Four Kinds of Learning in One Agent-Oriented Environment

Sergei NIRENBURG[a] , Marjorie McSHANE[a,1], Stephen BEALE[a] , Jesse ENGLISH[a],
Roberta CATIZONE[b]
[a] *University of Maryland Baltimore County*
ITE 325, 1000 Hilltop Circle
Baltimore, MD 21250
[b] *Onyx Consulting*
{sergei, marge, sbeale, english1}@umbc.edu
roberta@dcs.shef.ac.uk

Abstract. This paper briefly describes four kinds of learning carried out by intelligent agents in a computational environment facilitating joint activities of people and software agents. The types of learning and the applications we draw examples from are: learning by being told and learning by experience, as illustrated through a virtual patient application; learning by reasoning, as illustrated through a clinician's advisor application; and learning by reading, as illustrated by an ontology enhancement application. The agents carrying out these types of learning are modeled using cognitive modeling strategies that show marked parallels with how humans seem to learn.

Keywords. machine learning, learning by reading, virtual patient, agent network, decision-making, computational semantics

Introduction

In this paper we present four kinds of human-like learning carried out by cognitively modeled agents in an environment that can support various applications. The paper is not technical; rather, it presents a new, aggregated view of learning in the OntoAgent environment and will serve as an introduction to the systems to be demonstrated at the conference. Interested readers can find technical descriptions of each type of learning in referenced papers.

The core capabilities of the agents in this environment include the following: they are designed to operate in a hybrid network of human and artificial agents; they emulate human information processing capabilities by modeling conscious perception and action, which includes reasoning and decision making; they can communicate with people using natural language; they can incorporate a physiological model, making them what we call "double agents" with simulated bodies as well as simulated minds; they can be endowed with personality traits, preferences and psychological states that affect their perceived or subconscious preferences in decision-making; their means of perception include language understanding and interoception, which is the experiencing of sensations from one's body; their underlying principles, knowledge resources and processors are broad-coverage rather than geared at a particular application, which makes them – after a modicum of inevitable knowledge base refinement – portable to a variety of domains and application configurations; and the perception and action algorithms used by the agents are supported by and, in turn, augment the agents' memory of event, state and object instances to complement its ontological knowledge of event, state and object types.

What makes modeling such multi-faceted agents feasible is that all aspects of agent functioning are supported by the same ontologically-grounded knowledge substrate encoded in a single metalanguage [1]. That is, no matter how information enters an agent's consciousness – be it through dialog, reading texts, perceiving its own bodily signals, etc. – it must be automatically converted into the unambiguous metalanguage used for reasoning. Let us take, as an example, the case of language processing. If our agent receives the dialog input *You have achalasia*, it uses its knowledge of language, lexicon and ontology to convert that string into the semantic interpretation ACHALASIA-1 (EXPERIENCER HUMAN-1). (This is actually a shorthand version of the representation generated and used by the system.) The words in small caps are ontological concepts and the appended numbers indicate instances of those concepts. The instance HUMAN-1 would be coreferred in the agent's memory with itself.

As a testbed for the development of our extended cognitive architecture ("extended" because cognitive architectures typically do not include physiological simulation), we have been modeling agents in the medical domain. Our first application, Maryland Virtual Patient [1-3], is an environment for training medical personnel in patient management – diagnostics, treatment and care over time. Our second application, built on the same knowledge and processing substrate, is a CLinician's ADdvisor (CLAD), which is intended to assist practicing clinicians by providing advice and thus reducing their cognitive load. In parallel with these applications, we have been working on teaching our agents to learn ontology and lexicon by reading. This capability can be used at runtime to recover from so-called unexpected input and over a longer period of time to enhance the resources central to all agent communication.

To illustrate different types of agent learning, we select just one agent in one application for each type of learning. However, it must be emphasized that we are developing agent *capabilities* which are then applied to all agents in the network for which they are applicable.

It must be emphasized from the outset that the learning carried out in the OntoAgent environment does not readily lend itself to comparison with traditional machine learning paradigms because (a) it is not centrally statistical, (b) it relies on understanding the meaning of text and agent experiences and (c) the learned material is integrated into a multi-dimensional and ever-growing agent memory in which interpreted information is recorded in a metalanguage that supports automatic reasoning. Unfortunately, space does not permit a detailed comparison with traditional machine learning paradigms.

1. Learning by Being Told in Maryland Virtual Patient

Maryland Virtual Patient (MVP) is a cognitive simulation and training system whose goal is to provide medical practitioners with the opportunity to develop clinical decision-making skills by managing many highly differentiated artificial intelligent agents playing the role of virtual patients (VPs). These VPs can suffer from various diseases and combinations of diseases, and are capable of realistic physiological and cognitive responses even to unexpected actions on the part of the user. The system seeks to offer a breadth of experience not attainable in a live clinical setting over a corresponding period of time, and a depth of experience that is not currently available in interactive VP training systems. In short, trainees can learn by trial and error using a large number of patients that present with clinically relevant variations of each disease.

MVP is configured as a network of human and artificial agents. The human agent, who is typically a medical trainee seeking to improve his or her cognitive decision making skills, plays the role of the attending physician. The core artificial agent, the VP, is a knowledge-based model and simulation of a person suffering from one or more diseases. The VP is a "double agent" in that it models and simulates both the physiological and the cognitive functionality of a human. Physiologically, it undergoes both normal and pathological processes and responds realistically both to expected and to unexpected (e.g., by user error) internal and external stimuli. Cognitively, it experiences symptoms, has lifestyle preferences (a model of character traits), has dynamic memory and learning capabilities, has the ability to reason in a context-sensitive way, and can communicate with the human user about its personal history, symptoms and preferences for treatment. Other intelligent agents in the network include medical specialists and technicians and a tutor. Communication between the user and the intelligent agents is carried out in unrestricted English.

Two of the important cognitive features of virtual patients in MVP is their ability to learn and their ability to make decisions that reflect their personal preferences, character traits, etc. In fact, learning is often a prerequisite to making a decision: after all, even virtual patients should not make uninformed decisions. Table 1 shows a brief dialog between a virtual patient (VP) and the human user/doctor (D) that features the learning of ontology and lexicon in preparation for decision-making. We present the dialog in tabular form and not in the implemented, fully functional interface for reasons of space.

Table 1. A dialog in MVP and the ontological and lexical knowledge learned by the VP.

Dialog	Ontology learned	Lexicon learned
D: You have achalasia.	The concept ACHALASIA is learned and made a child of DISEASE.	The noun "achalasia" is learned and mapped to the concept ACHALASIA.
VP: Is it treatable? D: Yes.	The property TREATABLE in the ACHALASIA frame has its value set to 'yes'.	
D: I think you should have a Heller myotomy.	The concept HELLER-MYOTOMY is learned and made a child of MEDICAL-PROCEDURE. Its property TREATMENT-OPTION-FOR receives the filler HELLER-MYOTOMY.	The noun "Heller myotomy" is learned and mapped to the concept HELLER-MYOTOMY.
VP: What is that? D: It is a type of esophageal surgery.	The concept HELLER-MYOTOMY is moved in the ontology tree: it is made a child of SURGICAL-PROCEDURE. Also, the THEME of HELLER-MYOTOMY is specified as ESOPHAGUS.	
VP: Are there any other options? D: Yes, you could have a pneumatic dilation instead…	The concept PNEUMATIC-DILATION is learned and made a child of MEDICAL-PROCEDURE.	The noun "pneumatic dilation" is learned and mapped to the concept PNEUMATIC-DILATION.
D: (cont) which is an	PNEUMATIC-DILATION is moved from	

endoscopic procedure.	being a child of MEDICAL-PROCEDURE to being a child of ENDOSCOPY.
VP: Does it hurt? D: Not much.	The value of the property PAIN-LEVEL in PNEUMATIC-DILATION is set to .2 (on a scale of 0-1).

When the VP processes each of the doctor's utterances, it automatically creates text meaning representations that it then uses for purposes of reasoning and learning. Recall that we already saw the text meaning representation of the first sentence: ACHALASIA-1 (EXPERIENCER HUMAN-1). Semantically-oriented text analysis, of course, involves extensive reasoning about language and the world, and learning ontology and lexicon involves additional reasoning. For example, how does the VP know to make ACHALASIA a child of DISEASE? It combines its lexical knowledge of the possible meanings of the word *have* (one of which expects a disease as its direct object) with knowledge of the speech context (the VP is in the doctor's office) to hypothesize the meaning of the unknown word. A similar type of reasoning is used to suggest to the VP that a Heller myotomy is some sort of medical procedure. Our short dialog also shows two examples of belief revision: when the VP learns more about the nature of the procedures HELLER-MYOTOMY and PNEUMATIC-DILATION, it selects more specific parents for them, thereby permitting the inheritance of more specific property values.

2. Learning by Experience in Maryland Virtual Patient

Virtual patients in MVP have a simulated life that is sufficiently rich to support the needs of the teaching application. For example, they make medically relevant lifestyle decisions (whether or not to smoke, take their medication, etc.), they go to the doctor, converse and negotiate with him or her, have procedures carried out on them, and so on. They also experience symptoms of their disease through the process of interoception, which is the perception of the body's signals as perceived by the mind.

Many aspects of a VP's simulated life lead to learning, defined as populating the VP's memory with new facts about types and instances of objects and events in the world. For example, when the VP experiences symptoms over time, it remembers them as new and changing features of itself; when it has a test carried out, it has a certain perception about the test that might or might not correlate with what that doctor said to expect; when it takes medication or has a procedure carried out, it may or may not feel better as a result, thus learning that the given intervention was or was not effective; and so on. All of the VP's new experiences are recorded in memory using the same ontological metalanguage used throughout the system; as such, these memories are available as input to the same reasoners that work on information learned in all of the other ways presented in this paper.

3. Learning by Reasoning in CLinician's ADvisor

CLAD, a CLinician's ADvisor, is a system intended to decrease the cognitive load of clinicians caused by the vast amount of information available, and to improve overall patient outcome through providing high-value decision-making assistance. It is intended to support the work of a live clinician managing a live patient. It will assist clinicians by providing advice (along with its justification), answering questions, providing prognoses, carrying out administrative tasks (e.g., finding out if a given procedure is covered by the patient's insurance company), and so on.

CLAD is endowed with the same expert disease model that drives the physiological simulation in MVP. This means that it knows the salient features of each disease, the different courses a particular patient's disease can follow, the typical range of time frames for each stage of a disease (as conceptually delineated in the model), the different interventions that can be used and their range of affects on different patients, an so on. In short, it knows the theme and variations of diseases.

One of CLAD's jobs, paralleling that of a clinician, is to attempt to determine which "profile" a patient matches in order to select the best course of treatment. For example, one patient might be experiencing very slow disease progression and can be monitored at long intervals, while another is experiencing very fast disease progression and must be treated at once; one patient might have a history that puts him at risk for disease complications while another might have little likelihood for such complications; and so on. If CLAD knows nothing about a patient, it cannot know which profile he or she matches and can only reason in general terms, using population based statistics. However, as it learns more and more about the patient, through the information recorded in the patient chart, it can reason about the patient's most likely profile. This belief revision permits CLAD to provide more individually catered advice over time. For example, assume that a patient presents to the doctor complaining of difficulty swallowing and chest pain. CLAD hypothesizes that the patient has the disease achalasia and advises an upper endoscopy and a barium swallow. The doctor orders these tests and the results suggest a very early stage of achalasia, but without sufficient evidence for a definitive diagnosis. The doctor has the patient come for a follow-up and another barium swallow 6 months later. The patient's symptoms are about the same and the test results show only slight increase in relevant indicators. There is still insufficient evidence to posit a definitive diagnosis. Using its temporally sensitive disease model, CLAD concludes that the patient most likely has very slowly progressing achalasia (a revision from fully generalized achalasia) and no definitive diagnosis will be possible for 1.5 – 2 years. It suggests that the next follow-up be in 1.5 years unless the patient experiences some marked increase in symptoms.

4. Learning by Reading for Automatic Ontology and Lexicon Enhancement

The OntoAgent environment requires high-quality lexical and ontological knowledge bases that to date have been compiled manually. Manual acquisition of resources, however, is expensive – so expensive that some (e.g., [4]) consider it infeasible. One way to both alleviate the cost and also permit the system to better recover from so-called unexpected input during runtime processing is to prepare it to learn ontological concepts and lexical senses on the fly. The methodology we have been experimenting with is learning by reading [5, 6]. When our text processing agent encounters an unknown word, it behaves in a similar way as a person might: it uses the available properties of the word presented in the text to help it to guess what the word might mean. But rather than be limited to just the local context, it can search the Web for other contexts that use the word in an effort to learn other properties of it that will further constrain its meaning. Once it has compiled an inventory of property-value pairs, it matches them to existing concepts in the ontology in an attempt to find the most appropriate placement for it in the tree of inheritance. When the best placement has been found, the concept is provisionally added to the ontology (provisionally because it must be vetted by a person) and inherits fillers for all properties that were not explicitly defined during the machine learning process. In this way, a newly learned concept gets the fullest possible property-value profile, even though some of the values might be underspecified since they are inherited from the parent. Our work on learning by reading has not yet been incorporated into MVP or CLAD, though its incorporation and subsequent system-based testing is planned for the near future.

5. Closing Remarks

This paper has provided an informal overview of several types of human-like learning carried out by intelligent agents in our environment. The goal was not to explain the details of *how* this is done (which can be found in referenced works) but, rather, to show *what* is learned and *why*. The overarching goal of our program of research and development is to create intelligent agents that fulfill the original, lofty vision of artificial intelligence, but we seek to do that along a schedule that permits intermediate results to be of utility in near-term applications.

References

[1] S. Nirenburg and V. Raskin, *Ontological Semantics*, MIT Press, Cambridge, Mass., 2004.

[2] M. McShane, S. Nirenburg, S. Beale, B. Jarrell, and G. Fantry, Knowledge-based modeling and simulation of diseases with highly differentiated clinical manifestations, *Proceedings of the 11th Conference on Artificial Intelligence in Medicine*, 2007.

[3] S. Nirenburg, M. McShane, S. Beale, A simulated physiological/cognitive "double agent". *Proceedings of the Workshop on Naturally Inspired Cognitive Architectures at AAAI 2008 Fall Symposium*, 2008.

[4] C. Brewster, J. Iria, F. Ciravegna, and Y. Wilks, The ontology: Chimaera or Pegasus," In Proc. Dagstuhl Seminar on Machine Learning for the Semantic Web (2005), 13–18.

[5] J. C. English, *Learning by Reading: Automatic knowledge extraction through semantic analysis*, PhD Dissertation, University of Maryland Baltimore County, 2010.

[6] J. C. English and S. Nirenburg. 2010. Striking a Balance: Human and Computer Contributions to Learning through Semantic Analysis. Proceedings of ICSC-2010. Pittsburgh, PA.