

FUTURE RESEARCH

Alice Armstrong

Summary of Current Research

In my current research, I have developed a technique called mentoring which is applied to a statistical, adaptive learning engine called a Collective Learning Automaton (CLA). Mentoring reduces the number of random decisions made by the learning agent and improves learning performance on classification problems by assuming that a response that is effective in an unrelated situation may also be appropriate in the current, uncertain one. This process has its foundations in Piaget's theories of neonatal development.

Mentoring has been shown to dramatically improve learning, especially early in the learning process. It is effective in stable environments (in which the learning task does not change over time) and in unstable environments (in which the learning task shifts radically). Additionally, mentoring does not apply any feature comparison of either the inputs or the classes. This trait makes it a very powerful technique that can help compensate for poor feature selection.

In mentoring, an individual stimulant passes through several different stages as it becomes confident in one respondent. As a student, the stimulant will choose a mentor respondent, a respondent that has proven effective for other stimulants, instead of following the standard procedure of exploring the response range. When a student reaches a certain level of maturity, it will become independent of its mentor and will explore the response range on its own. When an independent becomes fully confident in a respondent, the stimulant elects the respondent to become a mentor and becomes a supporter of that mentor. The thresholds marking the transitions between these stages measure selection confidence, the confidence that a stimulant has that it should select the respondent with the largest probability.

Possible Extensions

Hierarchical CLAs to control the mentoring thresholds

Currently, all the mentoring thresholds are fixed during learning; however, informal observation suggests that adjusting these thresholds during learning could be very effective. But how to decide when and by how much to adjust these thresholds? A metric called *confident evaluation*, a combination of the average selection confidence of the entire CLA and the average evaluation of the responses made during a test period, could be useful in determining the need for adjustment.

An individual CLA does not have access to its confident evaluation because a CLA does not receive the results of tests periods. No learning happens during the test period and, therefore, no feedback is generated for the CLA. In order to incorporate the confident evaluation, an additional CLA could be added to use the confident evaluation as input and adjust the mentoring thresholds as its response. The original CLA would remain unchanged, using the stimulants and evaluation that come directly from the environment.

Environmental- and Learner-directed moderated collection length

The collection length determines how many responses the CLA makes between evaluations. In my research, the collection length is fixed during the entire learning process. This does not seem to reflect the way things happen in reality. A neonate receives almost constant feedback on every single action it makes. Think of the amount of attention that parents lavish on their newborns. As the child gets older, she will not need constant reminders, correction, and encouragement for tasks that she has already mastered such as successfully reaching for objects. In fact, too much feedback can become a distraction.

I am interested in devising a way to control the amount of feedback given to a learner based on its “age”. I want to consider environmental control methods (*i.e.* the parent deciding how much guidance to give) and internal control methods (*i.e.* allowing the learner to decide how much attention to pay to the environment).

Application to Pre-supervised Learning

Mentoring has only been applied in online learning problems in which the learner is evaluated periodically. I am curious to see if mentoring can be useful in a pre-supervised classifier. I have worked with the ALISA classifier, which utilizes a CLA, and I would like to add a module that would apply mentoring to the one-time classification.

Increasing the Effectiveness of human-derived evaluation

I am interested in the ability of mentoring to maximize the utility of very expensive human evaluation of a classifier. People’s time is so valuable that it is rarely useful to try to improve a classifier’s performance by having a technician critique the output; however, mentoring might be able to boost the efficacy of these corrections by improving the probability that the human-feedback would be used in new situations.

Combining Mentoring with Feature Vector Comparison

I designed mentoring to be as general as possible, which means that there is no feature comparison for either the inputs or the responses. However, many problems do lend themselves to feature comparison for clustering and generalization, and I look forward to seeing how mentoring can contribute to this process.

Other Areas of Interest*Natural Language Processing*

Recently, exciting advances have been made in statistical methods for analyzing natural language without using text that has been tagged with parts of speech. I am very interested in investigating this area in more depth.

Computer Science Education

Computer science requires students to think in a way that is new and not instinctive for most. I love watching students undergo the shift in perspective as they start to think of complex problems as a series of smaller problems. I explain this shift as being similar to the one I experienced when I learned to draw. Faces were very difficult for me to draw and it frustrated me. One day, I became aware that a face was not face, but rather a

collection of shapes, lights, and shadow. I could draw that. When students start to see how they can break down big problems, they acquire a new level of confidence in their abilities as problem solvers.

I want to become involved in studying and improving computer science education at the undergraduate level and I want to understand what must be done in primary and secondary education to encourage strong engineering students, especially young women.

Conclusion

My interest in machine learning was sparked by my love of both computer science and teaching. I am fascinated by all aspects of learning, and my current research interests reflect that.