doi: 10.1109/ISOEN.2017.7968888

The Sparse Structure of Natural Chemical Environments

Marlene D. Berke, David J. Field, Thomas A. Cleland
Cornell University, Ithaca, NY, 14853, USA
mdb293@cornell.edu, djf3@cornell.edu, tac29@cornell.edu

Abstract—Sparse representations, in which sources are represented by neural ensembles with small average activity ratios, have many desirable properties including ease of learning, the capacity for generalization, and noise resistance. All of these are characteristic of the biological olfactory system, despite activity patterns across receptor arrays often being highly non-sparse. This is possible because the underlying chemosensory data are "signal sparse" – a distinct property, dependent on environmental statistics, that enables sparse coding algorithms to construct representational sparseness in subsequent processing layers. Whereas this can be accomplished by simple rotation in some systems, chemosensory sparse coding requires curvature in the isoresponse surfaces for each neuron owing to the diversity of chemical sources (overcompleteness). We propose that this functional curvature arises from statistical learning processes instantiated in plastic networks of the early olfactory system. Using models of the chemosensory environment, we illustrate the application of sparse coding algorithms to the analysis of chemical signals.

***Index Terms*—**Sparse coding, source separation, curvature, manifold, olfaction

I. SPARSE REPRESENTATIONS

In neural systems, neuronal activity across a network is considered sparse if individual neurons in that assembly are rarely active and only a few neurons are active at a time. The sparseness of neural activity in this sense can be termed *representational sparseness*, a term encompassing both population sparseness (in space) and lifetime sparseness (in time) [1]. In contrast, *signal sparseness* refers to the sparse structure of the underlying environmental data sampled by a sensory system. That is, rather than describing the sparseness of a neural code embedded in the representational medium, signal sparseness describes the sparseness of the information that the code represents. The distinction is essential to an appreciation of the properties of environmental data and the strategies that sensory systems – natural or artificial – must employ to encode them efficiently.

Signal sparseness can be assessed in terms of the state space of a representational system, in which every possible state of a system can be represented by a unique point. For example, activity in a neural assembly can be depicted as a multidimensional state space with as many dimensions as there are neurons, and in which the location in each dimension corresponds to the activation level of that neuron. A system with two binary neurons (A and B) consequently has two state dimensions (one for each neuron) and four possible states (off/off, off/on, on/off, and on/on). The existence of these possible states in a representational system, however, does not mean that they necessarily all will occur; if a stimulus that activates neuron B is never encountered, only two states of this system will ever be visited. Redundancies in representation arise when some states are more probable than others in this way, and signal sparseness emerges from certain patterns of redundancy. Specifically, a signal has sparse structure when there are multiple *directions* in state space that have relatively high probability, such that the probability density function is "spiky" (Figure 1; [2]). In natural environments, sparse structure is represented by manifolds (lower dimensional surfaces) of these high probability regions; however, a signal can be considered sparse whenever (1) samples from the environment can theoretically be represented by a small number of descriptors (e.g., neurons) and (2) the set of active descriptors changes from sample to sample.

Importantly, while the two are related, signal sparseness does not necessarily produce representational sparseness. A population of 10 binary neurons can exhibit 2^10, or 1024, possible states. If only 50 of those states are ever evoked by sensory input, then the underlying sensory signal is sparse. However, no constraints have been placed on how many of those neurons are active at a time, or how frequently they are active; all 50 of these states could be represented by activity patterns in which 8 or more of the 10 neurons are active. In other words, *the representation of a sparse signal is not necessarily sparse*. In contrast, signals need to be sparse in order to be sparsely represented without loss (dropout methods and other "artificial sparseness" techniques are inherently lossy).

II. SPARSE CODING

Signal sparseness is often preserved as a sensory representation propagates across successive neural processing layers; however, this does not necessarily imply that representational sparseness will emerge. *Sparse coding*

978-1-5090-2392-9/17/$31.00 ©2017 IEEE

[2][3] is the process of transforming a sparse signal into a sparse representation via the application of specific algorithms (neural or otherwise). Although primary sensory representations (e.g., across photoreceptors or olfactory sensory neurons) are rarely sparse, subsequent processing – sparse coding – can reveal the sparse structure in the signal.

In idealized circumstances, sparse coding will isolate the sparse sources of the sensory data and associate a descriptor (e.g., a neuron) with each source (e.g., an odor-emitting object of potential interest), such that activity in that neuron indicates the presence of that source. Figure 2A depicts an illustrative case from Field (1994) with two orthogonal sources (hence, two-dimensional). The diagonals represent the distribution of the data in this data set. By aligning the sources with the data (*bottom left*), we produce a representation in which only one neuron responds to each source. That is, activity patterns in neurons (or, more realistically, relatively small ensembles of neurons) will come to selectively represent the odor of Stilton cheese, one's location in space, the face of one's mother, or the concept of Jennifer Aniston [4][5]. Notably, in both vision and audition, applying sparse coding to natural signals results in neurons that exhibit clear similarities to the properties of neurons in the early sensory pathway [3][6]. Additionally, sparse representations are often metabolically efficient, and their direct representation of signal sources may be the most effective way for local networks to produce a compact and meaningful deliverable for propagation across the brain.

Critically, natural systems do not reflect the idealized extreme in which one neuron corresponds to one source; such a mechanism likely would be fragile and impair the representation of similarity and uncertainty. Rather, natural systems sparsen representations to an extent that likely reflects a complex balance between the benefits and limitations of sparse representations and is instantiated by the neural plasticity algorithms underlying postsensory signal tranformations. Interestingly, the sparse coding algorithm itself also does not generate maximally sparse representation in higher dimensional systems; rather, it converges to incompletely sparse states (Figure 2B) depending on the structure of the system and its inputs [7]. The nature of this balance in both natural and artificial systems remains a topic of great interest.

III. Sparse Coding in Chemical Sensing

Because natural olfactory signals are almost certainly sparse, whereas primary olfactory representations are generally not population sparse, it follows that sparse coding techniques can provide important insights into olfactory coding. We here show that, under relatively simple assumptions, naturalistic odor signals have sparse structure. Importantly, signal sparseness in chemosensation does *not* refer to the concept of molecular sparseness [8] – the observation that natural odors comprise only a small subset of all available odorous molecules (as demonstrated by gas chromatography). First, molecules per se are not a directly relevant factor in odor representation [9], but, more importantly for present purposes, any appropriate measure of sparseness depends on a profile in which the natural olfactory environment is sampled from multiple locations and at multiple times. Signal sparseness depends on the statistics of the spatiotemporal odor environment. Accordingly, databases of sensor responses to batteries of chemical stimuli will be invaluable for practical applications of sparse coding to chemosensory device development.

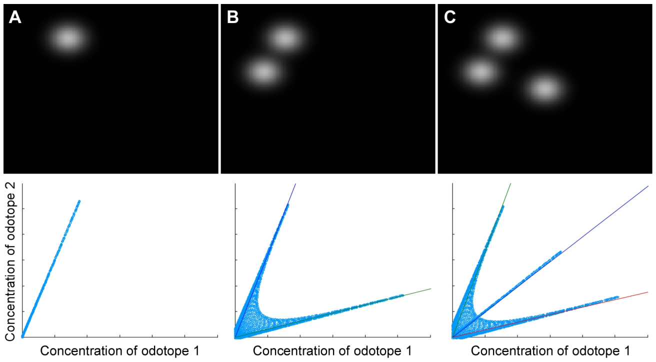

Fig. 1. Top panels. Depictions of odorant sources deployed in a two-dimensional environment. The centers of the white Gaussian blobs depict odor sources, with brightness values indicating odor concentrations at different distances from the source. Each of the three source odorants is a different mixture of the same two components (odotopes [9]). Bottom panels. Plots of 200 evenly distributed points drawn from the corresponding environment and representing the activation levels of receptors selective for odotope 1 (abscissa) and odotope 2 (ordinate). Because signals are sparse, these plots are "spiky". (A) A single source releasing an odorant "A" comprised of 25% odotope 1 and 75% odotope 2. (B) Two sources are depicted, the second comprised of 75% odotope 1 and 25% odotope 2. Rotating the axis via sparse coding would enable a single output unit to selectively respond to odorant A and identify its intensity, and another to do the same for odorant B. (C) Three sources are depicted, the third comprising 50% odotope 1 and 50% odotope 2. As the sources outnumber the number of receptors (the dimensionality of the system), this system is *overcomplete*. Consequently, sparse coding requires curvature of the output units' isoresponse functions in order to generate representational sparseness. "Webbing" between the spikes corresponds to spatial overlap between the odors emitted from different sources.

Figure 1 presents an illustrative model of a simple olfactory environment populated with distinct sources emitting odors constructed from different ratios of the same two components (for illustrative purposes, these components are *odotopes*, activating one sensory receptor each) [9]. Odor concentrations fall off as Gaussian functions of distance from their point sources (Figure 1, *upper panels*.. Sampling each of these three environments generates state spaces in which the distribution of probable states is far from uniform, instead assuming a "spiky" shape in which the spikes correspond to sources. (Whereas Gaussian decay is not a realistic distance function in most natural chemosensory environments, plume models generally do not disrupt this essential spikiness). With such signal-sparse

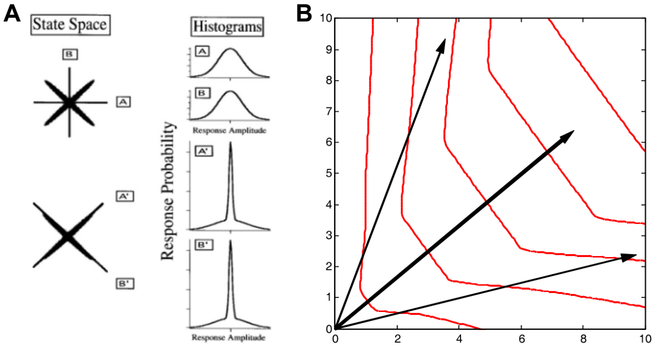

Fig. 2. (A) The basic form of sparse coding provides a way to rotate the representational axes such that they will align with sparse (spiky) signal data, thereby taking advantage of existing signal redundancy [2]. (*Top panels*) Sensory data (X-shaped distribution, comprising two orthogonal sources) are signal sparse, but both primary neuron A and primary neuron B respond to both sources with Gaussian receptive fields. Accordingly, identifying the presence and intensity of a source requires comparison of activity in both neurons. (*Bottom panels*) Sparse coding alters the receptive fields of higher-order sensory neurons (A', B') such that their sensitivities align with the signals emitted by individual sources. This axis rotation increases the kurtosis of the histograms of activation of the vectors, generating representational sparseness (*right panels*). (B) In chemical environments (and most realistic sensory environments), the number of sources exceeds the dimensionality of the sensory system (*overcompleteness*). Consequently, their representations cannot be orthogonal to one another, and sparse coding cannot be achieved by axis rotation alone. To isolate the sources and produce a sparse representation, neurons therefore must curve their isoresponse surfaces. Instead of basic, orthogonal isoresponse contours, the middle vector (boldface) has developed curved isoresponse contours such that it responds more weakly to data that lie close to neighboring vectors. Note that the sparse coding algorithm, which relies on the adaptive increase of kurtosis, has not generated perfect source identifiers (i.e., the isoresponse contours do not entirely avoid intersection with neighboring vectors); see text and [7] for details.

data, sparse coding algorithms are able to align downstream neurons (or ensembles) with particular sources and thereby produce a sparse representation of the data (Figure 1, *lower panels*). Based on model environments with different numbers of sources, different distances between sources, different molecular combinations and different spreads in the concentration of olfactory sources, we find that odor signals that are highly overlapping in quality can nevertheless be sparse, and that that sparse coding algorithms can, under incomplete mixing, find and identify the underlying sources, generating output exhibiting both population and lifetime sparseness. Indeed, whereas primary olfactory representations are rarely sparse, odor representations in the piriform cortex (roughly third-order, following the OSNs and the principal neurons of the olfactory bulb) exhibit highly sparse activity patterns that are reasonably diagnostic for odor sources, at least under laboratory conditions [10]. This attests to the existence of robust sparse coding algorithms embedded in early olfactory sensory processing.

Critically, unlike some simpler systems, olfactory sampling is strongly overcomplete; i.e., there exist many more odor sources than receptors. Consequently, traditional sparse coding algorithms based on axis rotation do not suffice. Curvature of the isoresponse surfaces of output units (e.g., second- or third-order principal neurons or their artificial equivalents) is required in order to generate sparse representations (Figure 2B). Such curvature could be constructed by a plastic network in which the selective allocation of feedback inhibition among second- or third-order principal neuron ensembles can be governed by the complex activation patterns evoked by identifiable odorant sources – precisely the properties hypothesized for the external plexiform layer network of the olfactory bulb [11]. With a sufficiently rich training set, sparse coding will factor strongly into the solution to the "hard problem" in chemosensation: source separation and identification based upon complex, degenerate, highly overlapping, well-mixed, and unpredictably occluded chemosensory inputs.

IV. REFERENCES

[1] B.D. Willmore and D.J. Tolhurst, "Characterizing the sparseness of neural codes," *Network*, Taylor & Francis, Abingdon (UK), vol. 12, pp. 255-270, 2001.

[2] D.J. Field, "What is the Goal of Sensory Coding?" *Neural Computation*, MIT Press, Cambridge, vol. 6, pp. 559-601, 1994.

[3] B.A. Olshausen and D.J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, NPG, London, vol. 381, pp. 607-609, 1996.

[4] R.Q. Quiroga, "Concept cells: the building blocks of declarative memory functions," *Nat Rev Neurosci*, Springer Nature, Berlin, vol. 13, pp. 587-597, 2012.

[5] M.R. Mehta, "From synaptic plasticity to spatial maps and sequence learning," *Hippocampus*, Wiley, Hoboken NJ, vol. 25, pp. 756-762, 2015.

[6] E.C. Smith and M.S. Lewicki, "Efficient auditory coding," *Nature,* NPG, London, vol. 439, pp. 978-982, 2006.

[7] J.R. Golden, K.P. Vilankar, M.C. Wu, and D.J. Field, "Conjectures regarding the nonlinear geometry of visual neurons," *Vision Research*, Elsevier, Amsterdam, vol. 120, pp. 74-92, 2016.

[8] Y. Zhang and T.O. Sharpee, "A robust feedforward model of the olfactory system," *PLoS Comput Biol*, PLoS, San Francisco, vol. 12, pp. e1004850, 2016.

[9] T.A. Cleland, "The construction of olfactory representations," In: *Mechanisms of information processing in the brain: encoding of information in neural populations*, ed. C. Holscher, M. Munk, Cambridge University Press, Cambridge (UK), pp. 247-280, 2008.

[10] J.S. Isaacson, "Odor representations in mammalian cortical circuits," *Curr Opin Neurobiol*, Elsevier, Amsterdam, vol. 20, pp. 328-331, 2010.

[11] A. Borthakur, T.A. Cleland, "A model of experience-dependent odor construction in the olfactory bulb," *Soc. Neurosci. Abstr.*, Society for Neuroscience, Washington DC, vol. 430.24, 2016.