

Final Report

**A Neurocognitive Approach to
Robotic Cause-Effect Reasoning During Learning**

N00014-19-1-2044

May 23, 2022

Rodolphe Gentili, Garrett Katz, James Reggia (PI)
University of Maryland and Syracuse University

Form SF 298 appears as the final page of this document.

DISTRIBUTION: Approved for public release: distribution unlimited.

Major goals and objectives of the project

The primary goal of this research was to create and critically evaluate a purely neuro-computational robotic system that uses causal inferences to learn to perform inspection and maintenance tasks from human demonstrations. Our emphasis is on producing a system that can learn to perform and generalize a task from a single demonstration, much as people do. Our target system is a neural network implementation similar to CERIL, a causal reasoning system that constructs plausible explanatory hypotheses for a demonstrator's actions during robot learning and which we previously implemented using traditional symbolic AI software. To achieve this challenging goal, our three research objectives are:

1. Create a neural virtual machine (NVM), a purely neurocomputational platform for implementing cognitive-level algorithms.
2. Use and enhance the NVM to implement a goal-directed neural system for robotic imitation learning based on cause-effect knowledge and inference.
3. Compare the performance of human subjects as they learn the same procedures as our robotic system to obtain a deeper understanding of imitation learning.

Results accomplished towards achieving these goals

Our results on each of the three objectives above over the period from January 2019 to December 2021 when the grant ended include the following.

Objective 1. Our first achievement was to implement and test the NVM, verifying its functionality and demonstrating how different neural representations of symbolic components influence its performance. We generated both theoretical results about the correctness and computational efficiency of the NVM, and experimental results establishing its effectiveness. Once the basic NVM was completed and tested on a battery of tasks, we then studied how to extend it via reinforcement learning of procedures from examples, i.e., using program induction (the resulting system is designated the NVM-RL). This study first showed that the NVM-RL could learn generic algorithmic procedures (such as reversing lists) entirely from scratch (no preprogrammed human knowledge), using only sparse reward signals and vanilla policy gradient reinforcement learning. We also extended the NVM-RL's learning capability to more sophisticated algorithms (such as list sorting) using more sophisticated RL methods (namely, proximal policy optimization). Finally, we demonstrated this NVM-RL approach on a robotic manipulation task (in simulation). In this application, the NVM-RL was first initialized with human-authored procedural knowledge about how to rearrange stacks of blocks with a robotic gripper, as a form of inductive bias. Importantly, this human-initialized NVM-RL was sub-optimal, such that the robot sometimes failed to complete the task successfully. By training this NVM-RL with additional self-practice and reinforcement (no additional human input), it was able to refine its neural representation of the procedural knowledge and improve its success rate on the task.

Objective 2. Because implementing a version of CERIL, our existing imitation learning system that we developed previously with ONR support, in a purely neurocomputational framework is so challenging, our approach to implementing the corresponding neural imitation learning system

was to extend the methods of the NVM in three steps: **a)** enhance our previous methods for neural simulation of working memory to now represent compositional data structures (trees, lists, networks of causal relations, etc.) and operations on these data structures; **b)** build on (a) to implement a high-level programming language called NeuroLISP (Common Lisp implemented in a purely neurocomputational framework) that is sufficiently expressive to allow one to represent high-level concepts relevant to imitation learning (cause-effect relations, conditional branching, variables, pointers, iteration, recursion, etc.); and then **c)** extend and use NeuroLISP to implement imitation learning algorithms modeled after CERIL. This incremental approach proved to be highly successful. The resulting NeuroCERIL reproduced CERIL's performance on a battery of test cases (the same demonstrations used to test the original CERIL system), and runtime memory usage was found to scale linearly with demonstration size. We completed and published the work on (a) and for (b), and we implemented the causal reasoning mechanisms of CERIL as NeuroCERIL and established via testing that it works correctly (a paper describing these latter results has been submitted and is under review).

Objective 3. We completed experiments with human subjects learning to perform a disk dock drive maintenance task to understand the errors people make and for comparison with robotic performance on the high-level planning underlying sequential tasks in general. The findings revealed that throughout practice, human performance (indexed by the Levenshtein distance LD) improved along with a reduction of mental workload which collectively led to an elevation cognitive-motor efficiency. Further, during practice human performance changed similarly for “chunkable” relative to “unchunkable” action sequences, while a much slower attenuation of mental workload and improvement of cognitive-motor efficiency was observed for the latter. Additional analysis examined potential synthetic mental workload correlates generated by CERIL, and found that their changes were reminiscent of those in humans. This complements our prior work which revealed that Levenshtein distance can be used to assess action sequence performance both in humans and humanoid robots. In parallel, during this period two experimental platforms were developed. The first was designed to examine the interaction between high and low level planning processes during a pipe maintenance task. Preliminary findings showed that during practice of “unchunkable” compared to “chunkable” action sequences, human performance and mental workload tended to change slower. In addition, although an elevation of the accuracy demand led to a further accuracy errors, it did not impact the improvement of action sequence performance. The second platform was developed due to the COVID-19 pandemic to conduct virtual (remote) data collection using various tasks (Tower of Hanoi, disk dock drive task, pipe task). This virtual platform was designed in collaboration with Dr. Jim Purlito from the Computer Science Department at UM. A thorough analysis of the usability and cognitive-motor performance when individuals executed various tasks using this virtual platform was conducted and compared to in-person performance. The results revealed that while the current platform was appropriate or needed some relatively straightforward enhancement for tasks using a simple control interface, tasks that needed more complex controls (e.g., pipe maintenance task) required further revisions, including potentially the use of immersive technology. Work in progress is completing data collection for the in-person pipe maintenance task and further examining synthetic mental workload.

During this period, we also acquired and assembled robotic hardware for the Poppy Ergo Jr arms with actuated grippers at Syracuse University. We completed implementing PyBullet simulation environments for our Poppy robot hardware: one environment for the Poppy humanoid (grippers not actuated), and another for the Poppy Ergo Jr arms (actuated grippers, but not attached to the humanoid torso). We completed initial motor planning algorithms that enable (in simulation) the humanoid to walk and the manipulator to pick up objects. Finally, we created a merged simulation environment for the full humanoid with actuated grippers, and we completed a motor controller that allows the full humanoid to keep balance while reaching for and manipulating objects (still in simulation).

Finally, at the beginning of this project we initiated a collaborative research effort with Tomo Furukawa's robotics research group at UVA with the goal of demonstrating a robotic system for inspection and maintenance tasks. The agreement was that we would port our CERIL imitation learning system to work with a simulated version of UVA's Tiger robot, and then transfer the software to Tomo's group, whereupon his group would develop a demonstration shipboard robotic system based on CERIL. We transferred this software to the UVA group for use with their physical robot at the end of the first year of this project.

Opportunities for training and professional development provided by the project

Several graduate students were supported on this project across our two campuses during the last three years. These students learned a great deal about robotics, human experimentation, neural network learning methods, and related issues. They are listed under Participants below.

Dissemination of results

The primary means of disseminating our research results has been via refereed publications in technical journals, conference proceedings, and other venues, as well as presentations at relevant AI, kinesiology, and engineering conferences. We also made our software available online to others by posting it on Github. Our publications, all of which cite ONR support, are listed below towards the end of this report.

Plans for the next reporting period

This is the final report from this project which ended in December 2021. However, we plan to send a white paper to ONR to extend CERIL to much more general tasks in the context of human-robot teaming. The idea is to interleave causal interpretation, planning and acting by a robot as it collaborates with a person on carrying out a task (e.g., device assembly or maintenance).

Technology Transfer

Our research team at Maryland and Syracuse transferred our CERIL imitation learning system to Professor Tomo Furukawa's robotics research team at the University of Virginia about two years ago. This version of CERIL was customized to work with a simulated version of UVA's Tiger robot. We also placed on GitHub (<https://www.github.com/garrettkatz/ceril>) an open-source version of CERIL.

Participants

University of Maryland

James Reggia, Prof., PI: neurocognitive systems, causal inference, machine learning
Rodolphe Gentili, Asst. Prof., Co-PI: computational neuroscience, cognitive motor control
Joshua Brule, PhD student
Greg Davis, PhD student
Isabelle Shuggi, PhD student
Alexandra Shaver, PhD student

Syracuse University

Garrett Katz, Asst. Prof., Co-PI: neural nets, cognitive robotics, AI and computer science
D. Patel, graduate student
Akshay, PhD student
Xulin Chen, PhD student
Borui He, PhD student
Rakesh Boyina, graduate student

Publications

Katz G, Davis G, Gentili R, Reggia J. A Programmable Neural Virtual Machine Based on a Fast Store-Erase Learning Rule, *Neural Networks*, 119, 2019, 10-30.

Brule J. Whittemore: An embedded domain specific language for causal programming, in *Beyond Curve Fitting: Causation, Counterfactuals, and Imagination-Based AI*, AAAI Spring Symposium, March 2019, Stanford, CA.

Oh H, Braun A, Reggia J, Gentili R. Fronto-Parietal Mirror Neuron System Modeling: Visuospatial Transformations Support Imitation Learning Independently of Imitator Perspective, *Human Movement Science*, 65, 2019, 121-141.

Hauge T, Katz G, Davis G, Jaquess K, Reinhard M, Costanzo M, Reggia J, Gentili R. A Novel Application of Levenshtein Distance for Assessment of High-Level Motor Planning Underlying Performance During Learning of Complex Action Sequences, *Journal of Motor Learning and Development*, 8, 2020, 67-86.

Reggia J, Katz G, Davis G. Artificial Conscious Intelligence, *Journal of Artificial Intelligence and Consciousness*, 7, 2020, 1-13.

Hauge T, Katz G, Davis G, Wang D, Reggia J, Gentili R. High-level Motor Planning Assessment During Performance of Complex Action Sequences in Humans and a Humanoid Robot, *International Journal of Social Robotics*, August 2020.

Katz G, Gupta K, Reggia J. Reinforcement-Based Program Induction in a Neural Virtual Machine, *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, Glasgow, United Kingdom, July 2020, 1-8.

Davis G, Katz G, Soranzo D, Costanzo M, Allen N, Reinhard M, Gentili R, Costanzo M, Reggia J. A Neurocomputational Model of Posttraumatic Stress Disorder, *Proc. International IEEE EMBS Conference on Neural Engineering (NERS'21)*, May 4-6, 2021, 107-110.

Davis G, Katz G, Gentili R, Reggia J. Compositional Memory in Attractor Neural Networks with One-Step Learning, *Neural Networks*, 138, 2021, 78-97.

Katz G, Akshay, Davis G, Gentili R, Reggia J. Tunable Neural Encoding of a Symbolic Robotic Manipulation Algorithm, *Frontiers in Neurorobotics*, 15, Dec. 14, 2021.

Reggia J, Katz G, Davis G, Gentili R. Avoiding Catastrophic Forgetting with Short-Term Memory, *Proceedings of the 23rd International Conference on Artificial Intelligence*, July 26-29, 2021.

Lin R, Katz G, Reggia J. Effectiveness of Proximal Policy Optimization Methods in Training a Neural Virtual Machine, *Proc. of the 23rd International Conf. on Artificial Intelligence*, July 26-29, 2021.

Davis G, Katz G, Gentili R, Reggia J. NeuroLISP: High-Level Symbolic Programming with Attractor Neural Networks, *Neural Networks*, 146, 2022, 200-219.

Katz G, Tahir N. Towards Automated Discovery of God-Like Folk Algorithms for Rubik's Cube, *Proc. of the 36th AAAI Conference on Artificial Intelligence*, February 22 - March 1, 2022.

Davis G. A Neurocomputational Model of Causal Reasoning and Compositional Working Memory for Imitation Learning, PhD Dissertation, May 2022.

Shaver A, Peri N, Mezebish R, Matthew G, Berson A, Gaskins C, Davis G, Katz G, Samuel I, Reinhard M, Costanzo M, Reggia J, Purtilo J, Gentili R. Assessment of a novel virtual environment for examining human cognitive-motor performance during execution of action sequences, *24th International Conf. on Human-Computer Interaction (HCI 2022)*, June 26 - July 1, 2022, in press.

Akshay, Chen X, He B, Katz G. Towards Human-like Learning Dynamics in a Simulated Humanoid Robot for Improved Human-Machine Teaming, *4th International Conference on Human-Computer Interaction (HCI 2022)*, June 26 - July 1, 2022, in press.

Davis G, Katz G, Gentili R, Reggia J. NeuroCERIL: Robotic Imitation Learning via Hierarchical Cause-Effect Reasoning in Programmable Attractor Neural Networks, 2022, submitted.

SF 298 follows on the next page.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 23/5/2022		2. REPORT TYPE Final		3. DATES COVERED (From - To) 01/01/2019 — 31/12/2021	
4. TITLE AND SUBTITLE A Neurocognitive Approach to Robotic Cause-Effect Reasoning During Learning				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER N00014-19-1-2044	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Reggia, James; Gentili, Rodolphe; Katz, Garrett				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Maryland College Park, MD 20742				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research Arlington, VA 22203				10. SPONSOR/MONITOR'S ACRONYM(S) ONR	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This project developed a purely neurocomputational cognitive architecture for robotic systems that are doing imitation learning (learning from demonstrations). The main results were: 1. implementation of a neural virtual machine; 2. implementation of a neural system that automatically creates explanations for a human demonstrator's intentions/goals by using cause-effect reasoning; and 3. comparative analysis of human versus robot behavioral activities during imitation learning.					
15. SUBJECT TERMS Robotic imitation learning, programmable neural networks, artificial intelligence, machine learning, learning from demonstrations					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 7	19a. NAME OF RESPONSIBLE PERSON ONR
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (include area code)