



Comment

It's a small dimensional world after all
Comment on “The unreasonable effectiveness of small neural ensembles in high-dimensional brain” by Alexander N. Gorban et al.

Gabriel Kreiman

Children's Hospital, Harvard Medical School, United States of America

Received 15 March 2019; accepted 25 March 2019

Available online 27 March 2019

Communicated by L. Perlovsky

The curse of dimensionality refers to the exponential growth in volume in high dimensional spaces, making efficient sampling and drawing conclusions from sparse sampling extremely challenging [1,4]. The difficulties of high-dimensional spaces are particularly prevalent in Neuroscience, where they are manifested in at least three different forms: (1) The number of possible stimuli and task conditions is infinite; (2) In large animals, the number of neurons is many orders of magnitude larger than the numbers that can be currently studied experimentally; (3) By and large, most modeling tools to think about neural data are founded on intuitions from small dimensional worlds.

Consider a specific example: the goal is to characterize how information is represented along ventral visual cortex in monkeys, the areas critical for visual object recognition, by presenting images and recording neuronal activity. The number of possible images is infinite. Even restricting the question to small image patches of size 100x100 pixels and 256 shades of gray, there are many more such images than the estimated number of stars in the observable universe. The number remains astronomically large even if we restrict the analysis to patches of images cropped from photographs on the internet to impose naturalistic constraints. Even without considering other variables such as the age, state, and goals of the animal, task demands, and other experimental conditions, characterizing this stimulus space is daunting to say the least. To make matters more complicated, to investigate the neural representation in visual cortex, traditional studies have relied on studying one neuron at a time; state-of-the-art techniques can push these numbers to hundreds and perhaps soon thousands of simultaneously recorded neurons. Yet, the number of neurons in macaque ventral visual cortex is on the order of several hundred million. These challenges have not prevented publication of a large number of studies about correlations and interpretation of the map between neuronal responses and visual stimuli. Are we deluding ourselves in thinking that we might be able to infer and model brain function by scrutinizing the activity of a handful of neurons in response to a small number of stimuli?

In an elegant article, Gorban and colleagues come to the rescue by explaining the *blessing of dimensionality* and how it may be relevant for Neuroscience studies [3]. Loosely speaking, the idea is that high-dimensional spaces might make data interpretation easier by virtue of correlations of the data dimensions that make a low-dimensional representation capture most of the variance. To build a geometrical intuition, the volume of a ball in n -dimensions is largely concentrated near the surface, an effect that grows with the dimensionality n (Fig. 4 in [3]). The authors refer

DOI of original article: <https://doi.org/10.1016/j.plrev.2018.09.005>.

E-mail address: Gabriel.kreiman@tch.harvard.edu.

to the *genuine* dimension of the data, which may reside in a relatively low-dimensional space. Astute dimensionality reduction techniques can be applied to apparently high dimensional data to extract those genuine dimensions and make complex problems solvable by simple elements like neurons. The authors propose that perception is based on two subsystems. A first step is involved in extracting adequate features from the data. In the machine learning community, such feature extraction ideas are often studied in the context of dimensionality reduction techniques, compressed sensing, and autoencoders. In the Neuroscience community, one could think about such feature extraction ideas in the context of the cascade of computations along the ventral visual cortex transforming the pixel-like representation in the retina into the high-level object representation in inferior temporal cortex. The second step is implemented by simple elements that can extract and represent abstract information in a sparse manner.

The ideas put forward by Gorban and colleagues are related to a recent theoretical study that aimed to examine the connection between task complexity and the dimensionality of dynamic neural representations [2]. Task complexity is defined by the volume of the task manifold. Task parameters need to be mapped onto the neural representation whose smoothness is defined by the neural population correlation lengths. Jointly, these two variables constrain how many linear dimensions the neural dynamics can explore. The theory thus predicts that increasing the number of recorded neurons may not lead to a large expansion of the dimensionality in the neural dynamics mapping onto the task. Rather, increasing the task demands will enable us to observe higher dimensional representations. In addition, Gao et al. capitalize on the theory of random projections, a form of dimensionality reduction, to explain why examining the activity of small numbers of neurons might still yield interpretable intuitions about large neural populations.

Over and over in Neuroscience, we find examples where a small number of dimensions can capture a large fraction of the explainable variable in a cognitive task (see for example the summary in Fig. 2 in [2]). Perhaps fortune favors the brave investigators that have courageously adventured into the terra incognita of the nervous system, armed with a meager electrode, and full of hope of deciphering the brain's inner secrets. The theoretical insights introduced by Gorban and colleagues as well as those of Gao and colleagues suggest that it may not be luck after all: complex systems of large numbers of interconnected neurons might actually hide gems of low-dimensional representations amenable to scientific scrutiny.

References

- [1] Bellman RE. *Dynamic programming*. Princeton, NJ: Princeton University Press; 1957.
- [2] Gao P, Trautmann E, Yu BM, Santhanam G, Ryu S, et al. A theory of multineuronal dimensionality, dynamics and measurement. 2018. bioRxiv.
- [3] Gorban A, Makarov V, Tyukin I. The unreasonable effectiveness of small neural ensembles in high-dimensional brain. *Phys Life Rev* 2019;29:55–88. <https://doi.org/10.1016/j.plrev.2018.09.005> [in this issue].
- [4] Vapnik VN. *The nature of statistical learning theory*. New York: Springer; 1995.