Signal-to-noise optimization: gaining insight into information processing in neural networks

Benjamín Pascual[1,2] and Salva Ardid [1,3*]

[1] Research Institute for Integrated Management of Coastal Areas, Universitat Politècnica de València

[2] PhD in Design, Manufacture and Management of Industrial Projects, Universitat Politècnica de València

[3] Department of Applied Physics, Universitat Politècnica de València

[*] Corresponding Author E-mail: sardid@upv.es

Abstract

The current mainstream view in neuroscience and machine learning is that neural networks compress representations into low-dimensional manifolds [Op de Beeck et al., 2001, Gao and Ganguli, 2015, Gallego et al., 2017, Ansuini et al., 2019, Recanatesi et al., 2019]. A recent study challenges this view, by arguing that neural networks benefit from high-dimensional representations [Elmoznino and Bonner, 2022].

In contrast to these positions, we argue that learning in deep neural networks optimizes signal-to-noise processing. According to this view, neural networks may benefit from feature compression and expansion to (i) increase signal processing and (ii) diminish noise, while (iii) mapping input representations into outputs categories. We speculate also that nonlinearities (e.g., in activation functions) facilitate this process.

A causal relationship must exist between the signal-to-noise ratio (SNR) and the behavioral performance of a network (e.g., in terms of classification accuracy) if SNRs are indeed optimized through learning. To test this hypothesis, we first adapted the SNR presented in [Sorscher et al., 2022], so it can be applied to neural representations associated with predictions of unseen data. We then computed the SNR to quantify the separability between category-based manifolds through different layers of neural processing, and tested the SNR with and without input noisy fluctuations, as well as with linear and nonlinear transformations (Linear, ReLU and Sigmoid).

Our results show that: (i) increasing noise fluctuations diminishes the SNR and the accuracy, whereas the dimensionality increases, and (ii) the accuracy increases with the SNR with nonlinear functions and with additional hidden layers. Undergoing analyses aim to test whether an inflection point exists that optimizes the SNR through changes in dimensionality. We speculate a potential benefit of an initial signal-aligned feature expansion followed by dimensionality reduction in latter processing stages close to the output representation, based on output categories being typically low-dimensional.

Keywords— deep neural networks, neural manifolds, high vs. low dimensionality, signal-to-noise processing

References

H. Op de Beeck, J. Wagemans, and R. Vogels. Inferotemporal neurons represent low-dimensional configurations of parameterized shapes. *Nat Neurosci*, 4(12):1244–1252, Dec 2001.

P. Gao and S. Ganguli. On simplicity and complexity in the brave new world of large-scale neuroscience. *Curr Opin Neurobiol*, 32:148–155, Jun 2015.

J. A. Gallego, M. G. Perich, L. E. Miller, and S. A. Solla. Neural Manifolds for the Control of Movement. *Neuron*, 94(5):978–984, Jun 2017.

A. Ansuini, A. Laio, J. H. Macke, and D. Zoccolan. Intrinsic dimension of data representations in deep neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.

S. Recanatesi, M. Farrell, M. Advani, T. Moore, G. Lajoie, and E. Shea-Brown. Dimensionality compression and expansion in deep neural networks. *arXiv preprint arXiv:1906.00443*, 2019.

E. Elmoznino and M. F. Bonner. High-performing neural network models of visual cortex benefit from high latent dimensionality. *bioRxiv*, pages 2022–07, 2022.

B. Sorscher, S. Ganguli, and H. Sompolinsky. Neural representational geometry underlies few-shot concept learning. *Proc Natl Acad Sci U S A*, 119(43):e2200800119, Oct 2022.