
How Ontologically Grounded Agents Learn Scripts through Instruction

Christian Arndt
Sergei Nirenburg
Sanjay Oruganti

ARNDT@RPI.EDU
NIRENS@RPI.EDU
ORUGAS2@RPI.EDU

Abstract

Scripts, also called complex events, are events with component subevents. Ontologically grounded agents use scripts to encode procedural knowledge. Scripts which represent mental procedures (agenda management, learning, etc.) are called **mental scripts**. **Object scripts** encode knowledge about procedures which involve interaction with an agent’s physical environment. We present recent advances in the methods that ontologically grounded agents can use to learn object scripts during operation. Specifically, we outline a new mental script which allows agents to learn task-relevant object scripts through natural language dialogue. We also present a demonstration of this mental script in which an agent assisting in a ship’s engine room learns to replace a fuse in real time.

1. Introduction

Enabling ontologically grounded agents to learn during operation will address two long-standing concerns in the field of symbolic AI research. First: allowing agents to acquire their own knowledge moderates concerns about the ‘knowledge bottleneck’—the cost of assembling high-quality, interpreted knowledge by hand. Second: symbolic AI systems have a reputation for being brittle. This reputation is often overblown; previous research has shown that well-equipped symbolic systems are adaptable to many new situations, including text inputs containing unknown words and concepts (McShane, Nirenburg, & English, 2024). Real-time learning methods (including the one presented here) are another tool that symbolic agents can use to engage with new information. Developing these learning methods is an important step in building robots which can meaningfully contribute to human-robot teams.

This research presents a new procedure that ontologically grounded agents can use to learn scripts through dialogue. Scripts, also called complex events, are events with one or more subevents. Scripts were introduced by Schank and Abelson (1977) as a way of encoding information about event sequences in ontologically grounded AI systems. This theory was further refined by Nirenburg and Raskin (2004). Symbolic AI agents use scripts to reason over and execute complex tasks.

In our framework, called the OntoAgent architecture, **mental scripts** are used to represent reasoning tasks like agenda management and learning. **Object scripts** represent physical tasks like changing an engine’s oil. These script categories are not formally different; instead, their

distinction is made to support our theory of agent metacognition. We present a new mental script, LEARN-SCRIPT, which describes a method for agents to learn object scripts through dialogue.

2. The Learning Script

LEARN-SCRIPT depends heavily on other agent systems which have been the subject of decades of prior research. The OntoAgent architecture’s semantic language understander, OntoSem, translates natural-language inputs into unambiguous, language-independent meaning representations (MRs) formulated to support agent reasoning (McShane & Nirenburg, 2021). The agenda management system determines what to do when prompted with a new MR and facilitates the execution of scripts (McShane, Nirenburg, & English, 2024).

Our research presents two new additions to this architecture. First: we lay out the mental script LEARN-SCRIPT, which describes how an agent can extract script-relevant information from an MR. And second: we augment the agenda manager to recognize new script-learning triggers and engage LEARN-SCRIPT when such triggers are found (Arndt, Nirenburg, & Oruganti, 2026).

3. Learning Demonstration

This learning script has been implemented in a demonstration system in which an agent learns how to replace a fuse, which can be viewed [here](#) (RPI LEIA Lab, 2026). Figure 1 is a screenshot from the demonstration video in which a physical robot operating on the OntoAgent architecture executes the newly learned fuse replacement script.



Figure 1. An agent executes a newly learned script for fuse replacement.

4. Conclusions and Future Work

The learning method presented above provides three main benefits to both physical robots and disembodied agents operating on the OntoAgent architecture. First: it enables agents to learn scripts right away. Second: both the script-learning method and the learned scripts are inspectable and enhanceable over time. Third: the method is generalizable across domains.

In the near-term, we plan to experiment with this learning method using LLMs as support technology. This will involve using LLMs as teachers by asking them for descriptions of events and using those descriptions as learning inputs. We also plan to test LLMs’ usefulness as knowledge verifiers by asking them to identify differences between a text describing a complex event and the learned script. We hope that this work will allow OntoAgents to quickly expand their library of actionable scripts.

Acknowledgements

We are deeply grateful to Dr. Marjorie McShane and Dr. Jesse English for their guidance, and to the Office of Naval Research for their support.

References

- Arndt, C. S, Nirenburg, S., & Oruganti, S. (2026). A Script for Learning Script: Automatic Acquisition of Procedural Knowledge in Ontologically Grounded Agents. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 48.
- McShane, M., & Nirenburg, S. (2021). 2: A Brief Overview of Natural Language Understanding by LEIAs. In *Linguistics for the Age of AI* (pp. 59-115). Essay, MIT Press.
- McShane, M., Nirenburg, S., & English, J. (2024). *Agents in the Long Game of AI: Computational Cognitive Modeling for Trustworthy, Hybrid AI*. MIT Press.
- Nirenburg, S. & Raskin, V. (2004). 7: The Static Knowledge Sources: Ontology, Fact Database and Lexicons. In *Ontological Semantics* (pp. 159-192). Essay, MIT Press.
- RPI LEIA Lab. (2026). *Interactive Demo Displaying an OntoAgent being Taught Via Instruction*. <https://www.leia-lab.com/harmonic/demo-teaching.html>
- Schank, R. C. & Abelson, R. P. (1977). *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*. Routledge.