Chapter 4

Computational neuroscience

Aapo Hyvärinen, Patrik Hoyer, Jarmo Hurri, Mika Inki

4.1 The statistical structure of natural images and visual representation

Our research concentrated on modelling visual perception using statistical models. This work was a direct continuation of what was reported in the previous biennial report.

Our basic approach is to build models of the statistical structure of the typical input that the perceptual system receives, and estimate the parameters of the model from realistic input, such as digital photographs of wild-life scenes, or digital video. We then describe the function of parts of the visual cortex as statistical estimation and inference in such models. Our modelling has been largely based on extensions of independent component analysis and blind source separation methods. The goal is typically to transform a data vector into components that are statistically independent or whose dependency structure is quite simple.

We have developed a number of new models that extend the now well-known results on independent components of natural images:

Non-negative sparse coding This model is specialized to analysis of data that is non-negative. More precisely, the data is usually positive and close to zero, and occasionally gets large positive values [1,2]. The model seems especially suited for analysis of higher-order features that are computed from outputs of lower non-negative features, such as complex cells.

Models of variance dynamics It is well-known that the independent components of natural images are not independent. Our previous work already modelled some of the dependencies that remain after ICA. Yet, no existing model have been able to estimate a full two-layer model of natural images, where the seoncd layer explains some of the dependencies left after the first linear layer. Previous research has only been able to estimate two-layer models when one of the layers has been fixed. We developed a model where two layers can be estimated based on the temporal structure of natural image sequences [4]. The layers correspond to simple and complex cells in the primary visual cortex. See Fig. 4.1 for an illustration of the main kinds of dependencies, and Fig. 4.2 for some dependencies estimated from real data.

Bubble coding We have proposed a unifying framework [6] for several models of the statistical structure of natural image sequences. The framework combines three properties: sparseness, temporal coherence, and energy correlations. It leads to models where the joint activation of the linear filters (simple cells) takes the form of "bubbles", which are regions of activity that are localized both in time and in space, space meaning the cortical surface or a grid on which the features are arranged. The concept of bubbles is closely related to invariant features such as those coded by complex cells; the principle is illustrated in Fig. 4.3.

Double-blind source separation These theoretical developments in biological modelling lead to the development of a new method of blind source separation [5]. The new method separates sources without the need for an explicit parametric model of their dependency structure. This is possible by some general assumptions on the structure of the dependencies: the sources are dependent only through their variances (general activity levels), and the variances of the sources have temporal correlations. The method can be called double-blind because of this additional blind aspect: We do not need to estimate

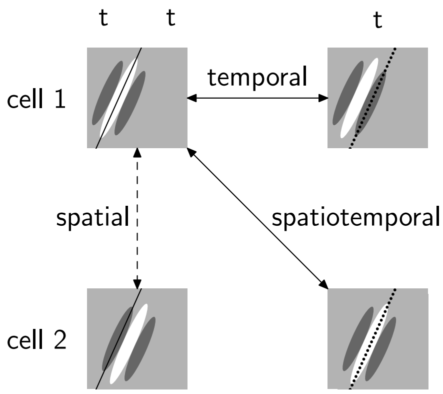

Figure 4.1: Illustration of the different types of dependencies found in natural image sequences. Consider a stimulus that consists of a line segment that moves accross the receptive fields of two linear neurons with receptive fields that have simile location, orientation and frequency. The outputs of a given neuron in two consecutive time steps are dependent. Further, two neurons with similar receptive fields have dependent outputs. Also, the outputs of similar neurons in consecutive time points are dependent.

Figure 4.2: Our two-layer model in [4] was able to estimate linear features and dependencies between the features. Each row shows the features with the highest and the lowest dependencies with respect to the reference feature on the left. Features with high dependencies code for similar orientation and frequeny. Features with low dependencies have very dissimilar parameters.

(or assume) a parametric model of the dependencies, which is in stark contrast to most previous methods.

Conditional and comparative statistics We have further investigated how the statistics are modified by conditioning by the value of one independent component [7]. If the components were really independent, this conditioning should not change anything, but our results show that it does. Comparison of the statistics of natural images with others kinds of images have also been performed [8].

Bayesian inference in the visual system On a more theoretical note, we proposed a model that explains some aspects of the response variability of neurons using the framework of Bayesian inference. It is proposed that the variability reflects a Monte Carlo sampling of the posterior probability distribution of perceptual parameters, given the input stimulus [3].

References

[1] P. O. Hoyer. Non-negative sparse coding. In *Proc. IEEE Workshop on Neural Networks for Signal Processing*, pages 557–565, 2002.

[2] P. O. Hoyer. Modeling receptive fields with non-negative sparse coding. In *Computational Neuroscience: Trends in Research 2003 (Proc. CNS 2002)*, 2003.

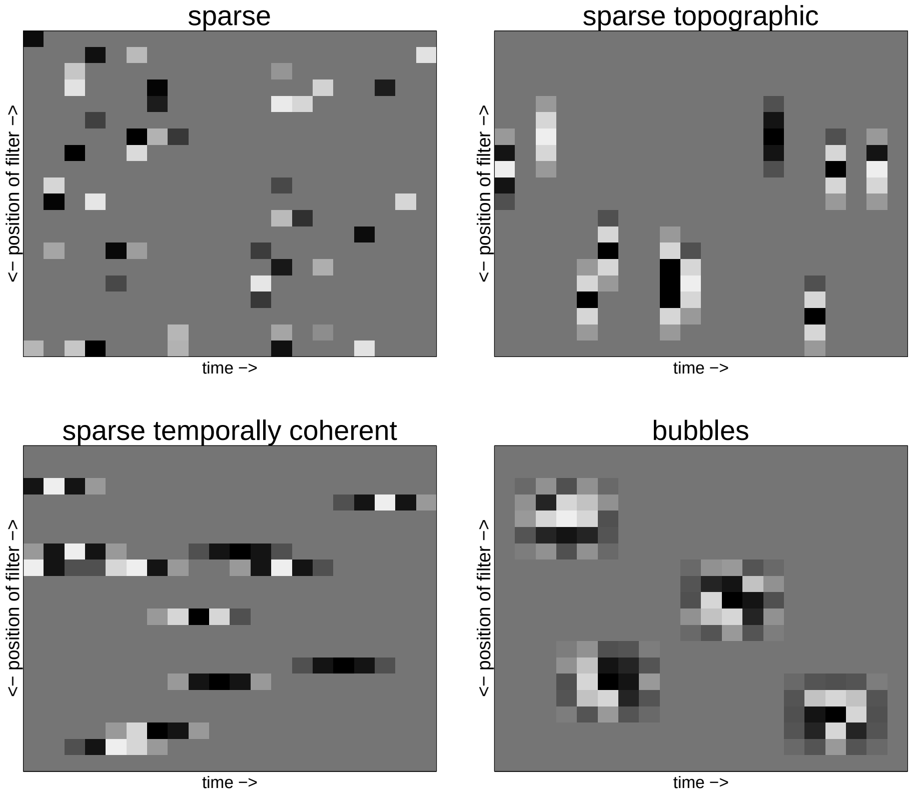

Figure 4.3: An illustration of a bubble representation. The plots show the outputs of filters as a function of time (horizontal axis) and the position of the filter on the topographic grid (vertical axis). Each pixel is the output of one unit at a given time point, gray being zero, white and black meaning positive and negative outputs. For simplicity, the topography is here one-dimensional. In the basic sparse representation, the filters are independent. In the topographic representation, the activations of the filters are also spatially grouped. In the representation that has temporal coherence, they are temporally grouped. The bubble representation combines all these aspects, leading to spatiotemporal activity bubbles. Note that the two latter types of representation require that the data has a temporal structure, unlike the two former ones.

[3] P. O. Hoyer and A. Hyvärinen. Interpreting neural response variablity as Monte Carlo sampling of the posterior. In *Advances in Neural Information Processing Systems*, volume 15, pages 277–284. MIT Press, 2003.

[4] J. Hurri and A. Hyvärinen. Temporal and spatiotemporal coherence in simple-cell responses: A generative model of natural image sequences. *Network: Computation in Neural Systems*, 14(3):527–551, 2003.

[5] A. Hyvärinen and J. Hurri. Blind separation of sources that have spatiotemporal variance dependencies. *Signal Processing*. In press.

[6] A. Hyvärinen, J. Hurri, and J. Väyrynen. Bubbles: A unifying framework for low-level statistical properties of natural image sequences. *J. of the Optical Society of America A*, 20(7):1237–1252, 2003.

[7] M. Inki. Examining the dependencies between ica features of image data. In *Proc. ICANN/ICONIP2003*, Istanbul, Turkey, 2003.

[8] M. Inki. ICA features of image data in one, two and three dimensions. In *Proc. Int. Symposium on Independent Component Analysis and Blind Signal Separation (ICA2003)*, pages 861–866, Nara, Japan, 2003.