WHICH SPARSITY PROBLEM DOES THE BRAIN SOLVE?

Ilias Rentzeperis [1*], *Luca Calatroni* [2], *Laurent Perrinet* [3], *Dario Prandi* [1]

[1] Universite Paris-Saclay, CNRS, CentraleSupelec, Laboratoire des Signaux et Systemes
[2] UCA, CNRS, INRIA, Laboratoire d'Informatique, Signaux et Systemes de Sophia Antipolis
[3] Aix Marseille Univ, CNRS, INT, Inst Neurosci Timone, Marseille, France
[*] ilias.rentzeperis@centralesupelec.fr

Experimental evidence suggests that activity in sensory cortices is sparse in that only few neurons, out of a large pool that could respond to sensed stimuli, are active at a time. Generative learning models that aim to replicate sensory systems could deviate from sparse activity patterns when representing noisy signals. We ask: are there biologically plausible implementations that will maintain sparse activations for different levels of noise while representing the underlying signal?

A family of generative algorithms modelling sensory systems represent a stimulus as a linear sum of an overcomplete dictionary of vectors with their corresponding coefficients taking the role of activations. Olshausen and Field [1] showed that a learning algorithm that is set to reconstruct natural images with sparse activations develops vectors that have three properties, also found in the receptive fields of neurons in the primary visual cortex, i.e. they are localized, bandpass, and oriented (Figure 1).

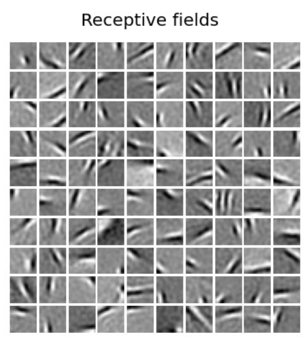

Figure 1: Feature vectors generated with sparse codes are similar to the receptive fields of V1 cells.

The properties emerge by solving an optimization problem which aims to minimize the square error between the actual image and the reconstructed one while keeping the sum of the activities (L1 norm) as small as possible (with the relative weight of the two tasks being controlled by a trade-off parameter, a positive scalar). Traditionally for this kind of problems the summation of activations is minimized (a convex optimization problem); but ideally, we would want to minimize the number of nonzero activations. What hinders us from that is that in the latter case the problem becomes nonconvex. Moreover, a fixed trade off parameter cannot accomodate the same sparsity for a wide range of noise levels.To remedy these effects we use an adaptable trade-off parameter and an Lp norm with p between 0 and 1. We discuss possible connections of these tools with brain mechanisms.

References

1. Olshausen, B. A. et al. Nature 381 (1996).