

Finding the optimal sparse, overcomplete model for natural images by model selection

Richard E. Turner, Pietro Berkes, and Maneesh Sahani

Gatsby Computational Neuroscience Unit, UCL

The principles that underlie the structure of receptive fields in the primary visual cortex are not well understood. One theory is that they emerge from information processing constraints, and that two basic principles in particular play a key role. The first principle is that of *sparsity*. Both neural firing rates and visual statistics are sparsely distributed, and sparse models for images have been successful in reproducing some of the characteristics of simple cell receptive fields (RFs) in V1 [1,2]. The second principle is *overcompleteness*. The number of neurons in V1 is 100–300 times larger than the number of neurons in the LGN. It has often been assumed that sparse, overcomplete codes might lend some computational advantage in the processing of visual information [3,4]. The goal of this work is to investigate this claim.

Many different sparse-overcomplete models for visual processing have been proposed. These have largely been evaluated on the basis of their correspondance with neural properties (RF frequency, orientation, and aspect ratio after learning), on their effectiveness in denoising natural images, or on the efficiency with which they can be used to encode natural images. However, only rarely have the questions about the degree of sparsity, the form of the sparsity, as well as the overcompleteness level, been addressed.

Here we formalise such questions of optimality in the context of Bayesian model selection, treating both the degree of sparsity and the extent of overcompleteness as parameters within a probabilistic model, that must be learnt from natural image data. In the Bayesian framework, models are compared based on their marginal likelihoods, a measure which reflects their ability to fit the data, but also incorporates a Bayesian equivalent of Occam's razor by automatically penalizing models with more parameters than are supported by the data. We compare different sparse coding models and show that the optimal model seems to be indeed very sparse but, perhaps surprisingly, only modestly overcomplete. Thus, according to our results, linear sparse coding models are not sufficient to explain the presence of an overcomplete code in the primary visual cortex.

References

- [1] B.A. Olshausen and D.J. Field (1996) Emergence of Simple-Cell Receptive Field Properties by Learning a Sparse Code for Natural Images. *Nature*, 381, 607–609.
- [2] A.J. Bell and T.J. Sejnowski (1997) The 'Independent Components' of natural scenes are edge filters. *Vision Research*, 37, 3327–3338.
- [3] B.A. Olshausen and D.J. Field (1997) Sparse Coding with an Overcomplete Basis Set: A Strategy Employed by V1? *Vision Research*, 37, 3311–3325.
- [4] H. Lee, A. Battle, R. Rajat, and A.Y. Ng (2007) Efficient sparse coding algorithms. In *NIPS* 19.