Population coding of sensory stimuli through latent variables

Martino Sorbaro[1,2], Michael E. Rule[1], Gerrit Hilgen[3], Evelyne Sernagor[3], and Matthias H. Hennig[1].
[1]ANC, School of Informatics, U. of Edinburgh, UK; [2]KTH, Stockholm, Sweden; [3]Institute of Neuroscience, U. of Newcastle, UK

Statistical criticality and latent-variable encoding

- Signs of "statistical criticality" have been observed in retinal activity in the form of Zipf laws and diverging specific heats [4].
- It has been proposed that inference is more likely to find critical models [3].
- Zipf laws are very widespread in nature, and have also been attributed to the presence of broadly distributed latent factors [1].
- It is not known whether statistical criticality has a functional or biological relevance, or if it is connected to dynamics.

We compare encoding of sensory stimuli in retinal ganglion cells and restricted Boltzmann machines (RBMs), and observe signatures of criticality. We study how to interpret signs of statistical criticality in latent-variable encoding.

Visual encoding in Restricted Boltzmann Machines (RBMs)

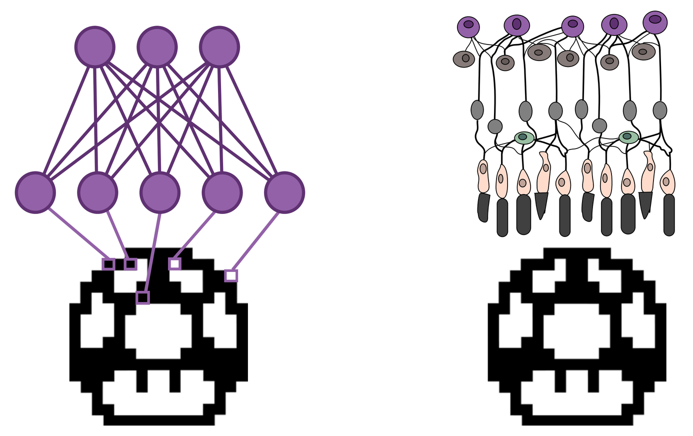

Figure 1: Analogy between a Restricted Boltzmann Machine (RBM), encoding the probability distribution of images into hidden units, and the retina, encoding an image to retinal ganglion cells (RGC) activity.

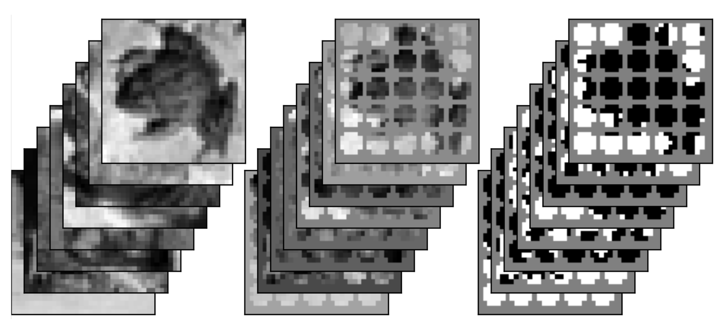

Figure 2: Processing of the CIFAR dataset into small binary patches, on which RBMs are trained.

Zipf laws in retinas and models

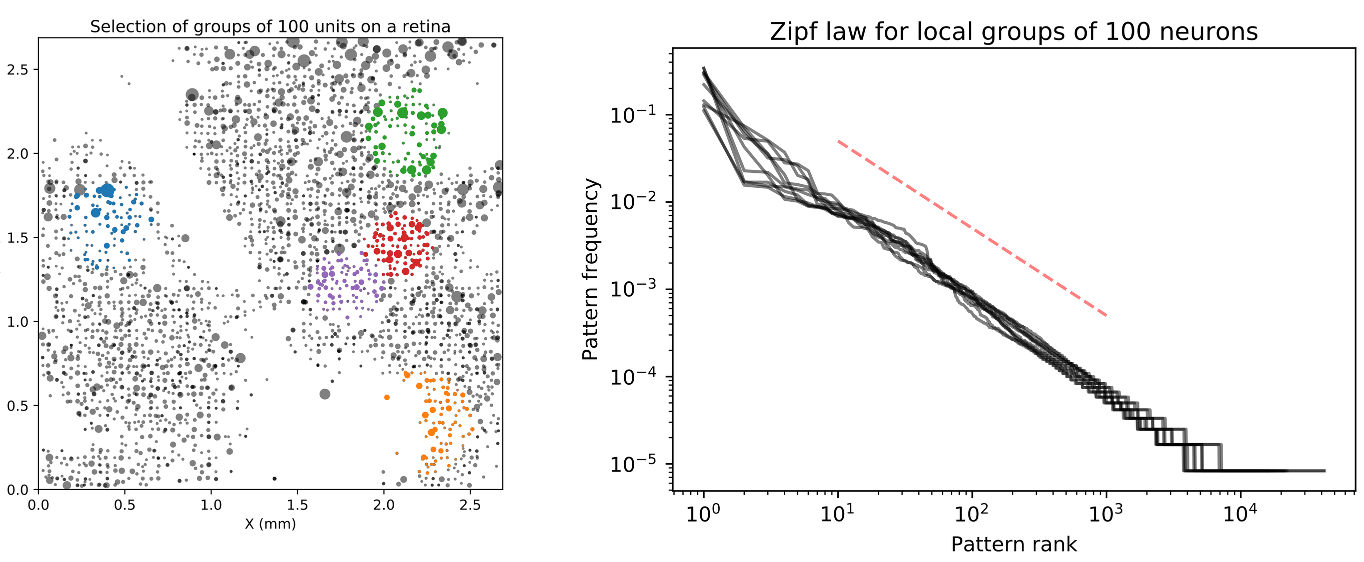

Figure 3: Left: The activity of a mouse retina is recorded on a multi-electrode array. Spikes are sorted using a custom algorithm [2] and local patches are used for analysis (coloured dots are examples). Right: Zipf laws are observed on any group of neurons in a retina stimulated by a random checkerboard. Each curve is a Zipf plot for 100 nearby neurons (activity binned in 10 ms bins).

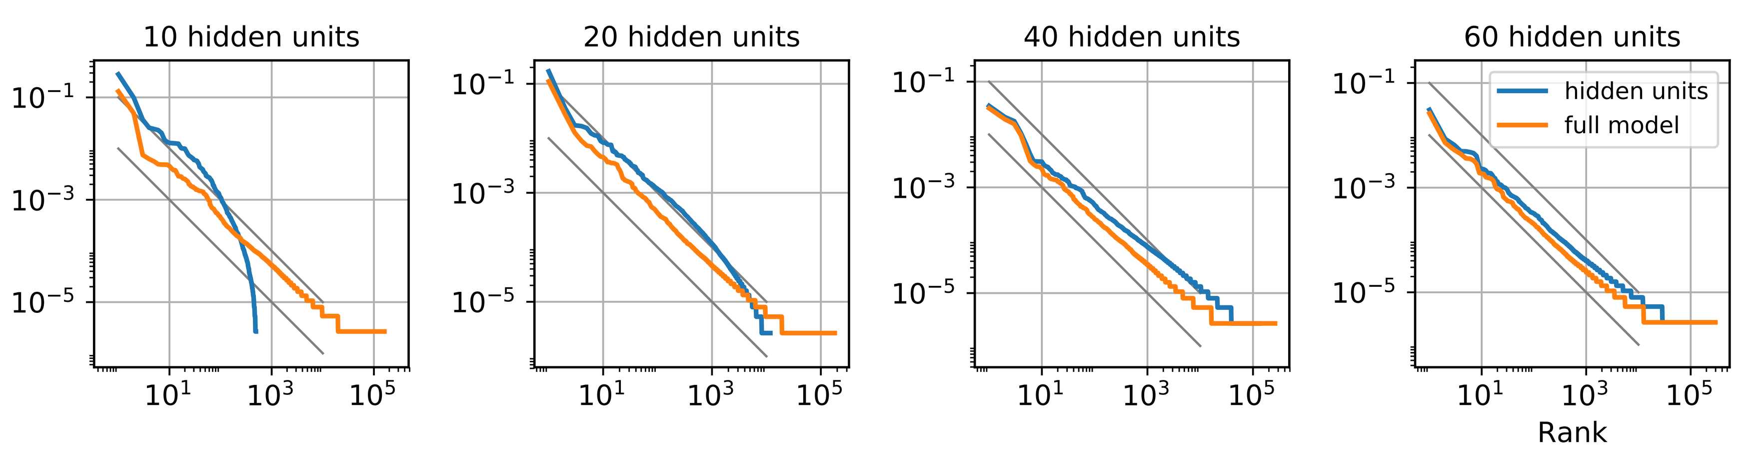

Figure 4: Samples from RBMs trained on CIFAR exhibit Zipf-like statistics as soon as they are complex enough to faithfully reproduce the data.

Fitted models lie close to the critical temperature

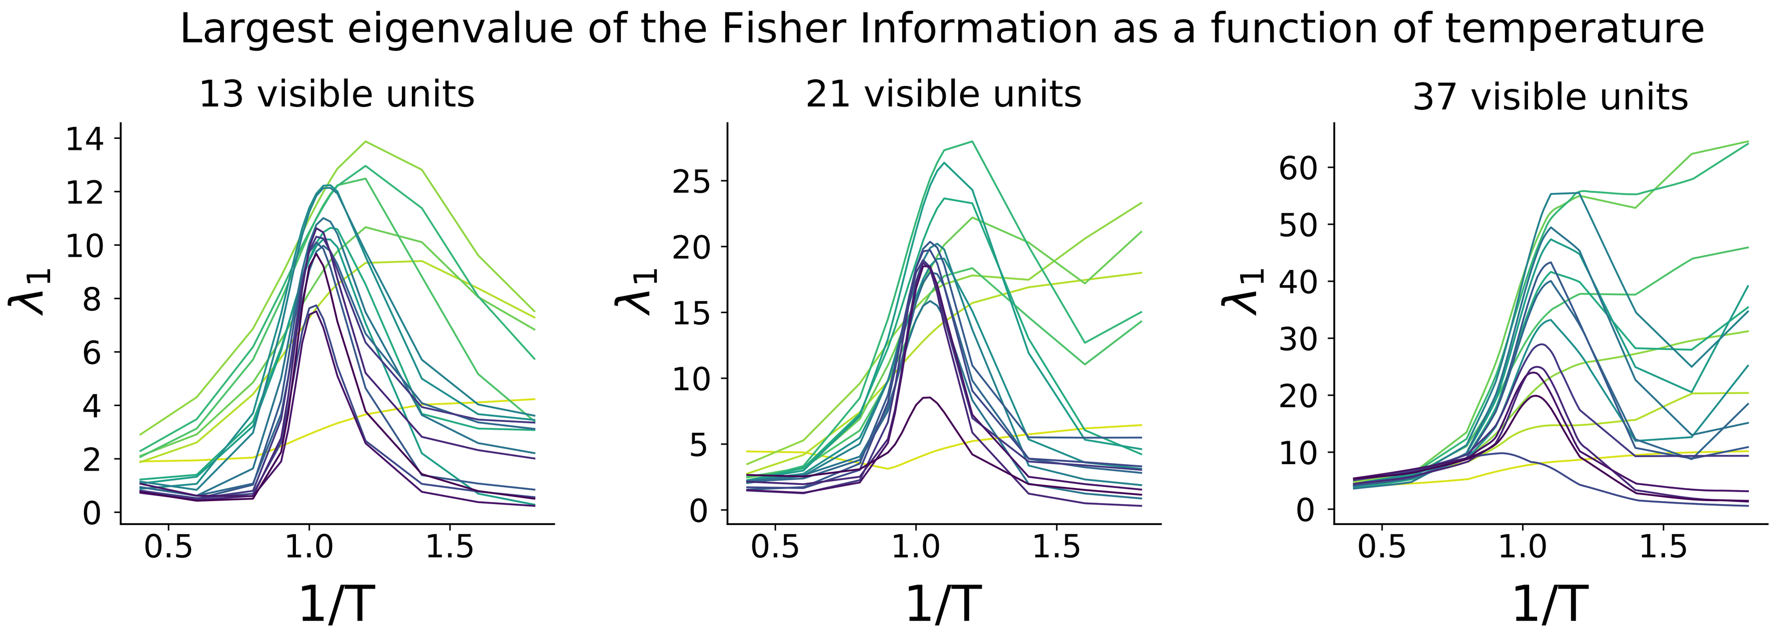

Figure 5: Largest eigenvalue of the Fisher information matrix as a function of temperature. $T = 1$ indicates the original temperature of the fit. The peak is a signature of criticality in the statistical sense.

Sufficiently large models learn sparse, decorrelated representations

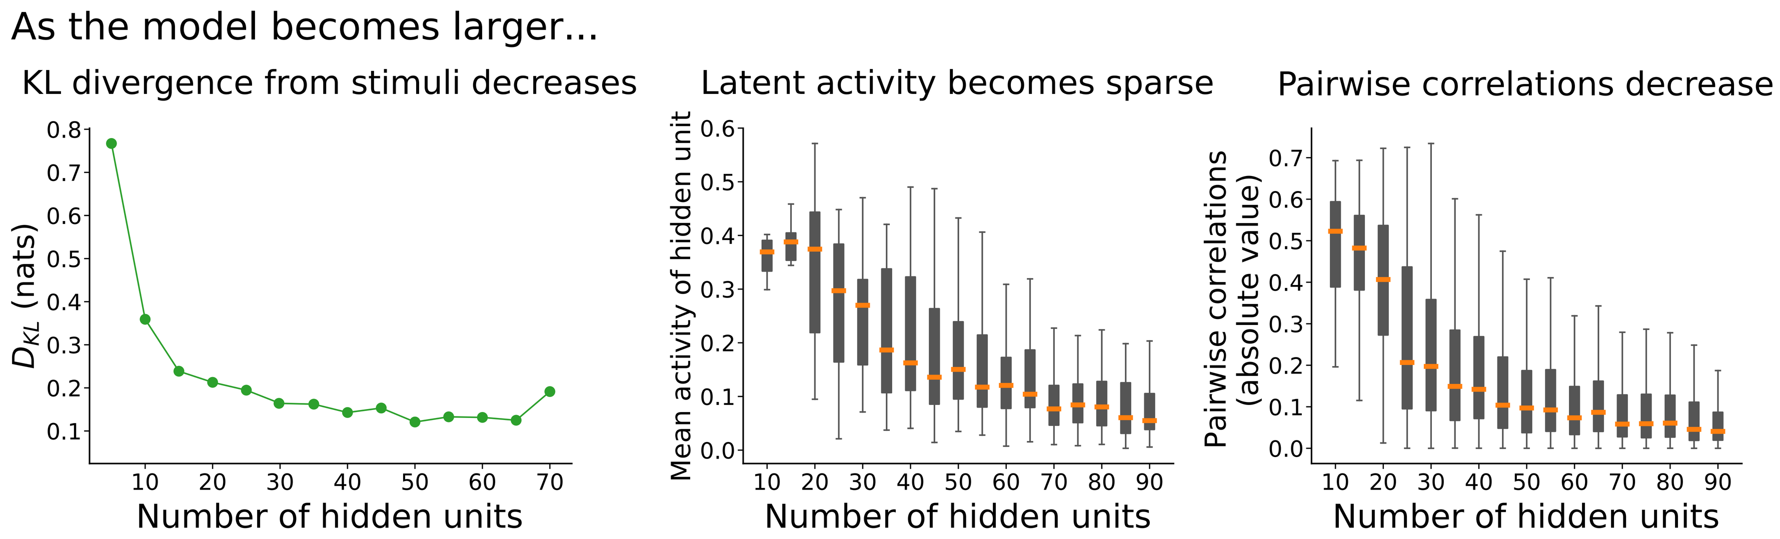

Figure 6: Left: The RBM learns a factorized generative model, and the distance between the generative distribution and the data reflects the quality of the model fit. Here, this distance is summarized in terms of the KL-divergence between the model distribution and the training data. After a certain hidden layer size, we achieve asymptotically good fits to the data. Middle: As the model gets larger and the quality of the fit increases, hidden units are less active, indicating sparse coding. Right: Correlations are attenuated in larger models, indicating successful learning of a factorized latent-variable model.

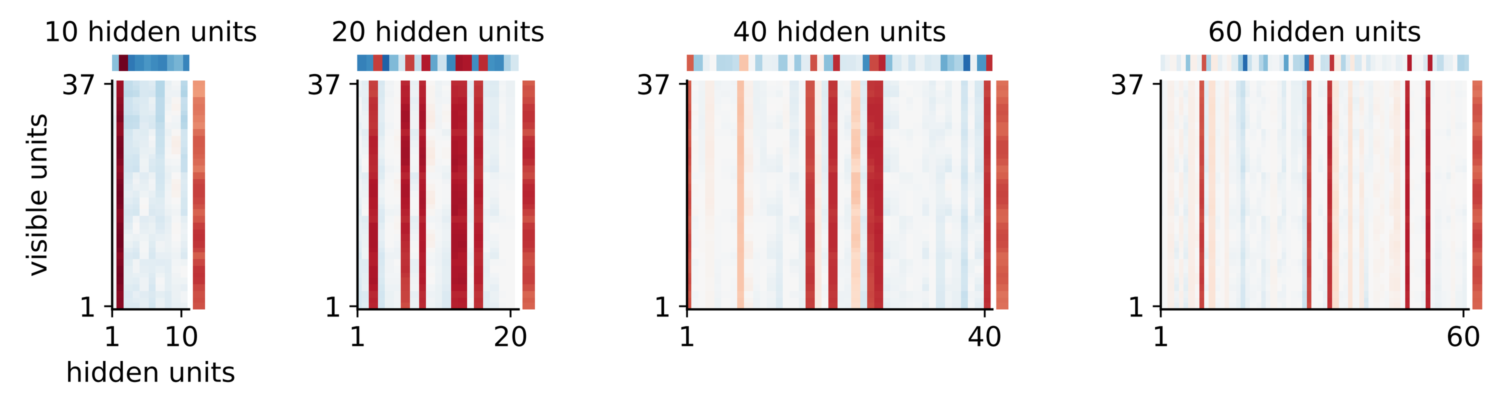

Figure 7: The first eigenvector of the Fisher information for various RBMs, which shows the most sensitive parameter direction. The addition of further hidden units after a good fit is reached doesn't add any sensitive parameters.

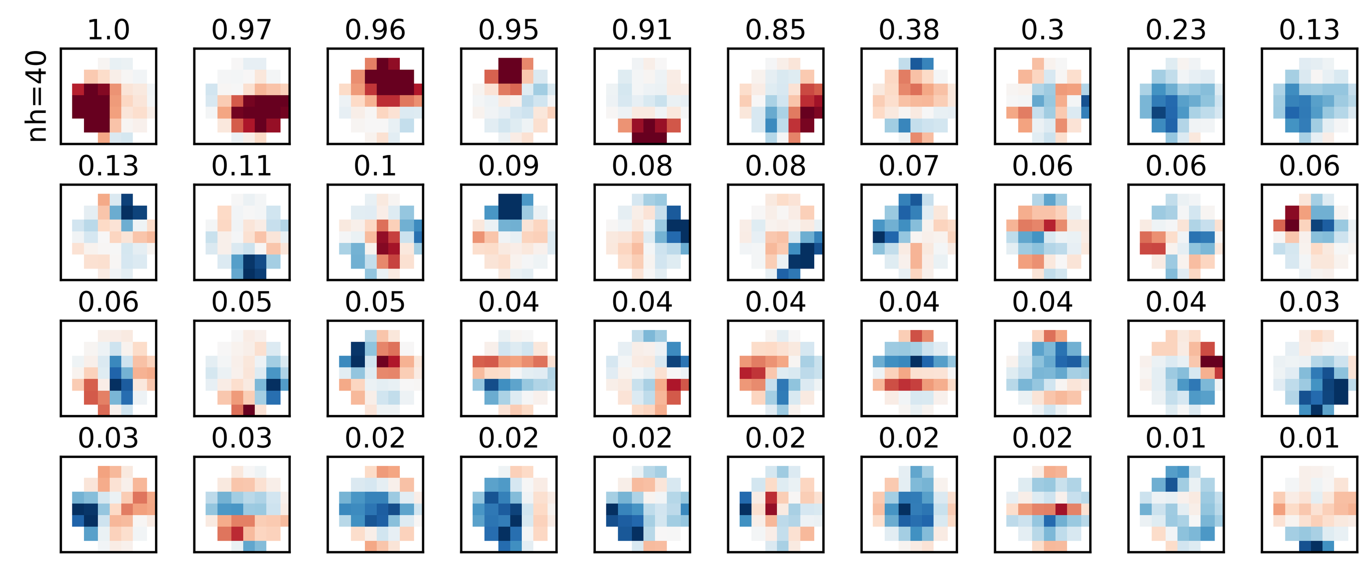

Figure 8: "Receptive fields" for the 40 hidden units of an RBM, sorted by their importance, as measured by Fisher information. RGC-like receptive fields correspond to more important units, and appear in smaller RBMs.

RBMs as a model for latent-variable encodings

- Optimal *latent-variable encoding* of visual stimuli seems to consistently yield models *near statistical criticality*. Poor fits (too few hidden units, under-fitting) do not exhibit this property.
- Critical RBMs mimic the retina in Zipf laws, sparsity, and decorrelation.
- Above the optimal model size, extra units are weakly constrained as measured by Fisher information. Receptive fields of excess units are less retina-like.

Questions and controversy

- *Is statistical criticality a general feature* of factorized latent variable models?
- Is criticality in the retina expected based *simply on optimal encoding*?

References

- L. Aitchison, N. Corradi, and P. E. Latham.
Zipfs law arises naturally when there are underlying, unobserved variables.
PLOS Computational Biology, 12(12):1–32, 12 2016.
- G. Hilgen, M. Sorbaro, S. Pirmoradian, J.-O. Muthmann, I. E. Kepiro, S. Ullo, C. J. Ramirez, A. P. Encinas, A. Maccione, L. Berdondini, et al.
Unsupervised spike sorting for large-scale, high-density multielectrode arrays.
Cell reports, 18(10):2521–2532, 2017.
- I. Mastromatteo and M. Marsili.
On the criticality of inferred models.
J. of Statistical Mechanics: Theory and Experiment, 2011(10):P10012, 2011.
- G. Tkačik, T. Mora, O. Marre, D. Amodei, S. E. Palmer, M. J. Berry, and W. Bialek.
Thermodynamics and signatures of criticality in a network of neurons.
Proceedings of the National Academy of Sciences, 112(37):11508–11513, 2015.

Acknowledgements

M. S. S. is supported by the Erasmus Mundus Doctoral Programmme in Neuroinformatics, the EPRSC Doctoral Training Centre in Neuroinformatics, and a Google European Doctoral Fellowship. M. R. is funded by EPSRC grant EP/L027208/1.

Contact email: martino.sorbaro@ed.ac.uk