Demonstration of PLOW: A Dialogue System for One-Shot Task Learning

James Allen, Nathanael Chambers, George Ferguson*, Lucian Galescu, Hyuckchul Jung, Mary Swift* and William Taysom

Florida Institute for Human and Machine Cognition, Pensacola, FL 32502

*Computer Science Department, University of Rochester, Rochester, NY 14627

Introduction

We describe a system that can learn new procedure models effectively from one demonstration by the user. Previous work to learn tasks through observing a demonstration (e.g., Lent & Laird, 2001) has required observing many examples of the same task. One-shot learning of tasks presents a significant challenge because the observed sequence is inherently incomplete – the user only performs the steps required for the current situation. Furthermore, their decision-making processes, which reflect the control structures in the procedure, are not revealed.

We will demonstrate a system called PLOW (Procedural Learning on the Web) that learns task knowledge through observation accompanied by a natural language “play-by-play”. Natural language (NL) alleviates many task learning problems by identifying (i) a useful level of abstraction of observed actions; (ii) parameter dependencies; (iii) hierarchical structure; (iv) semantic relationships between the task and the items involved in the actions; and (v) control constructs not otherwise observable. Various specialized reasoning modules in the system communicate and collaborate with each other to interpret the user’s intentions, build a task model based on the interpretation, and check consistency between the learned task and prior knowledge.

The play-by-play approach in NL enables our task learning system to build a task with high-level constructs that are not inferable from observed actions alone. In addition to the knowledge about task structure, NL also provides critical information to transform the observed actions into more robust and reliable executable forms. Our system learns how to find objects used in the task, unifying the linguistic information of the objects with the semantic representations of the user’s NL descriptions about them. The objects can then be reliably found in dynamic and

complex environments. See Jung et al (2006) and Chambers et al (2006) for more details on the PLOW system.

The PLOW System

PLOW learns tasks executable on the web involving actions such as navigation, information extraction and form filling, and can learn iterative steps that operate over lists of objects on pages. Figure 1 shows the system during learning a task to find publications for a specified author. Upper left is the Mozilla browser, in which the user can demonstrate action and the system can execute actions in a mixed-initiative fashion. The user may speak or type to the system (SR output is lower right), and PLOW combines knowledge from the language and the demonstrated actions to produce a parameterized procedure (described in generated natural language in the upper right corner). Figure 2 shows a complete training dialogue in which PLOW learns how to find article titles. To save space, simple acknowledgments by the system are not shown.

Figure 1: PLOW learning a task

U: Let me show you how to find someone's publications
 U: Go to this webpage
 Navigates to <http://scholar.google.com>
 U: Click on the advanced search link
 Clicks Advanced Scholar Search
 U: Put the author's name here
 Types the name in the Author text field
 U: Click the search button
 Clicks on the Search Scholar button
 U: OK here is the list of results
 Selects the table on page
 PLOW: Can you check that I understand the data.
 PLOW: *puts up a new window with the results.*
 U: OK this is good.
 U: this is the article title
 Highlights a title
 PLOW: *attempts to find all titles and displays them*

U: This is wrong.
 PLOW OK
 U: This is another article title
 Highlights another correct title
 PLOW: *recomputes all the titles and displays them*
 Switches back to Scholar window
 U: OK click next to get more results
 Clicks the Next link
 PLOW: When should I stop searching?
 U: Get thirty results
 PLOW: OK. I will do that
 PLOW: *iterates through pages until it gets 30 results*
 PLOW: I'm done.
 U: We're done with the demonstration
 U: Let's practice
 PLOW: what is the name of the author?
 Enters a new name
 PLOW *executes the procedure successfully!!*

Figure 2: Learning to find titles of an author's publications (U = user)

Evaluation

The PLOW system was evaluated by independent evaluators who considered four task learning systems developed in the CALO project. There were 16 human subjects who received training on each of the systems and who worked through a number of successful scripted training sessions with each. They were then given ten new problems, ranging from slight variations to problems they had seen to problems that were substantially new. They were free to choose which problems to work on and which system to use and the resulting tasks learned were tested with different settings of the parameters and scored out of a total of 4 points based on a complex predefined evaluation criteria (not known to the developers). The PLOW system did well in the test, not only receiving the highest average score on tasks learned by a system (figure 3) but also was strongly preferred by the users and selected more than half the time (figure 4).

Figure 3: Average score (out of 4)

The Demonstration

If we are allowed a presentation we will demonstrate PLOW live on a task selected by the audience. In addition, we would like to have the system available for an extended period of time during the conference so that attendees can spend time using the system to teach it simple tasks. The system runs on a laptop and all that is needed for a demo is internet access.

Acknowledgements & References

This work was supported by DARPA grant NBCH-D-03-0010 under a subcontract from SRI International, ONR grant N000140510314, and NSF grant 5-28096.
 Chambers, N. et al. (2006). *Using Semantics to Identify Web Objects*. Proceedings AAAI.
 Jung, H., J. Allen, et al. (2006). *One-Shot Procedure Learning from Instruction and Observation*. FLAIRS, Melbourne, FL.
 Lent, M. and Laird, J. (2001) *Learning Procedural Knowledge through Observation*, Proc. of the Intl Conf. on Knowledge Capture.

Figure 4: User preference for systems (55 trials)