AFRL-AFOSR-VA-TR-2019-0204

Interactive Task Learning

John Laird
REGENTS OF THE UNIVERSITY OF MICHIGAN

07/15/2019
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA2
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
12-07-2019	Final Report	15-04-2015; 14-04-2019

4. TITLE AND SUBTITLE Interactive Task Learning	**5a. CONTRACT NUMBER** FA9550-15-1-0157
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) John E. Laird	**5d. PROJECT NUMBER**
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Michigan Ann Arbor, MI 48109	**8. PERFORMING ORGANIZATION REPORT NUMBER**
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research	**10. SPONSOR/MONITOR'S ACRONYM(S)** AFOSR
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Distribution A - Approved for Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Our research on interactive task learning has emphasized the continued development of a natural language understanding system that interfaces with a human instructor, and the underlying task learning system. The language system uniquely combined relevant research: a construction grammar approach to representing linguistic knowledge; an incremental, single-path processing algorithm with local repair; a cognitive architecture as the computational platform; and embodiment in the robotic agent for grounding language to the agent's perception, action capabilities, and knowledge of the world. We extended this to process real world task instructions and the ability to handle multiple forms of ambiguity. The task learning research led to the development of an agent for learning all aspects of tasks, with emphasis on handling ambiguous scenarios. The agent can interactively learning over 55 games and puzzles, and transfers knowledge learned in one game to a similar game.

15. SUBJECT TERMS

Artificial Intelligence, Task Learning

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			John E. Laird
unclassified	unclassified	unclassified	UU	11	**19b. TELEPHONE NUMBER** *(Include area code)* 734 355-9613

Reset

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18
Adobe Professional 7.0

Final Report for Grant FA9550- 15-1-0157, "Interactive Task Learning"

John E. Laird
University of Michigan

I. Grounded Language Understanding in Support of Interactive Task Learning.

Research conducted by Peter Lindes under supervision from John Laird

An important question in cognitive science is how exactly humans understand natural language. An important question in artificial intelligence is how to enable autonomous artificial agents to communicate and cooperate with humans using natural language. In this research, we attack both of these problems jointly by building a language comprehension system embedded within a robotic agent using theoretical principles related to our knowledge of human language processing. This system enables the agent to learn and perform new tasks based on interactive natural language instruction.

For several years, our research group has been developing a software agent called Rosie as an autonomous agent for Interactive Task Learning (ITL; Laird et al., 2017). As part of this effort we have developed, and are continuing to develop, a natural language comprehension system which we call Lucia (Lindes, Mininger, Kirk, & Laird, 2017). Lucia has been developed both to provide the end-to-end grounded language comprehension that Rosie needs to be able to learn and perform new tasks, while doing this in a way that simulates the high-level cognitive abilities that humans use to understand language.

The fundamental purpose of Lucia is to transform each input sentence in natural language into a grounded meaning representation that the Rosie agent can act on to achieve its goals. We have designed its mechanisms in a way that is intended to model how the human mind represents and stores linguistic knowledge, how it processes new input dynamically based on this knowledge, how domain-general cognitive abilities are used in this processing, and how the resulting meanings are grounded dynamically to the agent's knowledge about the environment, its own capabilities, and its perception of the current situation in the world (Lindes, 2018).

In developing this design, we have drawn on several lines of cognitive science research. For representing knowledge of linguistic forms and how they relate to their meanings, we have found work on cognitive semantics and construction grammar to be particularly useful. Human language processing is known to be incremental and to build an interpretation of a sentence in small increments. We adopt the assumption that language processing is a learned skill that uses domain-general cognitive abilities, and that cognitive architectures can be used to simulate these abilities. Key to our approach is that the Lucia comprehension system is embedded within the Rosie agent and grounds the meanings of linguistic forms to the same knowledge that the rest of the agent uses to learn tasks and act in the physical world. All of this is being done within the Soar cognitive architecture.

We describe the work done for this project in three phases. The first phase, from April 2015 through May 2016, involved building a prototype system as a proof-of-concept for the approach we have taken. The second phase, from May 2016 through August 2017, involved expanding the capability of the system to cover a large test set of sentences relevant to the larger Rosie ITL project, including working out solutions to a number of forms of ambiguity. The third phase, from August 2017 through April 2019, has involved integrating fully the Lucia comprehension system with the Rosie agent, extending its coverage to more tasks, and exploring alternative ways to use the capabilities of the Soar architecture for language comprehension. Further work on this third phase is ongoing with follow on funding from AFOSR.

Phase 1, April 2015-May 2016 – Prototype

In this phase, we built a prototype, proof-of-concept system based on theoretical principles derived from various aspects of cognitive science. The theoretical principles include: a construction grammar approach to representing linguistic knowledge; an incremental, single-path processing algorithm with local repair; a cognitive architecture as the computational platform; and embodiment in the robotic agent for grounding language to the agent's perception, action capabilities, and knowledge of the world. This prototype worked well in small experiments, and has been the basis of work in the remaining phases.

We use the Embodied Construction Grammar (ECG) theory developed at UC Berkeley (Bergen & Chang, 2013; Feldman, 2006) to represent Lucia's knowledge of language. This theory fits our needs well as it draws on several decades of work in cognitive linguistics and ways of representing embodied meaning, as well as the general theory of construction grammar. To build this prototype we used the formal ECG grammar specification language (Bryant, 2008) to write a translator program that translates an ECG grammar into Soar production rules. Then we wrote a grammar in the high-level ECG language to handle a test set of 50 English sentences used to cover part of the range of language Rosie needs to learn its tasks.

The rules automatically generated from the grammar are augmented by hand-written Soar rules that comprise the basic infrastructure of the language comprehension algorithm. Together these two sets of rules make up a set of procedures for doing grounded, incremental semantic parsing and interpretation of all the sentences in the test set. This set of procedures successfully processes all 50 sentences to produce an actionable output message for each that matches a pre-defined gold standard. Initially, there were 487 rules generated automatically from the grammar, and 292 written by hand, or 62.5% to 37.5%.

As another experiment performed in this phase we translated the 50 sentences into Spanish and made modifications to the grammar to be able to correctly process the Spanish sentences. The English grammar generated a total of 487 Soar rules. To process the Spanish sentences we were able to share 319 of these rules and add 263 Spanish-specific rules.

Tangible Outputs in This Phase: A paper entitled "Toward Integrating Cognitive Linguistics and Cognitive Language Processing," and an associated presentation, was done as Peter Lindes's successful PhD qualifying exam. A paper entitled "Toward Integrating Cognitive Linguistics and Cognitive Language Processing" was presented at the 14th International Conference on Cognitive Modeling and published in its proceedings (Lindes & Laird, 2016). It won an Honorable Mention at that conference.

Phase 2, April 2016-August 2017 – Integration and Coverage

In this phase, work was done on extending the ECG grammar and Soar processing code to cover several dozen more complex sentences in our test set. These sentences included new vocabulary, sentences to define goals that require an additional level of semantic abstraction, and a number of sentences with conditional expressions or multiple clauses. The full test set of 207 sentences has been processed correctly for some time now, and a benchmark including 200 of these sentences was published (Lindes et al., 2017).

For the 207 sentences, there were 1078 rules generated automatically from the grammar, and 454 written by hand, or 70.4% to 29.6%. Compared to the numbers shown for the first 50 sentences, this shows some increase in the percentage of rules generated automatically from the grammar. A study of the details shows that it may be possible to convert a significant portion of the hand-coded rules into automatically generated ones by extending some features of the ECG language and making the processing algorithm more general. Also, the incremental process by which additions were made to the grammar to process specific new grammatical phenomena can serve as a model for studying possible models of incremental language acquisition.

One large piece of work in this phase was developing techniques for dealing with ambiguity. At each stage of incremental processing there may be more than one option of how to proceed: more than one sense of the current input word, or more than one possible composite construction to apply to the current state of the parse. We have used the principle that at each of these choice points the system chooses the best option available based on all the knowledge it has available at that time, thus allowing the system to proceed on a single analysis path. A variety of heuristics have been developed to deal with various cases of lexical, structural, and semantic ambiguity (Lindes & Laird, 2017a).

When a choice is made in order to proceed on a single path, sometimes later input shows that this was not really the correct choice. In these cases we adopt a strategy called "local repair," adapted from the techniques used by Lewis (1993). Here specific rules detect the problem and correct it by removing part of the comprehension state, an operation Lewis (1998) calls "snip," and then allowing the processing to proceed in a new way that includes the new information. Such a repair is not always possible if the point where the incorrect choice was made is too far back in the parse, thus not being "local" enough. In such cases Lucia shows a "garden-path effect" similar to the one in humans.

Another large piece of work was to integrate the Lucia language comprehension system with the rest of the Rosie ITL agent. A variety of software engineering issues had to be resolved, and now Rosie can use the Lucia parser to comprehend its natural language input.

Tangible Outputs in This Phase: A presentation entitled "Language Comprehension in Rosie" was presented at the 37th Soar Workshop. A paper entitled "Cognitive Modeling Approaches to Language Comprehension Using Construction Grammar" was presented at the AAAI Spring Symposium in April 2017 and published in its proceedings (Lindes & Laird, 2017b). A paper entitled "Ambiguity Resolution in a Cognitive Model of Language Comprehension" was presented at the 15th International Conference on Cognitive Modeling and published in its proceedings (Lindes & Laird, 2017a). A paper entitled "Grounding Language for Interactive Task Learning" was presented as a poster at the First Workshop on Language Grounding for Robots at the annual meeting of the Association for Computational Linguistics, and was published, along with supplementary materials, in its proceedings (Lindes et al., 2017).

Phase 3, August 2017-April 2019 – Architecture and Coverage Exploration

An important part of the research on Lucia is to explore different possible ways of representing and processing knowledge within the Soar cognitive architecture. The basic system, which we call System A, is implemented with all the grammar knowledge being translated into Soar production rules. Work is currently in progress on implementing a System B that will put the grammar data into Soar's semantic memory instead. This makes it possible to use the new spreading activation feature in semantic memory to bias retrievals based on situational context. At this point, several simple test cases have been run and confirmed that spreading activation can indeed bias these retrievals. Code is also working to elaborate the complex data structures needed to complete the comprehension steps. More work is needed to build the tool to translate the grammar into semantic memory structures and complete the full processing algorithm in this environment.

The overall ITL research project includes developing Rosie to be able to learn a variety of tasks involving navigation and manipulation in the real world, such as delivering objects to people in a building (Mininger & Laird, 2018), as well as a large set of games and puzzles (Kirk & Laird, 2016). A collection of interaction scripts for teaching these tasks has been assembled, and we are currently part-way through the process of extending Lucia to process correctly all of these scripts. Most of what needs to be added is a large number of new vocabulary words, and we have software tools partially developed to automate the addition of these words. In addition, some additional complex syntactic constructions are needed, as well as some Soar code to properly handle the semantics of linguistic expressions that refer to things in a hypothetical situation not currently visible in the real world.

Work has also been done on various analyses of how the existing system works. Grammatical productivity measures how well the grammar can recognize sentences it has not seen before. To model the creativity of human language, the grammar must have recursive features in order to produce an unbounded set of sentences from a finite set of grammatical knowledge. The Lucia grammar and processing do implement recursion, however to get a finite measure of productivity we must limit the depth of recursion. An analysis of the productivity of the Lucia grammar shows that, without using any recursion and with the relatively small vocabulary we have so far, the grammar can produce over 400,000 unique, syntactically correct referring expressions and over 1.4 billion declarative sentences. An analysis of processing dynamics is partially completed to show the statistics of the number of constructions built per input word and the number of Soar decision cycles used per construction.

An important new insight that has been gained from this work is that if construction grammar is a cognitively plausible model of how to represent human linguistic knowledge, and human processing is done incrementally with a single path as our model assumes and much psycholinguistic research suggests, then there must be a repeating cycle in the processing where each new construction is selected, integrated into the developing comprehension state, and grounded to the agent's knowledge. This insight has several implications.

Together, the work on using semantic memory to store linguistic knowledge and the principle of building constructions one at a time have implications for how a cognitive architecture must work to do language processing. A recent paper by experts in cognitive architectures (Laird, Lebiere, & Rosenbloom, 2017) describes a Standard Model of the Mind, which by community consensus has been renamed the Common Model of Cognition (CMC). We have done an analysis of how the Lucia model relates to the CMC (Lindes, 2018). Some key research questions for future work emerge from this analysis.

Our System B uses spreading activation in semantic memory to help select what construction to apply next. However, several retrievals from semantic memory are needed to fully integrate a construction into the comprehension state. If we leave all the detailed information in semantic memory, we violate the cognitive constraint that this detailed knowledge of language is not available to human consciousness. Furthermore, doing these multiple retrievals is too slow for real-time comprehension. Soar's chunking mechanism should be able to be used to speed up this process, but if everything is chunked the contextual bias from spreading activation is lost. This suggests chunking the integration part but not the selection part of each construction cycle. We have yet to implement and test this idea. There is also the question of how the skilled knowledge produced by Soar's chunking can be modified as new knowledge is built during language acquisition.

The model of comprehension using construction cycles in a general cognitive architecture makes predictions about how each cycle will reference a variety of different memories in a certain predictable sequence. We have made a small initial effort to see how this theory might correlate with some of the large body of research on measuring brain responses during human language comprehension (eg. Brennan & Hale, 2019).

A preliminary analysis of this comparison shows two significant changes that would have to be made to our model to better simulate brain responses (Lindes, 2019). First, activity in the brain tied to the processing of a particular word takes much longer than the average time between words in normal speech. This implies that the brain does some sort of parallel processing of words, in something like a cascade or pipeline. Implementing this in our model would be a significant change, and perhaps difficult to do in the current Soar architecture. Second, much research on brain responses to language emphasizes that the brain does a lot of prediction of what it expects to come next (eg. Bornkessel-Schlesewsky & Schlesewsky, 2019). Our model for System A has a very minimal implementation of prediction. With System B we can envision using partial matches of composite constructions as one way to do prediction and spreading activation to bias future lexical retrievals, but none of this has been implemented as yet. These suggestions point to important improvements that could be made to our model, and these will be a challenge for applying existing cognitive architectures. Much more remains to be explored along these lines.

Expected Progress in Ongoing Work: Four main aspects of this problem are expected to be completed in the next year. We will extend the grammar and processing algorithm as needed to cover the full set of approximately 500 sentences used for teaching Rosie the various tasks being worked on by the group. The System B using semantic memory will be completed. The analyses of both the grammar and the processing dynamics will be completed, along with some further exploration of how the processing data corresponds to empirical data on human sentence processing. Finally, the results of this work will be documented in two or three more published papers and Peter Lindes's doctoral dissertation.

Tangible Outputs in This Phase: A presentation entitled "Cognitive Language Comprehension in Rosie" was presented at the 38th Soar Workshop. A paper entitled "The Common Model of Cognition and humanlike language comprehension" was presented at the AAAI Fall Symposium last October and published in its proceedings (Lindes, 2018). A presentation entitled "Language Comprehension and the Frontiers of AI" was presented in a Computer Science Colloquium at Brigham Young University in October 2017. A presentation entitled "Lucia: A Cognitive Model of Human Language Comprehension" was presented in a Computer Science Colloquium at Brigham Young University in February 2019. An abstract entitled "Predictions of a Model of Language Comprehension Compared to Brain Data" was accepted in April this year for presentation as a poster at the 17th International Conference on Cognitive Modeling and publishing in its proceedings (Lindes, 2019).

References

Bergen, B. K., & Chang, N. C. (2013). Embodied Construction Grammar. In T. Hoffman & G. Trousdale (Eds.), *The Oxford Handbook of Construction Grammar* (pp. 168–190). Oxford University Press.

Bornkessel-Schlesewsky, I., & Schlesewsky, M. (2019). Toward a Neurobiologically Plausible Model of Language-Related, Negative Event-Related Potentials. *Frontiers in Psychology*, *10*(FEB). https://doi.org/10.3389/fpsyg.2019.00298

Brennan, J. R., & Hale, J. T. (2019). Hierarchical structure guides rapid linguistic predictions during naturalistic listening. *PLoS ONE*, *14*(1). https://doi.org/10.1371/journal.pone.0207741

Bryant, J. E. (2008). *Best-fit Constructional Analysis*. PhD dissertation in Computer Science, University of California at Berkeley.

Feldman, J. A. (2006). *From Molecule to Metaphor: A Neural Theory of Language*. Cambridge, MA: The MIT Press.

Kirk, J. R., & Laird, J. E. (2016). Learning General and Efficient Representations of Novel Games Through Interactive Instruction. *Advances in Cognitive Systems*, *4*.

Laird, J. E., Anderson, J., Forbus, K. D., Lebiere, C., Salvucci, D., Scheutz, M., … Kirk, J. R. (2017). Interactive Task Learning. *IEEE Intelligent Systems*, *32*(4), 6-21 (invited). https://doi.org/10.1109/MIS.2017.3121552

Laird, J. E., Lebiere, C., & Rosenbloom, P. S. (2017). A Standard Model of the Mind: Toward a Common Computational Framework across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics. *AI Magazine*, *38*(4), 13. https://doi.org/10.1609/aimag.v38i4.2744

Lewis, Richard Lawrence. (1993). *An Architecturally-based Theory of Human Sentence Comprehension*. PhD dissertation in Computer Science, Carnegie Mellon University.

Lewis, Richard L. (1998). Reanalysis and Limited Repair Parsing: Leaping off the Garden Path. In J. D. Fodor & F. Ferreira (Eds.), *Reanalysis in Sentence Processing. Studies in Theoretical Psycholinguistics, vol 21.* (pp. 247–284). https://doi.org/10.1007/978-94-015-9070-9_8

Lindes, P. (2018). The Common Model of Cognition and humanlike language comprehension. *Procedia Computer Science*, *145*, 765–772. https://doi.org/10.1016/j.procs.2018.11.032

Lindes, P. (2019). Predictions of a Model of Language Comprehension Compared to Brain Data. *Proceedings of the 17th International Conference on Cognitive Modeling (ICCM 2019)*.

Lindes, P., & Laird, J. E. (2016). Toward Integrating Cognitive Linguistics and Cognitive Language Processing. *Proceedings of the 14th International Conference on Cognitive Modeling (ICCM 2016)*. Retrieved from http://acs.ist.psu.edu/iccm2016/proceedings/lindes2016iccm.pdf

Lindes, P., & Laird, J. E. (2017a). Ambiguity Resolution in a Cognitive Model of Language

Comprehension. *Proceedings of the 15th International Conference on Cognitive Modeling (ICCM 2017)*.
Lindes, P., & Laird, J. E. (2017b). Cognitive Modeling Approaches to Language Comprehension Using Construction Grammar. *AAAI Spring Symposium on Computational Construction Grammar and Natural Language Understanding - Technical Report*, *SS-17-02*. Retrieved from https://www.aaai.org/ocs/index.php/SSS/SSS17/paper/view/15285
Lindes, P., Mininger, A., Kirk, J. R., & Laird, J. E. (2017). Grounding Language for Interactive Task Learning. *Proceedings of the First Workshop on Language Grounding for Robotics*, 1–9. https://doi.org/10.18653/v1/w17-2801
Mininger, A., & Laird, J. E. (2018). Interactively Learning a Blend of Goal-Based and Procedural Tasks. *Thirty-Second AAAI Conference on Artificial Intelligence*. Retrieved from https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/viewPaper/16903

II. Task Representation Learning

Research conducted by James Kirk under supervision from John Laird

Task Representation

Over the past years, we have worked to clarify the problem of learning through instruction the complete definition (from scratch) of a new task in a novel environment, focusing on discrete goal-oriented games and puzzles. We have clarified a critical aspect of the problem of learning new goal-oriented tasks from scratch through instruction. Based on Newell's original formulation of the problem space hypothesis, we have identified the core components of the problem space of a task, what we call the *task elements.* These include the actions, goals, failure conditions, and task-specific terms used to define games. We have shown that from a set of innate primitives, the building blocks of learning, our system can learn, through complex hierarchical composition, a large number of new task elements.

We have also created tools to display visualizations of the internally created world states and graphical representations of the internal structures learned for the task elements, in real time while the agent is learning (online). These visualizations allow an expert (one who is familiar with the learned representations) teacher to actively see the representations created for different interpretations and how they map (or fail to map) to the current world state.

Learning new Concepts

Over the years, we have dramatically increased the number and complexity of new terms, or concepts, that can be learned. In our initial work **[7]** Rosie could only learn a handle of new terms that were defined simply "If a location is not below a block then the location is clear."

Concepts can now be defined by complex compositions of previous learned concepts, and a learned concept can have multiple task or domain-specific definitions. **[4]** For example, the agent can learn a general concept that an opponent's piece is black, which is used in many different games. The agent can then learn that in one game instance, red blocks are used for the opponent's pieces. Moreover, it can learn that in a different game, red blocks can be used for its pieces. Thus, the mapping, or grounding, of terms to agent concepts can be many-to-many and context dependent. Furthermore, concepts can be defined over sets of objects and be constrained by functions (such as 'the number of red blocks'). To support these additions, the language system

had to be extended to handle multiple types of anaphoric reference, complex sentence constructions, conjunction, and quantifiers.
This hierarchal composition enabled the teaching of new concepts such as “attacking,” using a newly defined action to define a term. We have also expanded the concepts Rosie can learn by introducing new primitive concepts. Rosie is now precoded with primitive operators for subset of, product of, whether an object has a specific attribute or the same attribute as another object, and a new primitive action, *mark*, that creates a new value of an object by “writing” a string or value on it. This additional knowledge expanded the types of games Rosie could learn to include marking games such as Sudoku, KenKen, and Logi-5, and new variants of Tic-Tac-Toe that use marking ‘X’ instead of moving blocks.

Learning Heuristics
We also extended Rosie to teach it heuristics ("Prefer moving a block onto a corner location over moving the block onto an edge location." - Othello); search bounds ("The solution has six steps."); and opponent actions (“Your opponent can move a blue block onto a clear location.”) Though this has not been a focus of our research, it is extremely useful in reducing the search spaces of puzzles so that we can demonstrate that the agent can solve a task once it has learned it.

Ambiguity and Knowledge Transfer
Over the past year, we have worked on Rosie to extend its ability to effectively learn and transfer knowledge in more difficult learning scenarios, where ambiguity (due to many possible meanings) and learning distractors are present. Previous task learning systems have assumed a fixed set of words, or that new words and task elements can be directly mapped (one-to-one) to known primitives or subsymbolic representations in a single domain. However, as an agent learns many tasks in many different settings, there will inevitably be many-to-many mappings between words and meanings (the components of a task). In some cases, knowledge learned in previous tasks can interfere with a new task.

Our recent efforts **[1]** enable Rosie to create, analyze, and debug many different interpretations of task elements in order to handle scenarios where ambiguity and knowledge interference can negatively impact the ability to accurately learn and transfer knowledge.

In our approach, the agent learns to recognize the task elements by asking for if-then language definitions of the actions, goals, and failure conditions, creating internal declarative recognition structures from these definitions, and recursively learning all the supporting terms needed to ground(map) the structures to the world state. To ensure that Rosie correctly interprets an ambiguous situation, Rosie generates all possible declarative recognition structures, for each known meaning of the defining terms used. In order to determine the correct interpretation from the set of generated structures, Rosie leverages the external world state example. If the agent finds only one of the possible recognition structure can be detected, it learns this interpretation, otherwise if the agent finds that multiple structures from different interpretations can be satisfied (detected) in the current state, the agent analyzes each of the potentially matching interpretations and applies simple disambiguating strategies, either automatically, if possible, or through communication with the teacher.

This new ability of Rosie to create, analyze, and debug *multiple* hierarchical symbolic recognition structures extends Rosie so that it can learn to correctly recognize and apply the element of a tasks when many interpretations are possible. Furthermore, it enables Rosie to efficiently communicate to resolve ambiguity and select correct interpretations. Rosie can now resolve cases of knowledge interference, such as when two known meanings of the term “clear” result in two interpretations of an action (“move the block onto the clear location”) that both result in the detection of actions in the current state. By analyzing the differences between those interpretations, Rosie can create and communicate simple disambiguating questions such as “how many actions are present 1 or 2?” or “how many clear objects are present 1 or 5?” in order to determine the correct interpretation.

Expansion of learnable games

A major goal of ITL is to support task learning that is general: the agent is not limited to a small set of tasks that it can learn. We have attempted over the years to teach Rosie an increasing variety of different games in different settings pushing the total number of games that Rosie is capable of successfully learning. With the advancements in learning new concepts and using multiple hierarchical structures to handle ambiguity in knowledge transfer, we have made Rosie capable of learning many more games and capable of transfer knowledge between tasks even when there is knowledge interference. One of the exciting aspects of the work is we have reached a point where we no longer need to make changes to the system to add new games and puzzles.

Initially we were only able to teach Rosie a few games (Tower of Hanoi, Tic-Tac-Toe), which were then expanded into a small set of 11 games back in 2014. **[7]** This was expanded to in early supported work in 2016 to 17 games (with evaluations) **[4]**, to 40 games the beginning of this year **[1]**, and finally to 55 games that we can now teach Rosie.

These games, which include variants indicated by (total number) or (names), are Tower of Hanoi (3), N-Puzzle (4), Marking puzzles (Sudoku, Killer Sudoku, Jigsawdoku, KenKen, Product KenKen, Logi-5, Shuffle, Survo, Suko, Sujiko, Kakuro), Map 4-Coloring, Chess puzzles (N-Queens, N-Kings, N-Rooks, N-Knights, Knight’s tour, Knight swapping, 4 Corner knight swapping), Peg solitaires (2), Card solitaires (Golf, Pyramid, Tri Peaks), River crossing puzzles (Fox, Goose & Bean, Missionaries and Cannibals, Jealous Husbands, Jealous Wives, Family crossing), Traveling Salesman in a grid, 3x3 stone games (Tic-Tac-Toe, Three Men's Morris, Picaria, Nine Holes), Othello, Breakthough, Frogs and Toads, Eight men on a raft, Stacking Frogs (3), Blocks World (2), Mazes (simple, block pushing), Sokoban, Mahjong puzzle, and a sorting puzzle.

We have created a public archive as a resource for researchers in Interactive Task Learning that contains the teaching scripts and state representations for these games, as well as a video of Rosie learning. This is available online at www.umich.edu/~jrkirk/ijcai2019.html

Evaluations of Research

To evaluate transfer in task learning, we created instructions for 40 common games and puzzles and taught the agent sequences of the 40 games in 1000 randomly generated permutations. In each permutation, each game is taught, one after another, using scripts that simulate a teacher.

All 40 games are learned correctly in each permutation. The figure below shows the number of words, on average, used to teach each game in each position in the teaching order. At position 0, no other games have been taught, and at position 39 all other games have been taught. As more games are taught, the number of words required to teach a game decreases if there is transfer of the task elements shared between games. Games that have substantial conceptual overlap, such as Five-Puzzle and Eight-Puzzle, which share actions (slide) and learned predicates (clear, matched, adjacent), can be defined using very few words (only 31) at the end. The gradual decrease in the number of words is a reflection of the gradual increase in the probability that a related game is previously taught.

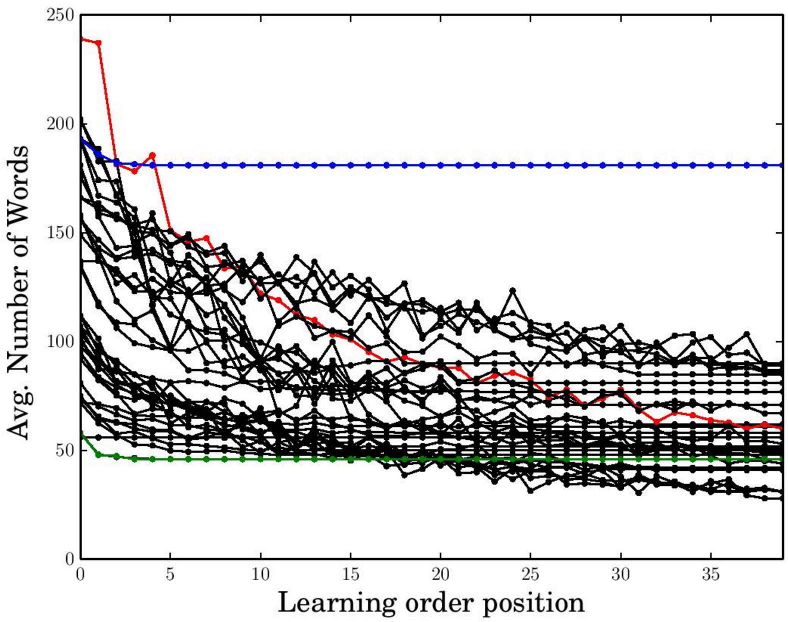

The red line highlighted is for Killer Sudoku, a Sudoku variant that has constraints about the sum of values in specified section (as in KenKen). The number of words required to initially teach (position 0) this puzzle is high due to the number of constraints in the puzzle. However, because of the overlap in concepts with the other tasks (Sudoku, KenKen), it benefits the most from knowledge transfer, with a decrease of more than a factor of three. The Frogs and Toads puzzle (blue) and Blocks World puzzle (green) show the least transfer because they share only *clear* with other tasks.

This experiment was then repeated with small clusters of games to further analyze knowledge transfer. These task clusters contain tasks that have a large conceptual overlap: Tic-Tac-Toe, Three Men's Morris, Nine Holes; Killer Sudoku, KenKen, Sudoku; and N Queens, N Rooks, and N Kings. The final cluster contains tasks with little overlap, with a single task selected from each of the other clusters. The figure above shows the results, again showing the number of words required to teach the task based on the position in the teaching order. Plots A-C show the dramatic effects of transfer in clusters of similar tasks, while Plot D shows almost no transfer between the unrelated tasks. This result is expected, but other task learning approaches that learn directly from subsymbolic representation have failed to replicate this type of task transfer that leads to dramatic learning speed up.

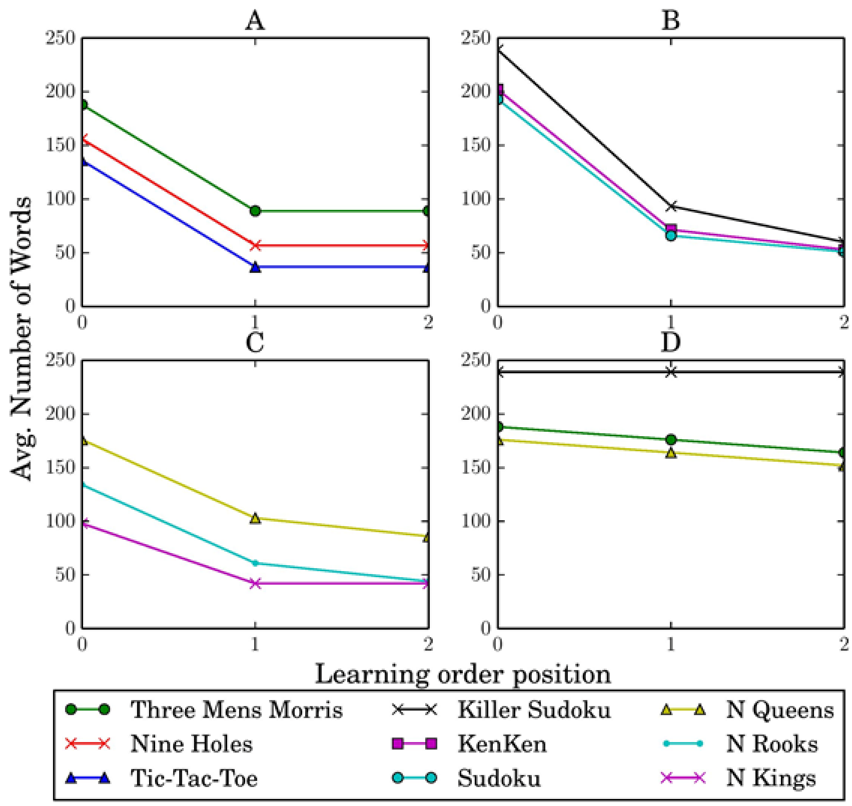

References

[1] Kirk, James R. and Laird, John E. "Learning Hierarchical Symbolic Representations to Support Interactive Task Learning and Knowledge Transfer." International Joint Conference on Artificial Intelligence. *IJCAI*, 2019.

[2] Peter Lindes, Aaron Mininger, James R. Kirk, and John E. Laird (2017). Grounding Language for Interactive Task Learning. Proceedings of the 1st Workshop on Language Grounding for Robotics at ACL.

[3] John E. Laird, Kevin Gluck, John Anderson, Kenneth D. Forbus Odest Chadwicke Jenkins, Christian Lebiere, Dario Salvucci, Matthias Scheutz, Andrea Thomaz, Greg Trafton, Robert E. Wray, Shiwali Mohan, and James R. Kirk (2017). Interactive Task Learning, IEEE Intelligent Systems, 32(4), 6-21, (invited).

[4] Kirk, James R. and Laird, John E. "Learning General and Efficient Representations of Novel Games Through Interactive Instruction." In Proceedings of the Fourth Annual Conference on Advances in Cognitive Systems, Evanston, Illinois, 2016.

[5] Kirk, J., Mininger, A., Laird, J. "Learning task goals interactively with visual demonstrations." Biologically Inspired Cognitive Architectures. New York, New York, 2016.

[6] Mohan, S., Kirk, J., Mininger, A., Laird, J. "Agent Requirements for Effective and Efficient Task-Oriented Dialog." AAAI Fall Symposium Series, North America, 2015.

[7] Kirk, James R. and Laird, John E. "Interactive Task Learning for Simple Games." Advances in Cognitive Systems, vol 3, pp. 13-30, 2014