Grant or Contract Number: N00014-18-1-2337
Date Prepared: July 19, 2021
Project Title: Extending Interactive Task Learning
Final Summary Report: June 16, 2018-June 15, 2021
Principle Investigator:
John E. Laird,
734 647-1761,
laird@umich.edu
University of Michigan

Distribution Statement
DISTRIBUTION A. Approved for public release: distribution unlimited.

ABSTRACT
Our goal is creating artificial agents that can interact with humans and learn completely new tasks through instruction. Solving this problem requires integrating many capabilities across AI. In this project, we identified three dimensions of task complexity (diverse types of actions, task formulations, and task modifiers), and implemented extensions that demonstrate greater learning capabilities for each dimension than previous work. First, we extended the representations and learning mechanism for innate tasks so the agent can learn tasks that utilize many different types of actions beyond physical object manipulation, such as communication and mental operations. Second, we implemented a novel goal-graph representation so that an instructor can formulate a task as achieving a goal and let the agent use planning to execute it, or can formulate the task as executing a procedure, or sequence of steps, when it is not easy to define a goal. Third, we added support for learning subtasks with various modifying clauses, such as temporal constraints, conditions, or looping structures. Our system has been used with various robotic platforms, and it combines all of these extensions while learning complex hierarchical tasks that cover extended periods of time and demonstrate significant flexibility.

1. Overview of Project

1.1. Objective:

The long-term goal of our research is to develop autonomous intelligent agents that have the cognitive capabilities of humans. One challenge is that humans can extend their capabilities through learning and adaptation by interacting with their environment and other humans. Existing robots and AI systems (agents) must be programmed by hand, requiring both subject matter experts and AI engineers. There are no ways of dynamically extend the capabilities of AI systems to new tasks or customize them to the specifics of a situation. As we move to military environments that include both humans and AI systems, we need the ability to quickly expand the capabilities of our AI systems. Our goal is to develop the underlying science and technology to support rapid instruction of AI systems through natural interaction – usually language. Before this project, we had developed a system that learns limited types of tasks interactively. The purpose of this project has been to greatly expand the diversity and complexity of tasks and knowledge that can be learned via Interactive Task Learning (ITL). Specifically, tasks that involve many different types of actions (e.g., communicative, perceptual, mental), modifiers (e.g., temporal, conditional, repetitive clauses), and formulations (e.g., goal-achievement, procedure, optimization).

1.2. Background:

In the five years leading up to this project, we researched how autonomous agents can learn new tasks through natural language interactions with humans, a capability that we call Interactive Task Learning (Laird et al. 2017). The goal of ITL is to provide a means for learning all aspects of a task through natural interactions with humans without programming. Our approach to interactive task learning emphasizes mixed-initiative instruction through natural language. When an agent is in a situation where it cannot make progress on a task, it will ask for help, and a human instructor provides whatever type of knowledge the agent needs through natural language. The language is typically grounded in the current situation, so that both the human and the agent can refer to relevant objects and relations in the environment. In some cases, the human can use instructions that refer to hypothetical situations.

Our work on interactive task learning led to the development of Rosie (Kirk & Laird, 2014; Mininger & Laird, 2016), a robotic agent implemented using the Soar cognitive architecture (Laird, 2012). Rosie learns many different tasks, from scratch, via natural language and demonstration. It has been taught over 60 different games and puzzles (Kirk & Laird, 2019) as well as navigation and delivery tasks, such as fetching and delivering objects. It accumulates knowledge over time, so that it can apply what it previously learned to new tasks, greatly reducing the amount of knowledge that must be provided by a human instructor. It learns hierarchical concepts that are defined both from primitives and previously learned concepts, and from the tasks it learns that are defined in terms of previously learned tasks. Under AFOSR funding, we expanded Rosie's capabilities for language understanding (Lindes & Laird 2017; Lindes, Mininger, Kirk, & Laird 2017) and hierarchical concept learning (Kirk & Laird, 2019). The language capabilities increase the range of instructions a human can use to teaching new tasks, greatly increasing the ease and robustness of instruction. Hierarchical concept learning allows for multiple groundings of the same concepts for different situations, as well as reuse of abstract concepts in new situation, greatly reducing what needs to be taught across multiple tasks.

Although we had previously made significant progress, there were still significant limits to Rosie's generality and capabilities. Before this research, we focused on tasks where there is a well-defined goal state (such as winning a game of Othello, or delivering a package), and only a subset of the available types of knowledge and reasoning. Furthermore, we constrained Rosie to simple pick and place actions, combined with simple indoor navigation.

In this project, we significantly extended the generality and adaptability of autonomous interactive task learning.

1. The first one was increasing the diversity of the innate tasks to many different types of actions (e.g., communicative, perceptual, mental). These are the building blocks from which more complex tasks are composed, so expanding this set of innate tasks will lead to learning tasks that involve broader types of actions.
2. The second area was adding more ways to formulate a task. We had demonstrated learning tasks formulated as achieving a goal or following a procedure, but we identified other ways to formulate a task that were worth adding.
3. The third type of task diversity was having different clauses that further modify a given task command, such as specifying a time, adding a conditional, or a repetition. Increasing this type of diversity will lead to greater customization when teaching a specific task.

By working on these three areas, we significantly extended the task learning capabilities of our agent. The point is to achieve performance on these tasks, through *instruction*, that exceeds previous efforts in terms of task complexity and underlying knowledge that have been developed through manual programming, and orders of magnitude faster than possible with manual programming.

2. Year 1: Activities and Accomplishments

Over the first year, we focused on the problem of learning and performing tasks with different kinds of modifiers. A basic version of a task (e.g. 'Move the box onto the table') can be modified with clauses that specify when, where, and how the task should be performed. If the agent can handle tasks with these modifiers, it can learn tasks that are more complex and varied. A key requirement is that it must be able to learn and reason over these types of modifiers, not simply execute them.

First, we have added support for temporal clauses that can specify the desired start, end, or duration of a task. Second, we have added support for conditional modifiers of both goals and actions. Conditional goals allow greater task variation by allowing the instructor to specify different goals for a task depending on the type of object. For example, the agent learns the task of storing an item where the storage location depends on the object ('If the fork is a utensil, then the goal is that the fork is in the drawer' vs 'If the soda is a drink then the goal is that the soda is in the fridge and the fridge is closed'). The agent can also learn conditional actions where the agent should only do the action if the condition is met. If the instructor uses the keyword 'whenever,' it indicates a global conditional action that should be performed any time the condition is met.

Below is an example patrol task that we can teach the agent to perform and that includes temporal and conditional modifiers.

Patrol Task

- Patrol the building at 15:00.
- Whenever you see trash then discard it.
- Go to the kitchen.
- If the fridge is open then close the fridge.
- Go to the copy room.
- If a person is visible then ask 'Can I do anything for you?'
  (If they give a command, do that command)
- Go to the main office.
- Say 'If you have a package, put it on the table'.
- Wait three minutes.
- If a package is on the table then deliver it to the mailroom.
- You are done

We also extended the task-learning agent to two new domains.

- The first was a simulated office environment in the AI2Thor simulator. This domain allowed us to learn much more complex tasks with interesting object interactions. For example, it includes a kitchen with a working refrigerator and microwave that change object temperature, drawers and cabinets that can hide objects, and object affordances such as cutting an apple or filling a cup.
- The second domain is a small commercial forklift robot called Cozmo that can drive around on a table and interact with special cubes. The types of tasks it can do is more limited, but it demonstrates our approach can be used with real-world robots in an environment that is easy to test and experiment with.

Our overall approach of using a cognitive architecture as the basis for developing an instructable agent proved to be effective for defining new tasks, both from scratch and by building on previously instructed tasks. Although the language is somewhat restricted, it is flexible enough to specify a variety of tasks, including goal-based tasks, where the instructor describes a goal, and the agent figures out a solution through planning; and procedural tasks, where the instructor gives step by step instructions. We also showed that these types of tasks can be intermixed and can transfer to new tasks as subtasks.

A significant theoretical contribution over this year was identifying and describing the different dimensions of task complexity described above. The ultimate goal of this effort is to develop a taxonomy of task learning capabilities that can be used to analyze a specific approach to ITL and serve as a basis of comparison between different approaches.

3. Year 2: Activities and Accomplishments

During year 2, we added significant capabilities that extended the task learning capabilities across all three dimensions of task diversity. First, we added more innate tasks, most notably mental actions that involve long-term memory. This means the agent learns tasks that involve storing and retrieving knowledge across long time periods. For example, it learns to recall a person's drink preference when serving a drink or recall where to store a particular object. We also improved the ability of the agent to find objects or people by adding these mental actions, such as recalling where an object was last seen or recalling a person's office to go look there. Having knowledge-based actions greatly increases the efficiency of the task learning as the instructor does not have to individually teach each case. As a theoretical contribution, while adding mental actions we developed a novel approach to learn action models by utilizing the agent's past experiences stored in its episodic memory. This makes the learned tasks more robust to novel variations without the need for further training.

The second dimension of task diversity that we tackled was having diverse task formulations. Over the past year, we added two additional formulations: opportunistic tasks and composite tasks. An opportunistic task is one that should be initiated whenever some condition in the environment is met. For example, we teach the agent to throw away trash when litter is seen on the ground or pull the fire alarm if a fire is detected. A composite task is one where the task is to carry out a set of subtasks, but the ordering is not important. For example, we teach the agent to tidy the kitchen by closing any cabinets, clearing plates from the table, and storing any condiments. Unlike a procedure, here the ordering is not important. We had also identified RL tasks as another formulation but decided to first add these simpler ones before adding RL, as that will involve major extensions.

We have also made progress on the third dimension of task diversity, additional task modifiers, by adding support for repetitious clauses (while, do-until, and do N times). We demonstrated the agent learning an interior guard task where it patrols a sequence of rooms in a loop until relieved. One novel capability that was added was the ability to interrupt the agent's current task to insert additional steps or instructions. This is the first example of allowing the instructor to modify a previously learned task; a crucial capability.

We also made progress in teaching causal relationships. This required extending the agent's internal knowledge representation so that it had declarative representations of the causal relationships. This is necessary for use in means-ends analysis. We have a prototype implemented where the AI agent learns simple causal relationships through interactive task learning, both accurately and efficiently.

We also expanded the environments and task scenarios that the agent can learn in. We had planned to work with the AI2Thor simulator, but it proved too cumbersome to extend and add new features, so we decided to improve an in-house robot simulator instead. Of particular interest was having a simulated environment where we could teach a sentry robot some indoor military procedures, such as interior guard and patrolling a building. The target tasks would be developed through collaboration with contacts at ONR and NRL. In this simulator, we created a barracks environment to learn an interior guard task. Using the Marine Corps' handbook on interior guard as a reference, the agent learned to patrol a series of rooms until relieved by an officer, cleaning and inspecting each room in different ways. This task is the most complex one demonstrated so far and involves 36 instructions over 15 minutes of continuous learning. In the same simulator, we developed an office environment to test tasks involving an office assistant robot domain. Finally, we successfully have our agent work with the Cozmo toy robot, but difficulties in localization limited the effectiveness of this platform.

4. Year 3: Activities and Accomplishments

During year three, we made significant improvements and extensions to the agent's task learning capabilities along the three dimensions of task diversity we had identified. First, although we did not add any new types of innate tasks, we did significantly improve their generality and flexibility by unifying the representations for learned and innate tasks. This means that all tasks have the same representations in procedural and declarative memory, regardless of whether they were innate or learned. The major advantage is that innate tasks can utilize the existing planning and learning mechanisms. We then ran an experiment that evaluated the learning performance of the agent with and without this capability when performing 100 random variations of a move task. The results showed that unifying the task representations significantly improved the generality of learned rules, improved the planning efficiency, and increased the agent's reactivity.

Another contribution was adding a new type of task formulation: a maintenance goal. This type of task involves maintaining some goal condition over time, for example, ensuring there are always three water bottles on a table or making sure a door stays closed. The main distinction is that the task does not end when the goal condition is achieved. It requires an additional termination condition (e.g., until the meeting is over) or an explicit statement from the instructor. We also developed several demonstrations that showed several different ways the agent could blend aspects of different formulations within a single learned task. This improved the robustness of the agent and supported our argument that a general task learning agent should support multiple types of formulations. One such demonstration involved teaching the agent the task of giving a tour of the lab, which utilized every formulation within a single task hierarchy.

We also made several accomplishments in the area of diverse task modifiers. First, we added several types of temporal clauses specifying either the start or end of a task as a duration, clock time, or condition. Second, although last year we had done some proof-of-concept work in the area of repetitious clauses (while, do-until, and iterative loops), in the third year we reimplemented that functionality in a more rigorous form. The main theoretical contribution was developing the goal graph: a hybrid task representation that combines both procedural and state-based goals and that is flexible enough to include complex control flow such as branches and loops. This is a general representation stored in the agent's long-term semantic memory. To execute a task, the agent follows a path of subgoals from a start to an end node.

Adding all of the different types of modifiers drastically increased the number of ways a learned task can be used as part of a parent task. To demonstrate this, we performed a demonstration where the agent was taught a task through a single training example, and then we tested it by combining that task with nine different modifiers (including temporal, conditional, and repetitious clauses). The agent was able to adapt the learned task knowledge to those nine cases without needing any additional instruction.

5. Personnel

Years 1-3
Principal investigator: John E. Laird

Team Members (graduate students):
Aaron Mininger
Elizabeth Goeddel
Mazin Assanie
Note: Aaron Mininger received his Ph.D. in early 2021. Elizabeth Goeddel and Mazin Assanie will receive their Ph.D.'s in Fall 2021 which will acknowledge this grant for support.

6. Opportunities for training and professional development

- All graduate students had weekly meetings with the PI.
- Weekly research meeting where Ph.D. students read and discuss research papers and present their work to other students.
- We have yearly Soar workshops in which knowledge of our work is transmitted. In June 2018, 2019, 2020, 2021.
- Professor Laird gave a remote presentation on this research to students at the Naval Postgraduate School (April 2019, May 2020).
- We establish the Virtual International Symposia of Cognitive Architecture (VISCA), June 2020 (>350 registrants), 2021 (> 250 registrants). Videos of all the talks are available online. https://visca.engin.umich.edu/.

7. Dissemination of Research

- Soar Workshop, 2018, 2019, 2020, 2021.
- Robotics: Science and Systems, 2018. "Exploiting symbolic task constraint flexibility for improved tabletop manipulation action trajectories in cluttered environments" Mamanatov and Laird
- Aaron Mininger and John E. Laird (2019). Using Domain Knowledge to Correct Anchoring Errors in a Cognitive Architecture Advances in Cognitive Systems.

- Aaron Mininger and John E. Laird (2019). Using Domain Knowledge to Correct Anchoring Errors in a Cognitive Architecture, *Advances in Cognitive Systems*
- Mininger, A. (2021). Expanding Task Diversity in Explanation-Based Interactive Task Learning, *Dissertation.*
- Professor Laird gave the following presentations that referenced this research:
    - "Interactive Task Learning," NIPS Workshop on Learning by Instruction, 2018.
    - "Learning Fast and Slow," Carnegie Mellon University Distinguished Lecture Series, Sept. 2018.
    - "Interactive Task Learning," to Naval Postgraduate School (April 2019)
    - "Cognitive Architecture, Soar, and Interactive Task Learning." Cognitive Science Symposium, University of Michigan, November 2018.
    - "Interactive Task Learning," Autonomous Agents and Multi-agent Systems (AAMAS), May 2019.
    - Cognitive Science Symposium: How Does Current AI Stack Up Against Human Intelligence? Cognitive Science Society Workshop: Everyday Activity
    - Advances in Cognitive Systems Conference: Invited talk and paper presentation. August 2019
    - Speaker at MOR-MS Forum in Taiwan. This is a meeting on AI applications to modeling and simulation, August 2019.
    - Speaker at ONR Sciences of Autonomy meeting and presented summary of this research. August 2019
    - Presented at the Fall HPT&E Technical Review.
    - Speaker at the University of Michigan AI Seminar: Interactive Task Learning, Nov. 2019.
    - Speaker at Ann Arbor Friday AI Seminar: Interactive Task Learning, Nov. 2019.
    - Plenary Speaker at the International Association for Dentistry Research. May 2020.
    - Invited Speaker, "Interactive Task Learning," to Naval Postgraduate School (May 18, 2020)
    - Introductory Speaker, "Introduction to Cognitive Architecture", Virtual International Symposia of Cognitive Architecture (VISCA) (June 2020).
    - The symposium on Democratizing AI, Oct. 6, 2020
    - The conference on Biologically Inspired Cognitive Architectures, Oct. 9, 2020
    - The workshop on Integrated Execution and Goal Reasoning at ICAPS, Oct. 21, 2020
    - OSD Study: Future Directions of Human Machine Teaming, Nov. 15, 2020;
    - AAAI Doctoral Consortium, February 2-3, 2021
    - Human-Robot Interaction Pioneers Workshop, March 8, 2021
    - AAAI Doctoral Consortium, February 2-3, 2021
- The Soar cognitive architecture, which includes software developed as part of this research project, is freely available as open source and is widely distributed with documentation and tutorials. ~300 people have signed up to be on our distribution list.
- We have a web site with videos of our research https://soargroup.github.io/rosie/

- Rosie is freely available with a BSD license.

8. Technology Transfer

All of our research is incorporated into the Soar architecture, which is freely available, open source. The continual extension of Soar (as well as its stability, efficiency, and cost-effectiveness) makes it an attractive system for applications of cognitive architecture.
Soar is being used for multiple research projects at:

- the 711th Human Performance Wing of AFRL at Wright-Patterson AFB
- Naval Postgraduate School
- AFIT
- Navy project on Assessing Cognitive Architectures for Autonomous Mission Management on Unmanned Naval Systems (Naval Surface Warfare Center & Penn State Advanced Research Lab)
- Soar Technology, Inc. for many DoD projects
- more than 20 addition institutions worldwide
- Talk on task learning and our research to Naval Postgraduate School (April 2019, May 2020)

We have a yearly workshop and tutorial in which knowledge of our work is transmitted.

9. Products, Publications, Patents, License Agreements, etc.

Soar Cognitive Architecture Website
Title: Soar
URL: Soar.eecs.umich.edu
Description: The homepage for the Soar Cognitive Architecture. Includes all software, documentation, tutorials, and publications of the Soar research group. Soar is a general cognitive architecture for developing systems that exhibit intelligent behavior. Researchers all over the world, both from the fields of artificial intelligence and cognitive science, are using Soar for a variety of tasks. It has been in use since 1983, evolving through many different versions to where it is now Soar, Version 9.6

10. Awards received

- John E. Laird received the Herbert A. Simon award for Advances in Cognitive Systems, 2019.
- John Laird: University of Michigan, College of Engineering, 2019-2020 Stephen S. Attwood Award, the most prestigious award that the College of Engineering bestows, in recognition of his "extraordinary achievement in teaching, research, service, and other activities that have brought distinction to the College and University."

11. Point of Contact in Navy

Greg Trafton, Head
Intelligent Systems Section
Navy Center for Applied Research in Artificial Intelligence
Code 5515

Naval Research Labs
Washington DC 20375
Last interaction was July 2019.

12. Acknowledgment/Disclaimer

This work was sponsored by the Office of Naval Research (ONR), under grant (or contract) number N00014-18-1-2337. The views and conclusions contained herein are those of the authors only and should not be interpreted as representing those of ONR, the U.S. Navy or the U.S. Government.

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED START DATE	END DATE
07/22/2021	Final Report	05/01/2018	04/30/2021

4. TITLE AND SUBTITLE
Extending Interactive Task Learning

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
N00014-18-1-2337		
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Laird, John E.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
University of Michigan 2260 Hayward St. Ann Arbor, MI 48109-2121	

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)
Office of Naval Research 875 N. Randolph Street, Suite 1425 Arlington, VA 2203-1995	ONR	

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION A. Approved for public release: distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
Our goal is creating artificial agents that can interact with humans and learn completely new tasks through instruction. Solving this problem requires integrating many capabilities across AI. In this project, we identified three dimensions of task complexity (diverse types of actions, task formulations, and task modifiers), and implemented extensions that demonstrate greater learning capabilities for each dimension than previous work. First, we extended the representations and learning mechanism for innate tasks so the agent can learn tasks that utilize many di

15. SUBJECT TERMS
Artificial Intelligence, Interactive Task Learning

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	C. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
U	U	U	UU	11

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
Laird, John E.	(734) 647-1761

PREVIOUS EDITION IS OBSOLETE.

STANDARD FORM 298 (REV. 5/2020)
Prescribed by ANSI Std. Z39.18