In Proceedings of the Sixth International Conference on Cognitive Modeling, 412-413. Mahwah, NJ: Lawrence Earlbaum.

Generalization and Discrimination in a Semantic Network Trained with Semi-Supervised Learning

Rebecca J. Robare (RobareOwl@att.net)
Psychology/Experimental Cognition, Graduate Center, CUNY
Department of Psychology, City College of New York
138th St. at Convent Ave
New York, NY 10031

Keywords: semi-supervised learning, semantics, language learning

Abstract

The goal of the proposed research is to explore semantic learning in an artificial neural network trained with a semi-supervised learning paradigm. Semi-supervised learning permits both labeled and unlabeled data to be used in training the network, giving the training procedure greater ecological validity than seen in fully supervised learning, and improving the generalization ability of the network beyond that seen in networks trained with full supervision. As there is evidence that both supervised and unsupervised learning are necessary for learning word meaning, semi-supervised learning can provide a balanced approach to investigating semantic learning.

Semi-supervised learning can be defined as the use of both labeled and unlabeled data in the training of an artificial neural network. Such learning takes the form of expectation-maximization (EM) algorithms (i.e., Nigam, McCallum, Thrun, & Mitchell, 2000), support vector machines (SVM) (i.e., Chen, Wang, & Dong, 2003), and a variety of other forms. In the computer sciences, these algorithms have been used successfully in applications for various practical purposes, especially text classification in documents or Web sites. Semi-supervised learning has the following advantages over more traditional, fully supervised learning: (1) it is less costly in terms of time and labor to train from unlabeled examples than to manually label many items of information (Nigam et al., 2000); (2) the trained classifier or network shows better generalization, that is, it is more successful at appropriately classifying novel items of information (Anagnostopoulos et al., 2003); and (3) overfitting of the data is thereby avoided (Nigam et al., 2000). Improved generalization is, perhaps, the most important of these advantages. In fully supervised learning, when performance on a set of patterns is perfect, there will be no generalization; the network will not correctly classify novel patterns. Using unlabeled input patterns can improve generalization performance at test because the network is not overtrained on the specific labeled examples included in the training set (Anagnostopoulos et al., 2003; Chen, Wang, & Dong, 2003). Semi-supervised learning has the following advantages over learning mechanisms that are entirely unsupervised: (1) class labels can be learned; and (2) training exemplars that are similar on a physical dimension but should appropriately be classified differently, can be (Nigam et al., 2000).

Semi-supervised learning has proven successful in computer science applications. However, it has not been applied to any of the problems in cognition toward which artificial neural networks are often used. In connectionist psychology, learning is either fully supervised or unsupervised; the "environment" in the form of the "teacher" or labeled training examples provides either perfect feedback or none (O'Reilly & Munakata, 2000). This absence of semi-supervised learning is a gap in psychology's simulation of and understanding of human cognition. I offer the following two justifications for researching semi-supervised learning in psychology.

First, environments are inconstant and may not provide feedback consistently or perfectly. One example of this is language. It is often claimed by linguists that the feedback present in a language learner's environment is not sufficient to prevent learners from making linguistic mistakes that they do not, in fact, make (Pinker, 1999). And although some lexical items, such as some nouns and verbs, may be learned without supervision (Roy, 2000), it seems unlikely that generative syntax can be learned without some sort of supervision mechanism (Bloom, 2000). Semi-supervised learning is an elegant possible solution to this problem.

Second, semi-supervised learning is at least as biologically plausible as supervised or unsupervised learning. Mechanisms for both have been found within the brain, and these are generally assumed to operate over all synapses (O'Reilly & Munakata, 2000). In other words, both kinds of learning can operate over the same synapse simultaneously. If this is the case, then observing that the synapse is plastic as a result of both kinds of learning, under various environmental conditions, this could constitute semi-supervised learning at the synaptic level.

Preliminary research on the utility of semi-supervision has been encouraging. For the purposes of constructing neural networks, semi-supervised learning can be used with learning algorithms already applied in psychology. I have used the Leabra algorithm (O'Reilly & Munakata, 2000; O'Reilly, 2001) sucessfully with semi-supervised learning. Leabra is ideal for this purpose because it combines error-driven and Hebbian learning, thus providing a framework in which semi-supervised learning can operate. Furthermore, Leabra was designed with the intent that it be biologically plausible. Applying semi-supervised learning in this structure, therefore, can be seen to not contradict current knowledge about synaptic learning.

Using the PDP++ software (O'Reilly et al., 1995) it is easy to construct a data set for training the network that includes both labeled and unlabeled data. It is the use of unlabeled data in this context that is novel to psychology; earlier networks have used only labeled or unlabeled data, never both. Even Leabra has been used only to apply Hebbian learning to data that is otherwise learned in a fully supervised fashion. Simple Leabra networks trained by semi-supervised learning so far have shown task learning as good as that of fully supervised networks and generalization that is better in a task requiring a non-linearly separable discrimination.

My dissertation project will be to explore the potential of semi-supervised learning as a substrate for language learning. Specifically, I will be exploring semantic learning occurring in a supervised or unsupervised versus a semi-supervised fashion. There is reason to believe that semantic learning requires both types of learning processes (Bloom, 2000), so semi-supervised learning very naturally assumes this role. The overarching intention is to make this learning as realistic as possible, and to this end it is anticipated that word meanings will be trained using inputs that represent sentences, "images" or some sort of corresponding visual input (see Roy, 2000), and other words, thus giving the training set a context that approximates in a general way the context of the language learner. It is not a good approximation; but it is a better approximation than a network that merely trains associative word-to-word mappings and attempts to make inferences about semantic organization.

According to Bloom (2000), sometimes the words a language learner hears are presented in a context in which the meaning can be extracted, such as when the mother says "cookie" and gives the child a cookie, and sometimes when it cannot be, as when the caregiver says, "Do you want a cookie?" when no cookies are present. Note that the first of these scenarios can be interpreted as a correlation-detection situation (unsupervised learning) and the second as a situation of reinforcing the expectation of a cookie with the eventual presence of a cookie (supervised learning). Neither type of learning on its own will suffice. Two distinct but interdependent learning abilities are required to extract the correct information from each scenario and put that information together to determine the meaning of "cookie." Semi-supervised learning is ideal in this situation; it allows the accurate simulation of both learning mechanisms and environmental conditions, which are necessary for the development of precise models of language acquisition and other cognitive processes.

References

Anagnostopoulos, G. C., Bharadwaj, M., Georgiopoulos, M., Verzi, S. J., & Heileman, G. L. (2003). Exemplar-based pattern recognition via semi-supervised learning. *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks (ISCNN '03).* Portland, Oregon. Retrieved 28 March 2003 from pegasus.cc.ucf.edu/~mbharadw/thesis/gca-ssl-icjnn-2003.pdf.

Bloom, P. (2000). *How children learn the meanings of words*. Cambridge, MA: MIT Press.

Chen, Y., Wang, G., & Dong, S. (2003). Learning with progressive transductive support vector machine. *Pattern Recognition Letters 24*, 1845-1856.

Nigam, K., McCallum, A. K., Thrun, S., & Mitchell, T. M. (2000). Text classification from labeled and unlabeled documents using EM. *Machine Learning 39*, 103-134.

O'Reilly, R. C. (2001). Generalization in interactive networks: The benefits of inhibitory competition and Hebbian learning. *Neural Computation 13*, 1199-1241.

O'Reilly, R. C., Dawson, C. K., McClelland, J. L., & Carnegie Mellon University. PDP++.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explanations in cognitive neuroscience.* Cambridge, MA: MIT Press.

Roy, D. (2000). Learning from multimodal observations. *IEEE International Conference on Multimedia and Expo (I)*, 579-582. Retrieved 7 May 2003 from dkroy.www.media.mit.edu/people/dkroy/.