Slow feature analysis yields a rich repertoire of complex cell properties

Pietro Berkes and Laurenz Wiskott

Institute for Theoretical Biology, Humboldt University Berlin

p.berkes@biologie.hu-berlin.de, http://itb.biologie.hu-berlin.de/∼berkes

1 The slowness principle

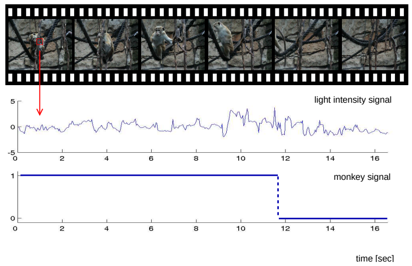

In this work, we investigate **slowness** as a coding principle for the primary visual cortex (V1):

- The input signals to the cortex originate from the sensory cells by **raw, local measurements of the environment**.
- Such measurements are extremely sensitive to small changes in the state of both the environment and the observer, and **vary** thus **on a timescale faster than that of the environment itself**.
- The **slowness principle** assumes that **the cortex extracts slow signals out of its fast varying input in order to reconstruct information about the environment**.
- Slow features are likely to reflect the properties of the environment and are in addition **invariant** or at least robust **to frequent transformations** of the sensory input.

2 Slow Feature Analysis

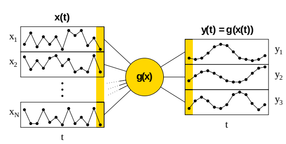

The problem of extracting slow features from time sequences can be formally stated as follows [5]:

- Given an **input signal** $\mathbf{x}(t)$, find an **input-output function** $\mathbf{g}(\mathbf{x})$ (in the simulations presented here a polynomial of degree 2).
- The function generates the **output signal** $\mathbf{y}(t) = \mathbf{g}(\mathbf{x}(t))$.
- The output signal should vary slowly, i.e. **minimize** $\langle \dot{y}_j^2 \rangle$.
- The output signal should carry much information, i.e. $\langle y_j \rangle = 0$, $\langle y_j^2 \rangle = 1$, and $\langle y_{j'} y_j \rangle = 0 \quad \forall j' < j$.

This optimization problem can be solved with **slow feature analysis (SFA)** [5], an unsupervised algorithm based on an eigenvector approach.

3 Input data

Our data source consisted of 36 gray-valued natural images. We constructed image sequences by moving a window over the images by translation, rotation, and zoom and subsequently rescaling the frames to a standard size of 16×16 pixels. To include temporal information, the input vectors to SFA were formed by the pixel intensities of two consecutive frames at times t and $t + \mathrm{d}t$.

4 Results - Physiological-like experiments

SFA learns a set of units that consist in second degree polynomials that applied to our input visual stimuli have the most slowly varying output. The units are ordered by slowness (the first unit being the slowest) and their outputs are mutually uncorrelated. **The functions g_j can be interpreted as non-linear spatio-temporal receptive fields of neurons in V1 and tested with input stimuli such as linear sinus gratings much like in neurophysiological experiments.**

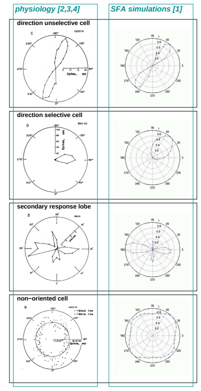

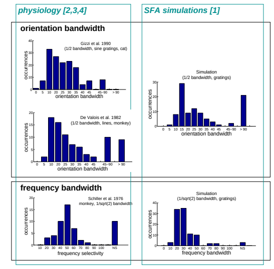

5 Results - Optimal stimuli

We can also analytically compute for each unit the **optimal excitatory stimulus** $\mathbf{S}^+$ and the **optimal inhibitory stimulus** $\mathbf{S}^-$, i.e. the input that elicits the strongest and the weakest output from the unit, respectively. This is in analogy to the physiological practice of characterizing a neuron by the stimulus to which the neuron responds best.

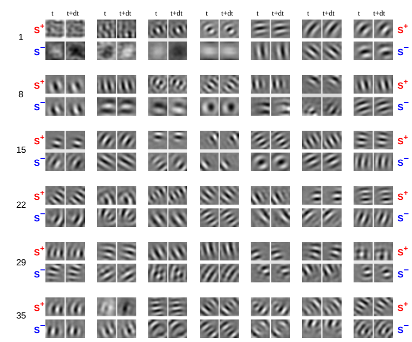

6 Results - Response images

We use **test images** in order to study the activity of the units to a large spectrum of stimuli (e.g. in the example on the left, to all orientations and a wide range of frequencies).

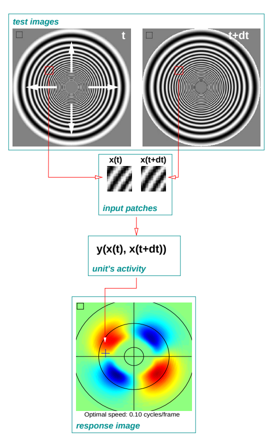

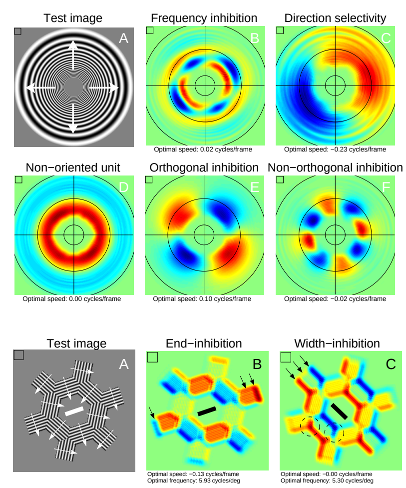

7 Conclusions

We have shown that slowness leads to a great variety of complex cell properties found also in physiological experiments, namely **edge detection, phase-shift invariance, active inhibition, non-orthogonal inhibition, direction selectivity, end-inhibition and side-inhibition**. Our results demonstrate that such a rich repertoire of receptive field properties can be accounted for by a single unsupervised learning principle.

Additional material, papers and other informations are available at
`http://itb.biologie.hu-berlin.de/~berkes`

References

[1] Pietro Berkes and Laurenz Wiskott. Slow feature analysis yields a rich repertoire of complex-cell properties. Cognitive Sciences EPrint Archive (CogPrint) 2804, http://cogprints.ecs.soton.ac.uk/archive/00002804/, 2003.

[2] R.L. De Valois, E.W. Yund, and N. Hepler. The orientation and direction selectivity of cells in macaque visual cortex. *Vision Res.*, 22(5):531–44, 1982.

[3] Martin S. Gizzi, Ephraim Katz, Robert A. Schumer, and J. Anthony Movshon. Selectivity for orientation and direction of motion of single neurons in cat striate and extrastriate visual cortex. *Journal of Neurophysiology*, 63(6):1529–1543, June 1990.

[4] P.H. Schiller, B.L. Finlay, and S.F. Volman. Quantitative studies of single-cell properties in monkey striate cortex. II. Orientation specifi city and ocular dominance. *J. Neurophysiol.*, 39(6):1320–1333, 1976.

[5] Laurenz Wiskott and Terrence Sejnowski. Slow feature analysis: Unsupervised learning of invariances. *Neural Computation*, 14(4):715–770, 2002.