

Curricula and Metrics to Investigate Human-Like Learning

Jacob Beal and Paul Robertson and Robert Laddaga

BBN Technologies
10 Moulton St
Cambridge, MA, USA 02138

Abstract


In order to learn from naive instructors, machines must learn more like how humans learn. We are organizing a “Bootstrapped Learning Cup” competition, in which competitors attempt to create the best learning agent for a curriculum whose focus is known but whose specifics are not. By focusing each competition on a particular factoring of the larger problem of human-like learning, we hope to simultaneously identify productive decompositions of the larger problem and components that can eventually be integrated to solve it. To this end, we seek to measure learner autonomy with “spectrum curricula” that measure learning against an incrementally varied set of curricula, ranging from extremely telegraphic to overly detailed.

If this program of competitions is successful, it will lead to a revolutionary change in the deployability of machine intelligence, by allowing human curricula to be used in configuring a system for an application area and by allowing machines to “culturally” adapt to the specifics of their deployment. The tools and models developed during this effort may also lead to significant improvements in our models of human learning and cognition.


Learning from Naive Instructors

If humans and machines are to work together in challenging environments, then machines must be able to learn from humans in a manner similar to how humans learn from one another. The environments that humans act in are highly diverse and constantly changing—whether in the office, at home, or on a battlefield. Even a machine with an exhaustive background knowledge must be customized for the particular environment in which it acts, just as a secretary must become accustomed to the peculiarities of a particular office’s policies and customs or a soldier learn the particular patterns and cautions of a deployment area.

At present, effectively instructing a machine generally requires that one be a highly trained expert in a field like programming, knowledge engineering, or machine learning. Typical practice is to either insert the knowledge directly—a delicate, expensive and error-prone process—or else to simply expose the machine to a vast pile of data and hope that it



(a) Blocks World



(b) RoboCup

Figure 1: Some possible domains for Bootstrapped Learning competition curricula.

has learned what one desired.¹

This is not how humans teach one another. Outside of formal classroom settings, humans teach one another informally, imprecisely, and highly effectively (and even a classroom is extremely informal and imprecise compared to the current requirements for machine instruction). Cultural knowledge, like how to change a flat tire, what to do when you’re sick, or fun games to play together, passes easily from person to person in a population, even when the teacher has no training as an instructor.

For humans and machines to work together routinely, machines must have this human-like capacity to learn from naive instructors. Neither the homemaker nor the soldier

¹There are, of course, notable exceptions, such as near-miss learning (Winston 1970).

have the time and inclination to bother with a machine partner otherwise. Although machines are far from human capabilities, much might be accomplished even by meeting the humans partway, as with Graffiti character recognition (MacKenzie and Zhang 1997). Even a partial success could be revolutionary, allowing machines to adapt “culturally” to the environments in which they are deployed and also allowing the vast reservoirs of human curricula to be used in configuring machines for application domains.

The Bootstrapped Learning Competition

We are beginning an investigation of this area by organizing a “Bootstrapped Learning” competition, modelled after the RoboCup competition and building off of infrastructure developed by DARPA’s Bootstrapped Learning program (DARPA IPTO Retrieved Nov 5 2008a).

The Bootstrapped Learning program is seeking to make systems naively instructable through a combination of spoken language, gestures, and demonstration. For purposes of investigation, the interaction of the student with the instructor and its environment has been constrained to symbolic messages in a formal interlingua. A curriculum is formed of a series of training and testing lessons arranged in a “ladder” leading from simply concepts up to an ultimate “capstone” concept, and each lesson is designated to use one or more out of a collection of specified “Natural Instruction Methods” (NIMs) such as “Example of Procedure” or “Feedback Using Explanation.”

We are taking the framework developed thus far by this program and coupling it with additional documentation, our own curricula, and a semi-competent base learner for competitors to build off. Competitors can then take their own work and add learning components with particular specialties—either in type of learning or type of teaching—and have the base learner allow them to take a narrow focus and still remain competitive.

Methods of Evaluation

With a more well-defined problem, such as character recognition, evaluation can be a straightforward measure of error rate, and ordinary engineering methods applied to make progress. Human-like learning, however, is so broad an area and likely so tangled a problem that any near-term work must address only simplified versions of the problem.

Some of the many problems inherent in learning from naive humans include:

- missing and noisy inputs
- ambiguity of instruction (ranging from individual references to unstated assumptions about the material)
- literal incorrectness of the teacher
- interpretation of natural input (language, gesture, hands-on, observed actions)
- characterizing the cultural “protocols” used by human instructors in natural instruction

We are thus faced with an age-old question of AI: how can we make a problem tractable enough for progress, without

turning it into a toy whose solution has no scientific merit? This problem is particularly acute since we are organizing a competition, and the evaluation method may strongly shape the decisions participants make about what ideas to investigate.

In our judgement, there are three basic methods for evaluating the scientific merit of an idea in this domain:

- The tests can be made extremely complicated and difficult, so that even partially solving them clearly implies that something important has been learned—either that the system with the partial solution has an important idea within it, or that there is some appropriate representation under which the problem is easier than it appeared.
- Given a strong hypothesis about the means by which a machine ought to be able to solve the problem, a fairly simple test may serve to verify the predictions of the hypothesis relative to other, competing hypotheses.
- Given a strong hypothesis about the key differences between human capabilities and current machine capabilities, a fairly simple test can show progress targeted toward those capabilities.

The first method is typical of major competitions such as the DARPA Urban Challenge (DARPA IPTO Retrieved Nov 5 2008b). This approach effectively favors systems integration to the detriment of innovation in components, because a significant failure in any single part of the system will generally cripple it with respect to the holistic performance measure. We believe that it is premature to look for integration in this problem, as many of the components needed are still immature or do not exist.

The second method is appropriate when the same team is designing the measure and the solution hypothesis, but not applicable for a competition where the point is to encourage investigation of many possible solutions.

The third method, on the other hand, lets us focus investigation on narrow sub-problems, simultaneously investigating how to factor the larger problem and developing ideas and tools that will ultimately contribute to its solution.

Each competition, then, will be centered around a hypothesis that some particular axis where human learners and current machine learners are different is both important and isolatable. The curriculum and scoring for the competition will be centered around testing for human-like behavior on that axis. Although the curriculum will not be known to competitors, they will be given information about how it tests the axis of interest. The results will thus give evidence toward hypotheses at two different levels:

- If the competition entries include elegant solutions that perform decently, then this axis of investigation is likely to be a good factoring of the larger human-like learning problem.
- The performance of particular entries reflects the quality of their investigators’ hypotheses about the narrow sub-problem.

When both indication of a good factoring and high-quality solutions exist, we intend to refine the entries with the best combination of elegance and performance into components

that are added to the base system as tools to be drawn upon by future competitors. These tools would also be good candidates for cognitive science investigation of whether human behavior is consistent with the operation of such a tool, and also for eventual integration into a broad model of human-like learning.

Tentative Hypothesis and Metrics

For the first contest we are organizing, intended to run in 2009, we have tentatively chosen to focus on autonomy in the learner. We hypothesize that an important quality of human-like learning is that the student is always actively hypothesizing about structures in the environment. If this is the case, then the role of the teacher is generally to help guide the student's choice between competing models, rather than to give a model to the student. Thus, for example, the main contribution of an electronics teacher would not be telling a student about Kirchhoff's Laws, but the set of examples and graduated problems that help a student sort out ambiguities and misunderstandings in how to apply them.

If this hypothesis is the case, then we would expect a student exhibiting such "autonomy" to have the following properties:

- Learning can be advanced by any signal that clearly favors one model out of a set of possibilities, even if the contents of the signal are ambiguous or incorrect.
- A student will "leap ahead" and fill in small gaps in knowledge without any intervention from the teacher.
- Easily perceptible affordances in the environment will trigger learning without any intervention from the teacher.

Spectrum Curricula

We are therefore developing "spectrum curricula" that measure incremental changes in ability associated with particular aspects of instruction-based learning. A spectrum curriculum is a collection of teaching lessons that all target the same concept using the same natural instruction method. The lessons in the collection vary incrementally with regards to one aspect of interest (such as autonomy), from highly challenging to entirely absent.

At the challenging extreme, the aspect is emphasized so strongly that even humans may find the lessons too difficult, while at the zero-challenge end even simple machine learning approaches should be sufficient and a human may find the lesson challenging due to its artificiality. For example, a spectrum emphasizing teacher incorrectness might range from the teacher making a total hash of the lesson to the teacher being perfectly logically precise and correct in every statement and action.

The curriculum begins with a thorough test measuring the student's uninstructed ability, then progresses from most to least challenging. Between each teaching lesson, the student is tested again (using different but equivalent questions). The product of such a curriculum is thus a curve showing how the student's cumulative understanding improves as the aspect varies.

Testing an instruction-based learning system against several curricula for the same aspect should produce a good

measure of the system's performance with respect to that aspect. No system should be expected to perform well on the most challenging lessons, but the performance of well-designed system should degrade gracefully as the challenge increases. Moreover, a spectrum that ranges from too challenging to zero challenge should also be able to provide a fine-grained measure of performance on which incremental progress towards success can be easily tracked.

Example Curricula for Learner Autonomy

In order to focus on learner autonomy, we are developing curricula with varying levels of ambiguity in instruction in the RoboCup and Blocks World domains—in particular ambiguity that comes from information omitted from the curriculum by the teacher. Our ambiguity spectra thus begin with teaching lessons that are extremely telegraphic, providing almost no information to the student, and end up extremely detailed, providing masses of information to the student.

At the extreme telegraphic end, learning may be difficult even for humans. As more detail is added, the lesson should become easy for a human to learn from. Eventually, as the level of detail continues to increase, learning should become easy for conventional machine learning algorithms. At the high detail end, however, it may once again be hard for humans to learn from due to the excess of data—for example, a human may become confused or believe that the concept being taught must be more complicated than it is because the instruction is so detailed.

We estimate that 10 spectrum curricula should be sufficient to provide a good measure of capability in coping with ambiguity in instruction. There are many pieces of knowledge in the RoboCup and Blocks World domains that make good spectrums, such as:

- Learning "out of bounds" in RoboCup by example: from a single gesture to many randomly scattered examples.
- Learning to cope with the Sussman Anomaly in Blocks World by telling: from just indicating the critical block to explaining every step in the process.
- Learning when to pass in RoboCup keep-away by feedback: from pointing out the key opponent to giving the precise formula.

The "out of bounds" curriculum, for example, consists of a seven step spectrum. In order from most detailed to least detailed (with the additional assumptions expected to be necessary in parentheses) the steps are:

1. 50 random locations, half expected to be out of bounds
2. 20 random locations scattered near the boundary, half expected to be out of bounds (homogeneity away from examples).
3. 4 pairs of examples, one in and one out on each boundary line (four pairs means four sides, probably a rectangle).
4. Two examples (inside or outside) at opposite corners (the shape is a rectangle and the examples are at its corners).
5. Point to the boundary line and give one example in and one out (the visible line is the decision boundary).

6. Point to the boundary line and give one example (with a boolean function, outside the line will be opposite to inside).
7. Just point to the boundary line (play is likely to be legal inside a line).

Testing asks about ten random locations, half expected to be out of bounds. The student is evaluated by starting with the most telegraphic (“Just the line”) example, producing a curve of seven points.

Contributions

In order to learn from naive instructors, machines must learn more like how humans learn. We are organizing a “Bootstrapped Learning Cup” competition, in which competitors attempt to create the best learning agent for a curriculum whose focus is known but whose specifics are not. By focusing each competition on a particular factoring of the larger problem of human-like learning, we hope to simultaneously identify productive decompositions of the larger problem and components that can eventually be integrated to solve it.

If successful, this effort will lead to a revolutionary change in the deployability of machine intelligence, by allowing human curricula to be used in configuring a system for an application area and by allowing machines to “culturally” adapt to the specifics of their deployment. The tools and models developed during this effort may also lead to significant improvements in our models of human learning and cognition.

References

- DARPA IPTO. (Retrieved Nov. 5, 2008)a. Bootstrapped learning. <http://www.darpa.mil/ipto/programs/bl/bl.asp>.
- DARPA IPTO. (Retrieved Nov. 5, 2008)b. Darpa urban challenge. <http://www.darpa.mil/GRANDCHALLENGE/>.
- MacKenzie, I. S., and Zhang, S. X. 1997. The immediate usability of graffiti. In *Graphics Interface '97*, 129–137.
- Winston, P. 1970. *Learning Structural Descriptions from Examples*. Ph.D. Dissertation, MIT.