Learning in a Generative Model with Competitive Combination Is Approximated by (Soft-)Winner-Take-All Networks

Jörg Lücke and Maneesh Sahani

Gatsby Computational Neuroscience Unit, UCL

In recent years, algorithms such as independent components analysis (ICA) [1] and sparse coding (SC) [2] have been used to describe the statistics of the natural environment, and the components extracted by these methods have been linked to sensory neuronal response properties. Stated in the language of probabilistic generative models (see e.g. [3]), ICA and SC describe sensory data as a linear superposition of learnt components. For many types of data, including images, this assumed linear cooperation between generative causes is unrealistic. Alternative, more competitive, generative models have also been proposed: in [4] hidden causes are combined by noisy-or, and in [5] a still more competitive scheme is described. Here, we formulate an extreme case of competition, in which the strongest generative input to an observed node alone determines its value. Thus, where ICA and SC use a sum, we use a $\max$. Such a rule has the property of selecting, for each observed node, a single generative cause to determine that node's value. In the case of image data, combination by occlusion shares this selective property.

Whilst exact maximum-likelihood learning of the parameters of such a model is intractable, we show that efficient approximations to expectation-maximisation (EM) can be found in the case of sparsely active hidden sources. One of these approximations is shown to be equivalent to a neural network model with a generalised soft-max activation function, and a simple Hebbian Δ-rule with divisive weight normalisation. Thus, we show that learning in winner-take-all (WTA) type networks, e.g. [6-9], may be interpreted as approximate maximisation of a data likelihood. In the limit of very sparse input, we recover the classical soft-max function, which is commonly used for clustering. This observation may help to explain how such WTA networks can successfully resolve components as well as determine clusters in data where appropriate [6-9].

Using the benchmark "bars test" [10], we numerically verify the accuracy of the approximate update rules and the corresponding neural network. These experiments show that the generative approach is competitive with results obtained by other methods.

Acknowledgments
This work was funded by the Gatsby Charitable Foundation.

References

[1] A. J. Bell and T. J. Sejnowski. *Vision Research*, 37:3327 – 3338, 1997.
[2] B. A. Olshausen and D. J. Field. *Nature*, 381:607 – 609 , 1996.
[3] P. Dayan and L. F. Abbott. Theoretical Neuroscience. MIT Press, Cambridge, 2001.
[4] E. Saund. *Neural Computation*, 7:51 – 71, 1995.
[5] P. Dayan and R. Zemel. *Neural Computation*, 7:565 – 579, 1995.
[6] J. Lücke and C. von der Malsburg. *Neural Computation*, 16:501 – 533, 2004.
[7] M. W. Spratling and M. H. Johnson. *Neural Computation*, 14:2157 – 2179, 2002.
[8] R. C. O'Reilly. *Neural Computation*, 13:1199 – 1241, 2001.
[9] J. Lücke. *Neural Networks* 17:1377 – 89, 2004.
[10] P. Földiák. *Biological Cybernetics*, 64:165 – 170, 1990.