explore NTL as a limit case of how interactive task learning can work without language, as many of the same NTL mechanisms are likely active during multimodal (verbal + nonverbal) interactive task learning. As Laird et al. (2017) observe, “many challenges remain in developing cognitive models of task learning. One of the most important is moving beyond instruction following to other common forms of learning, such as learning by example” (p. 13).

### 4.2 NTL in people who have language difficulties

There are many clinical populations in which individuals have difficulties in using or understanding language, including, for example: individuals with an acquired aphasia (cognitive impairment specific to language) due to strokes, traumatic brain injuries, tumors, or other neurological conditions (National Aphasia Association, 2019); and children with language disorders that may or may not be tied to comorbid conditions such as, “Down syndrome, fragile X syndrome, autism spectrum disorder..., and being deaf or hard of hearing” (National Academies of Sciences, Engineering, and Medicine, 2016). Individuals with language difficulties may rely more on nonverbal task learning mechanisms like those in NTL than do individuals with full language ability. In some cases, individuals with severe language impairments may have access only to nonverbal task learn-

ing mechanisms. Thus, developing cognitive models of NTL will be critical for understanding how task learning works in these populations.

### 4.3 NTL in non-human animals

Most cognitive research with non-human animals does not expect its subjects to learn new tasks through one training example. Instead, animals often undergo intensive and repetitive reinforcement-based training sessions to learn new tasks. For example, one study trained pigeons to perform transitive reasoning by pecking differently shaped buttons; training including autoshaping (training the pigeon to peck at buttons, i.e., essentially teaching an aspect of the task output format) followed by 125 sessions of 40+ stimulus trials with reinforcement (Von Fersen et al., 1991). After all of this, only 4 of the 6 pigeon subjects had learned the task to a criterion of 80% accuracy.

Sometimes, however, an animal shows successful performance on the very first trial of a new task. How can this be? From a study of orangutans at Zoo Atlanta (Talbot et al., 2015, p. 180):

One interesting possibility is that experience with cognitive testing leads to improved performance on this type of task. While our sample was too small to test this further, it is worth noting that our most successful subject, Madu, who distinguished familiar individuals on the first presentation, had more extensive testing history with the matching-to-sample paradigm than the other orangutans.

Other studies also find that individual animals more experienced with cognitive testing in general often learn faster and/or perform better on new tasks than less experienced individuals do (Brosnan et al. 2011; Parr et al. 2000; Vonk & Hamilton 2014, as cited in Talbot et al. 2015). Congruent with this idea, it turns out that the pigeons in the transitive reasoning study were “experimentally naive,” meaning they had not previously participated in these kinds of cognitive research studies and learned these kinds of tasks. Across-task transfer of task formats and goals could explain instances of successful NTL observed in some animal subjects.

## 5. Nonverbal Intelligence Tests as a Task Domain for NTL

Most human intelligence tests use spoken or written instructions to inform test-takers how to solve test problems, i.e., the task definition is given explicitly. Such tests would be a good proving ground for techniques in interactive task learning that use both verbal and nonverbal information.

Nonverbal intelligence tests, in contrast, are specifically designed to avoid verbal instructions altogether, so that they can be taken by individuals with language difficulties (DeThorne & Schaefer, 2004).<sup>2</sup> The Leiter International Performance Scale (Leiter), the Universal Nonverbal Intelligence Test (UNIT), and the Test of Nonverbal Intelligence (TONI) are all examples of such nonverbal intelligence tests.

In these tests, examiners initially show test-takers a simple example problem and its solution. Test-takers must learn the task definition (e.g., matching shapes, finding one shape in another, completing a visual pattern, etc.) by observing the example, and then use this knowledge to solve a

<sup>2</sup>Note that the term *nonverbal intelligence test* is sometimes used to refer to a “nonverbal test of intelligence” that does not require verbal instructions, as defined here, but other times is used to refer to a “test of nonverbal intelligence” that taps into nonverbal reasoning abilities but may require verbal instructions to be administered.






Figure 6: Two examples of the kinds of tasks often found on nonverbal intelligence tests—a “matrix reasoning” task (left) and an “odd-one-out” task (right)—each with a solved example problem ( $E$ ) and two unsolved test problems ( $p_1$  and  $p_2$ ).

series of more difficult test problems. No reinforcement is given for the test problems, though tests are often stopped once a test-taker has gotten a certain number of consecutive problems wrong. In some cases, examinees are supposed to use gestures and facial expressions to help the test-taker learn each task (DeThorne & Schaefer, 2004), which falls somewhat outside our narrow definition of NTL but is worth noting as a potentially important source of social cues.

Figure 6 illustrates the types of tasks that might be included in a nonverbal intelligence test. Most are the usual sorts of memory and reasoning tasks that one would expect to see on an intelligence test, but both the example problems and the sequences of test problems are cleverly designed to facilitate successful initial task learning and continued adaptation of the task definition as needed, as problems become more complex.

Using nonverbal intelligence tests as a task domain for NTL research is promising for several reasons. There are a relatively large number of tests and problems; for example, one version of the Leiter test has 20 different subtests, each representing a different task. The tests are easy to administer to human participants and there is also a significant body of already published studies that use these tests in a wide variety of human studies, making comparisons between AI and human performance relatively straightforward and potentially quite rich in research implications.

In addition, these tests have been carefully designed and refined over many decades, and so the example problems and test problem sequences are very well crafted to facilitate NTL, i.e., learning the task definition from the single provided example problem and solution and successfully updating the task definition across the sequence of test problems. Thus, NTL systems would be well poised to answer questions about inter-problem learning on such tests, which have not been studied very much in prior AI research on intelligence tests. Psychology research has shown that inter-problem learning contributes significantly to people’s ideas of the task definition; one study demonstrated how providing different initial problems could bias children to give different answers to identical later problems, where the later problems were designed to have two different “correct” answers depending on two possible interpretations of the task goal (Kirby & Lawson, 1983).

NTL research on such tests would also be well poised to answer questions about what kinds of background knowledge are needed by a test-taker, and how different sets of background knowledge might lead to different outcomes, despite equivalent “cognitive” abilities. For instance, despite historical arguments that certain nonverbal reasoning tasks were “pure” measures of intelligence not requiring any particular background knowledge, and thus could be used as “culture-free” or

“culture-fair” tests of intelligence, these claims are no longer widely accepted, as evidence shows a significant effect of culture and background knowledge in such tasks (Lohman, 2005).

NTL research could also provide insights into trajectories of child development. For example, the task shown on the left of Figure 6 is similar to the Raven’s Progressive Matrices family of tests, which technically call for verbal instructions but are sometimes also administered in nonverbal fashion. The Colored Progressive Matrices is the easiest version of the test and is frequently used as a measure of intelligence for children, the elderly, or individuals with cognitive or developmental disabilities (Raven et al., 1998). Interestingly, the CPM test manual observes that up to the ages of 4-5, some (but not all) children “have grasped the idea that they have to fill the gap in the pattern,” i.e., discovered the CPM task goal (Raven et al., 1998, p. CPM-13). Some children who have grasped this concept nevertheless fail to understand more complex relationships among problem elements that appear on later problems in the test, beyond just visual similarity. To give an idea of CPM difficulty, six-year-olds will typically answer about 40% of the problems correctly, eight-year-olds will answer about 60% correctly, and ten-year-olds will answer about 75% correctly. To what extent do these differences reflect developmental shifts in visuospatial reasoning abilities, like mental rotation, versus task learning abilities, like being able to represent and adapt the task goal in increasingly sophisticated ways? And to what extent might these two sets of abilities be developmentally linked? These are fascinating questions, and NTL research could undoubtedly contribute interesting computational insights for these questions.

As a final note, almost all current nonverbal intelligence tests fall within the domain of visuospatial reasoning. (In other words, they are essentially nonverbal tests of visuospatial intelligence.) Some efforts have been made to create tactile versions of tests (Reid, 2002; Rich & Anderson, 1965) for use in populations with visual impairment but without language difficulty, and so verbal instructions can be given. Could nonverbal test designs be adapted to other reasoning domains or modalities, such as using auditory stimuli? This is an important open question, not just for understanding other forms of reasoning in the general population, but also for improving the reach of available testing procedures for special populations. Working on such tests would thus represent a valuable extension of NTL research in AI.

## 6. Related Work

Nonverbal task learning (NTL) brings together major themes and challenges from several other areas of AI and cognitive science research, which we describe below.

**Learning by observation.** Within learning by observation, sub-areas of research include learning different aspects of the task definition, such as the task goal, initial conditions, constraints on actions, action policies, etc. Many papers study how to use techniques for analogy and generalization to extract essential information about a problem-solving procedure from one initial problem in order to solve new problems (Tecuci & Kodratoff, 1990; Wilson & Scheutz, 2014).

However, the particular problem of learning the task goal by observation is especially difficult, because the agent must somehow generalize a goal concept from concrete examples. As one paper on learning by observation pointed out, “A tough question is where the task criterion comes from. Ideally, it should be learned from observation. The learner should infer the intent of the teacher.

This is very difficult, and we defer addressing this question by manually specifying a task criterion” (Bentivegna et al., 2004, p. 166).

This problem is made easier if the goal can be defined as a concrete instead of general concept: for example, in one study of one-shot imitation learning, tasks had different goals that were learned by observation, where each goal was defined as a concrete state that had to be reached from an arbitrary starting state (Duan et al., 2017). Likewise, there are many approaches for learning the task goal when it can be represented as an objective function over states, as in inverse reinforcement learning (Abbeel & Ng, 2004). Other research has looked at how a combination of observation and instruction can be used to obtain more general task goal concepts from a concrete starting point, with instruction playing a key role in identifying what parts of the goal concept should be generalized (Kirk & Laird, 2014). A very recent paper uses prior conceptual knowledge to infer somewhat abstract goal concepts from single, visually observed solved examples (Lázaro-Gredilla et al., 2019).

**Analogy and transfer learning.** By its nature, learning by observation is inherently analogy-based: something is learned by observing an example—the analogical source—and then some part of what was learned must be transferred to enable solving a new problem instance—the analogical target. Research in virtually all areas of cognitive science emphasizes the important roles of analogy and transfer in human intelligence (e.g., Gentner et al., 2001; Kolodner, 1992; Lakoff & Johnson, 2008; Nersessian, 2008, , etc.), including in the specific context of nonverbal reasoning and intelligence testing (Campione et al., 1985). An interesting methodological critique has pointed out that, while some human studies do fail to find evidence of transfer, such studies often use artificial tasks that have been stripped of any inherent meaning, in order to control for effects of background knowledge (Campione & Brown, 1984). “It is difficult to see how subjects *could* demonstrate lateral transfer within an essentially arbitrary domain. Transfer is frequently assumed to be a consequence of comprehension...of understanding how the form of the solution is related to, or follows from, the structure of the problem” (Campione & Brown, 1984, p. 271). In the case of NTL, this observation raises a very important question: what background or domain knowledge is necessary to enable effective transfer, and thus successful performance, on new tasks?

**Explanation-based learning.** The need for background knowledge to enable more effective transfer in NTL is reminiscent of another area of AI research: explanation-based learning (EBL). In EBL, domain knowledge is used to “explain” relationships between features and labels in the context of supervised learning, thus promoting more effective generalization from a fewer number of examples (DeJong & Lim, 2011). The connection between EBL and interactive task learning has been made previously (Mohan & Laird, 2014).

## 7. Conclusion

We have presented a discussion of the problem of nonverbal task learning, in which agents must learn the definition of a new task through a single solved example problem, and then generalize this definition to solve additional new test problems. AI research on NTL would be valuable for advancing the capabilities of interactive task learning systems, as well as for improving our understanding of NTL-like processes in human cognition, in both neurotypical populations as well as populations

with various cognitive conditions. Nonverbal intelligence tests, designed to be administered without any verbal instructions, provide an ideal testbed for NTL research, and we encourage the AI community to take a look at these very interesting tests.

## Acknowledgements

Thanks to the reviewers for their comments and pointers to relevant literature. Thanks also to Calum Hartley whose research on verbal and nonverbal cognition in autism catalyzed this line of inquiry into nonverbal task learning. This work has been funded in part by NSF Award #1730044.

## References

- Abbeel, P., & Ng, A. Y. (2004). Apprenticeship learning via inverse reinforcement learning. *Proceedings of the twenty-first international conference on Machine learning* (p. 1). ACM.
- Bentivegna, D. C., Atkeson, C. G., & Cheng, G. (2004). Learning tasks from observation and practice. *Robotics and Autonomous Systems*, *47*, 163–169.
- Brosnan, S. F., Parrish, A., Beran, M. J., Flemming, T., Heimbauer, L., Talbot, C. F., Lambeth, S. P., Schapiro, S. J., & Wilson, B. J. (2011). Responses to the assurance game in monkeys, apes, and humans using equivalent procedures. *Proceedings of the National Academy of Sciences*, *108*, 3442–3447.
- Campione, J. C., & Brown, A. L. (1984). Learning ability and transfer propensity as sources of individual differences in intelligence. *Learning and cognition in the mentally retarded*, (pp. 265–293).
- Campione, J. C., Brown, A. L., Ferrara, R. A., Jones, R. S., & Steinberg, E. (1985). Breakdowns in flexible use of information: Intelligence-related differences in transfer following equivalent learning performance. *Intelligence*, *9*, 297–315.
- DeJong, G., & Lim, S. H. (2011). Explanation-based learning. In *Encyclopedia of machine learning*, 388–392. Springer.
- DeThorne, L. S., & Schaefer, B. A. (2004). A guide to child nonverbal iq measures. *American Journal of Speech-Language Pathology*, *13*, 275–290.
- Duan, Y., Andrychowicz, M., Stadie, B., Ho, O. J., Schneider, J., Sutskever, I., Abbeel, P., & Zaremba, W. (2017). One-shot imitation learning. *Advances in neural information processing systems* (pp. 1087–1098).
- Gentner, D., Holyoak, K. J., & Kokinov, B. N. (2001). *The analogical mind: Perspectives from cognitive science*. MIT press.
- Hinrichs, T. R., & Forbus, K. D. (2014). X goes first: Teaching simple games through multimodal interaction. *Advances in Cognitive Systems*, *3*, 31–46.
- Kirby, J. R., & Lawson, M. J. (1983). Effects of strategy training on progressive matrices performance. *Contemporary Educational Psychology*, *8*, 127–140.

- Kirk, J., Mininger, A., & Laird, J. (2016). Learning task goals interactively with visual demonstrations. *Biologically Inspired Cognitive Architectures*, 18, 1–8.
- Kirk, J. R., & Laird, J. (2014). Interactive task learning for simple games. *Advances in Cognitive Systems*, 3, 13–30.
- Kolodner, J. L. (1992). An introduction to case-based reasoning. *Artificial intelligence review*, 6, 3–34.
- Laird, J. E., et al. (2017). Interactive task learning. *IEEE Intelligent Systems*, 32, 6–21.
- Lakoff, G., & Johnson, M. (2008). *Metaphors we live by*. University of Chicago press.
- Lázaro-Gredilla, M., Lin, D., Guntupalli, J. S., & George, D. (2019). Beyond imitation: Zero-shot task transfer on robots by learning concepts as cognitive programs. *Science Robotics*, 4, eaav3150.
- Lohman, D. F. (2005). The role of nonverbal ability tests in identifying academically gifted students: An aptitude perspective. *Gifted Child Quarterly*, 49, 111–138.
- Mohan, S., & Laird, J. E. (2014). Learning goal-oriented hierarchical tasks from situated interactive instruction. *AAAI* (pp. 387–394).
- National Academies of Sciences, Engineering, and Medicine (2016). Childhood speech and language disorders in the general U.S. population. In S. Rosenbaum & P. Simon (Eds.), *Speech and language disorders in children: Implications for the social security administration's supplemental security income program*.
- National Aphasia Association (2019). Aphasia FAQs. From <https://www.aphasia.org/aphasia-faqs/>.
- Nersessian, N. J. (2008). *Creating scientific concepts*. MIT press.
- Parr, L. A., Winslow, J. T., Hopkins, W. D., & de Waal, F. (2000). Recognizing facial cues: individual discrimination by chimpanzees (pan troglodytes) and rhesus monkeys (macaca mulatta). *Journal of Comparative Psychology*, 114, 47.
- Raven, J., Raven, J. C., & Court, J. H. (1998). *Manual for raven's progressive matrices and vocabulary scales*. Harcourt Assessment, Inc.
- Reid, J. V. (2002). Testing nonverbal intelligence of working-age visually impaired adults: Evaluation of the Adapted Kohs Block Design Test. *Journal of Visual Impairment & Blindness*, 96.
- Rich, C. C., & Anderson, R. P. (1965). A tactual form of the progressive matrices for use with blind children. *Journal of Counseling & Development*, 43, 912–919.
- Talbot, C. F., Mayo, L., Stoinski, T., & Brosnan, S. F. (2015). Face discriminations by orangutans (pongo spp.) vary as a function of familiarity. *Evolutionary Psychological Science*, 1, 172–182.
- Tecuci, G., & Kodratoff, Y. (1990). Apprenticeship learning in imperfect domain theories. In *Machine learning*, 514–551. Elsevier.
- Von Fersen, L., Wynne, C., Delius, J. D., & Staddon, J. (1991). Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 17, 334.

- Vonk, J., & Hamilton, J. (2014). Orangutans (*Pongo abelii*) and a gorilla (*Gorilla gorilla gorilla*) match features in familiar and unfamiliar individuals. *Animal cognition*, *17*, 1089–1105.
- Wilson, J. R., & Scheutz, M. (2014). Analogical generalization of activities from single demonstration. *Ibero-American Conference on Artificial Intelligence* (pp. 494–505). Springer.
- Zoo Atlanta (2014). Great apes use touchscreen computers behind-the-scenes at Zoo Atlanta. From <https://www.youtube.com/watch?v=EPfhMU5HQBc>.