In Proceedings of ICCM - 2007- Eighth International Conference on Cognitive Modeling. 305 - 305. Oxford, UK: Taylor & Francis/Psychology Press.

Interactions Between Data Labeling and Ratio of Hebbian to Error-Driven Learning in Mixed-Model Networks

Rebecca J. Robare (RobareOwl@gmail.com)
Graduate Center, City University of New York
770 Bronx River Road #A64
Bronxville, NY 10708

Cognitive models using mixed-model (that is, combining both task and model learning) learning algorithms, such as Leabra (O'Reilly, 2001), have typically been simulated with supervised-only paradigms. Despite decades of research on learning from partial data (e.g., Dempster, Laird, & Rubin, 1997), only recently have paradigms such as semi-supervised and active learning been applied to more psychologically-oriented models (Love, Medin, & Gureckis, 2004; Robare, 2004). This has been driven, at least in part, by a concern for creating models more reflective of real-world learning environments (Elwin et al., 2007; Li, Farkas, & MacWhinney, 2004; Robare, 2004), which should make mixed-model algorithms and semi-supervised and active learning paradigms more popular in the future.

As paradigms like semi-supervised learning are applied to connectionist networks, often used to model cognitive phenomena in an explicitly brain-based fashion, some unexpected interactions have been observed. Leabra users know that the ratio of task to model learning (the *lmix* variable in the PDP++ software; O'Reilly et al., 1995) interacts with learning rate, *k*-winners-take-all inhibition (kWTA), and the number of epochs needed to train the network (O'Reilly & Munakata, 2000). The ratio of labeled to unlabeled data also interacts with these variables. Although similar interactions have been explored in earlier research on semi-supervised learning and its related paradigms (e.g., Nigam et al., 2000), particularly in the doman of training duration, related as it is to the cost reduction for which semi-supervised learning was initially developed, with the application of semi-supervised learning to connectionist models it is worthwhile to study these interactions anew. There are two primary reasons for this.

First, when the goal of cognitive modeling is not the solving of a computational problem or the development of an artificial intelligence but the explication of human cognition, a different set of considerations apply. Time and cost of human labor, a driving force in other semi-supervised learning research (e.g., Nigam et al., 2000; Yarowsky, 1995), becomes less relevant (though cost may be important in other ways; Antrobus, personal communication, March 2007). Neuroscience provides constraints absent from purely computational research (Sejnowski, 1986), and O'Reilly (1998) has argued persuasively for other constraints on connectionist networks that follow from the requirement for biological plausibility. These constraints mean, or may mean, that the labeled:unlabeled data ratio interacts differently with other variables in a biologically plausible connectionist network than in other kinds of cognitive architectures.

Second, because the drive for using semi-supervised learning in psychologically-oriented models comes not from neuroscience-based constraints but from environment-based constraints, the specific interactions of labeled:unlabeled data ratios with other network variables may have implications for the way we study feedback and reinforcement in real-world environments. Only recently has the notion of selective feedback moved outside the realm of respondent (Rescorla & Wagner, 1972) and operant (Hemmes & Eckerman, 1972) conditioning into cognitive research (Elwin et al., 2007). If we understand, in our models, the relationship between the labeled:unlabeled data ratio and the variables with which it interacts, we will be better able to predict the results of empirical studies on partial and selective feedback and be able to direct observations on the nature of feedback in real-world environments.

The current work focuses on the interaction between the labeled:unlabeled data ratio and the task learning: model learning ratio as applied in Leabra. The Leabra algorithm is a natural fit for semi-supervised paradigms, and the interaction between these two variables is particularly important. The changes wrought in the brain by contact with environmental feedback that is inconstant, incomplete, or absent (Elwin et al., 2007; Li, Farkas, & MacWhinney, 2004; Love, Medin, & Gureckis, 2004; Robare, 2004) will depend on the parameters in the brain that determine changes in synaptic strength.

For example, holding constant kWTA at 0.1 and learning rate at 0.01, for a task learning:model learning ratio of 19:1 (that is, 5% of the weight change comes from Hebbian learning), adding unlabeled data to the training set for a 2:1 ratio of labeled: unlabeled data lowers learning performance, whereas a 3:4 ratio of labeled: unlabeled data improves it. However, when 7.5% of the weight change is from Hebbian learning (a task:model ratio of 12.3:1), the 2:1 labeled:unlabeled data ratio causes less of a decrement than at the 19:1 task:model ratio. These data points, gathered over a series of simulations exploring lexical acquisition and category formation in semi-supervised models, demonstrate the urgent need for more systematic observation of this interaction. Even this casual sequence indicates a complex relationship between the contribution of Hebbian learning to weight-change calculation and the amount of labeled versus unlabeled data presented to the network. Thus we see that at higher amounts of Hebbian learning, less unlabeled data are necessary for performance to improve, whereas more unlabeled data will be required in networks that rely less on Hebbian learning. The precise curve of this relationship has yet to be determined, as has the presence of potential ceiling and floor effects.

Acknowledgements

I thank Dr. John Antrobus for his constant assistance and advice regarding my work.

References

Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological) 39*, 1-38.

Elwin, E., Juslin, P., Olsson, H., & Enkvist, T. (2007). Constructivist coding: Learning from selecive feedback. *Psychological Science 18*, 105-110.

Hemmes, N. S., & Eckerman, D. A. (1972). Positive interaction (induction) in multiple variable-interval, differential-reinforcement-of-high-rate schedules. *Journal of the Experimental Analysis of Behavior 17*, 51-57.

Li, P., Farkas, I., & MacWhinney, B. (2004). Early lexical development in a self-organizing neural network. *Neural Networks 17*, 1345-1362.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of category learning. *Psychological Review 111*, 309-332.

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. (1998). Six principles for biologically-based computational models of cortical cognition. *Trends in Cognitive Sciences, 2*, 455-462.

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. (1995). *PDP++*.

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

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variation in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Theory and research*. New York: Appleton-Century-Crofts.

Robare, R. J. (2004). Generalization and discrimination in a semantic network trained with semi-supervised learning. In *Proceedings of the Sixth International Conference on Cognitive Modeling* (pp. 412-413). Mahwah, NJ: Lawrence Erlbaum.

Sejnowski, T. J. (1986). Open questions about computation in cerebral cortex. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group, *Parallel distributed processing, Vol. 2: Psychological and biological models*. Cambridge, MA: MIT Press.

Yarowsky, D. (1995). Unsupervised word sense disambiguation rivaling supervised methods. *Proceedings of the Annual Meeting of the Association for Computational Linguistics* (pp. 189-196).