

Understanding and Characterizing Olfactory Responses

BY NIKOLAOS A. LASKARIS,
EFSTRATIOS K. KOSMIDIS,
DEJAN VUCINIC,
AND RYOTA HOMMA

The olfactory bulb (OB) is the first stage of olfactory information processing in the brain. On its way to the cortex, odor information is encoded in spatiotemporal maps of activated loci in the bulb, and these are known as input maps. Using optical recording techniques, experimental neuroscientists can reveal the input maps by applying response-dependent fluorescent dyes to the OB of rodents and imaging the odorant-evoked responses. A hard signal analysis task then emerges (i.e., to handle the resulting sequence of spatial patterns). We suggest that the intricate spatiotemporal patterning of event-related dynamics can be described fully by using a manifold learning approach that merges both the spatial and temporal aspects of the response into a single visualization scheme.

The OB Response Dynamics

Olfaction is one of the most ancient and primal senses for many animals. Humans can recognize and remember about 10,000 odors, while dogs and rodents certainly have larger repertoires. Research to elucidate the underlying neural mechanisms spans across several levels of organization, from the molecular level, where odor-receptor interactions are considered, to the system level, where the complicated cortical networks are investigated. No satisfactory theory exists at present to explain how higher olfactory centers process different odor stimuli to yield a conscious percept of a particular smell. On the other hand, following the discovery of olfactory receptor proteins by Linda Buck and Richard Axel [1], great insights have been gained concerning the initial stages of olfactory information encoding. Olfactory receptor neurons are located on the epithelium in the roof of the two nasal cavities and covered by a mucous layer, where they come in contact with volatile odorant molecules. Each olfactory receptor cell expresses only one type of odorant receptor protein, and each receptor can detect a limited number of odorant substances. The receptor neurons transmit information directly to distinct microdomains, known as glomeruli, in the OB, the first structure in the brain dedicated to olfaction. In mice, there are ~2,000 glomeruli, with each one receiving input from ~10,000 sensory neurons. Despite these large numbers, the network is highly organized. Receptor cells carrying the same type of receptor are scattered throughout the epithelium, but they send their nerve processes to the same glomerulus. This classical scheme of convergence is highly specific in the sense that a given glomerulus receives input only from one type of receptor. The output cells of the bulb, the mitral cells, pick up information from the glomeruli after some pretreatment from an extensive network of local interneurons and transmit it to higher olfactory centers for further processing.

An odorant typically activates many types of sensory neurons, and, as a result, a substantial subset of all glomeruli can be activated even at moderate odorant concentrations. The pattern of activated glomeruli is odor and concentration specific in rodents, and it is referred to as the odor's input map. To experimentally visualize input maps, standard event-related paradigms are adopted from animal electrophysiology and employed in an optical recording setup. The sensory neurons are loaded with a calcium-sensitive dye, and odorant-evoked responses are imaged from their axon terminals that enter the corresponding

*A Manifold Learning
Approach Based
on Optical Recordings*

glomeruli in the OB [2]. Changes in internal calcium concentration reflect neuronal activation and are followed by amplitude changes of the fluorescent dye's emitted light. The signals are acquired in real time using a charge-coupled device camera with high spatial and temporal resolution. The data are collected in the form of brief movies containing the OB response dynamics. This three-dimensional (3-D) information volume is traced for stimulus representation, given that experimental manipulations have compensated for disturbances of physiological origin (such as ongoing respiratory and vascular rhythms or head movements) and appropriate preprocessing

has eliminated other irrelevant signals (such as dye bleaching or slow intrinsic signals). It is a time series of patterned activity and instantaneous input maps, which encode both the type and concentration of odor in a not yet fully understood format.

Since the role of spatial encoding has been emphasized in the related literature, the standard procedure for analyzing an odor-evoked response starts with the visual inspection of a sub-sampled sequence of spatial maps [3], [4]. Loci of maximal activity are then detected and recognized as candidates for well-activated glomeruli [see Figure 1(g)]. Odorant information can be contained in the combination of activated glomeruli,

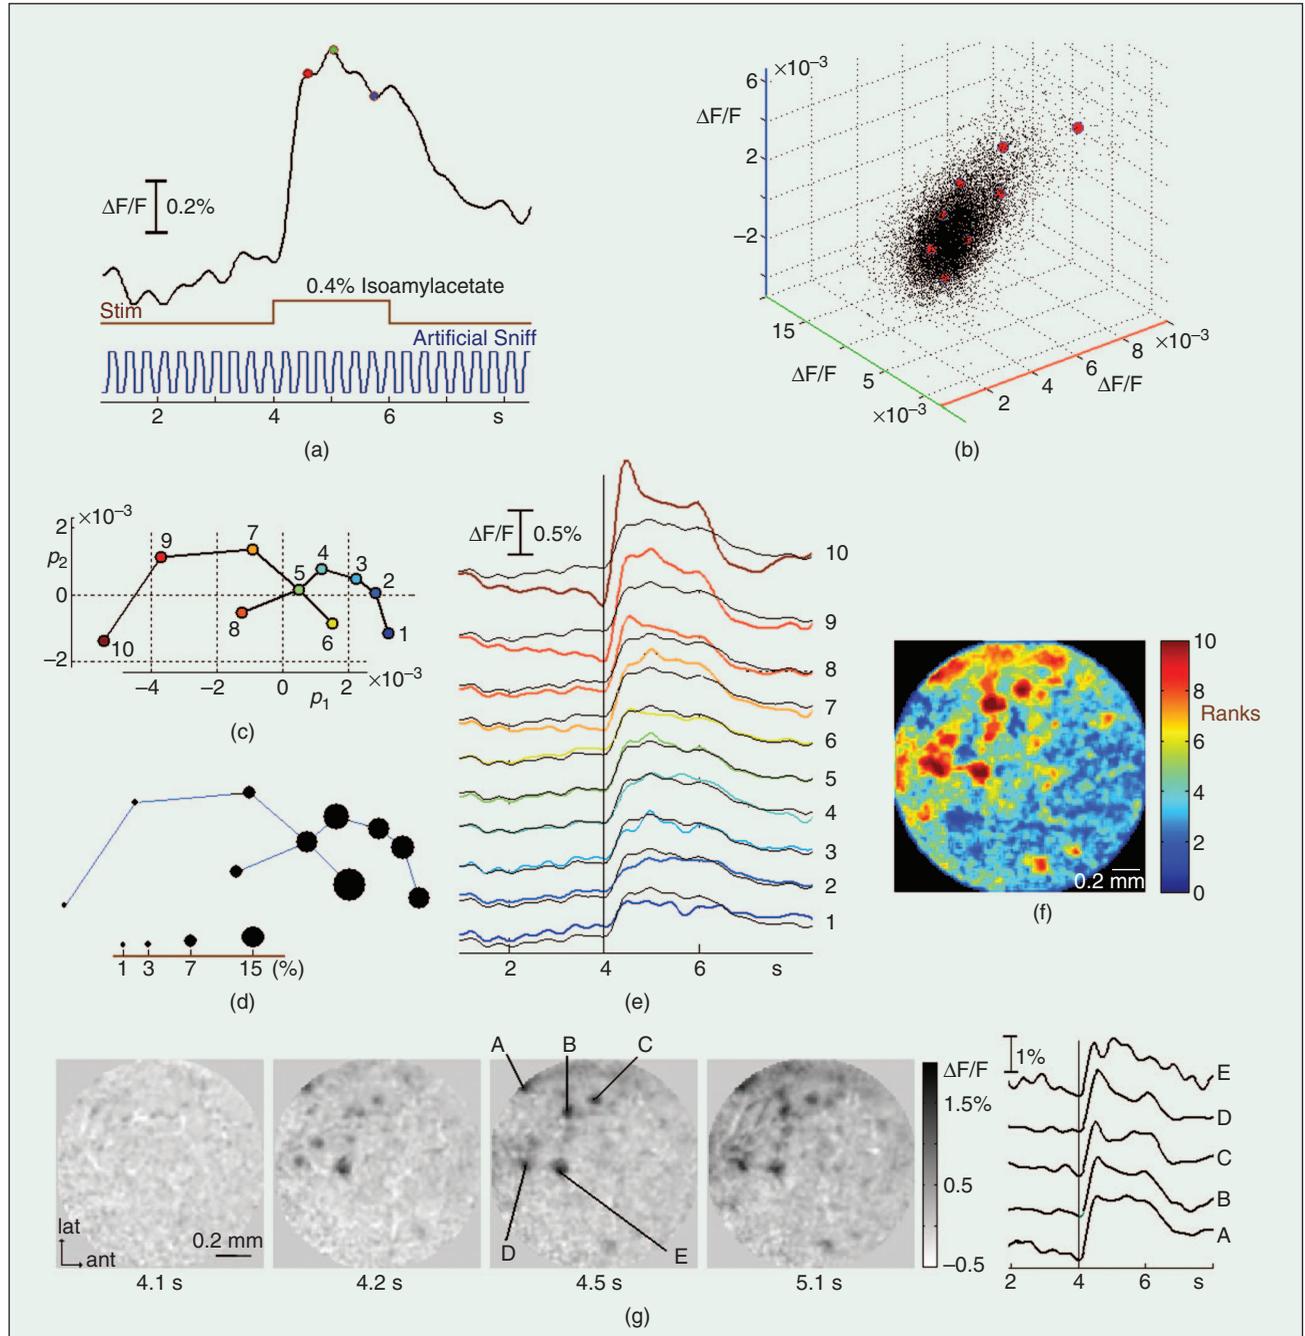


Fig. 1. (a)–(f) The different steps for learning and representing the OB response manifold. The profile in (a) has been computed by averaging temporal patterns across all pixels and is over plotted in (e). (g) Conventional analysis relying merely on relative signal strength. Each spatial map is an average of five consecutive frames around the depicted time. The five time courses have been produced by integrating, using a Gaussian kernel, the activity around the depicted pixels.

their spatial position, as well as their relative response intensity. To explore also the possibility of a coexisting temporal code, sometimes a dedicated response profile is assigned to each of the identified glomeruli and compared with each other's. The whole approach is rather empirical. It demands tight supervision and is, therefore, prone to the user's misinterpretation. The main disadvantage of this approach stems from the inability to handle, in a systematic way, comparisons between different odor representations or recording conditions. The number of different combinations to be contrasted grows rapidly, making the manual analysis insufficient and restricting the whole exploration to a qualitative level. In recent studies, this limitation was overcome using ad hoc procedures (e.g., [3]).

It is the scope of this article to suggest a well-founded methodology [5] for analyzing the OB response dynamics that will enable unraveling the full, multifaceted character of the encoding scheme. The imaged OB region is considered as a spatially extended dynamical system with the tendency to self-organize in response to odor presentation [6]. Using dense spatial sampling, we register these relative changes of activity as a function of time. In this way, a wealth of data is gathered that needs to be summarized, prioritized, and finally presented in an orderly manner so as to be easily understood by domain experts.

Instead of attempting to identify glomeruli from the beginning, we collect all the time courses that correspond to individual pixels, some of which may form spatially connected, coherently activated groups that could be assigned to individual glomeruli. In other words, and in contrast to the traditional approach, we choose to keep the full resolution that our discretization offers (i.e., all pixels and uncollapsed latencies). The ensemble of pixels is embedded in a high-dimensional information space, which is spanned by the latencies, where the pixels take the form of a point swarm, the geometrical shape of which is descriptive of the response variability in an area of the OB. Since the emerged data structure is expected to be low dimensional and possibly forming a nonlinear manifold (i.e., constrained hypersurface), an efficient learning technique is first employed to approximate it by means of prototypes. The geometrical cross-prototypes' relationships are then computed, and a distance-preserving technique is applied to them within the scope of a low-dimensional space, usually a plane. This results in a parsimonious representation of the temporal variations in the OB response by means of a reduced graph, the nodes of which correspond to prototypes, while the deterministic walk on it offers a natural parameterization of response variation. Next, an ordering of the prototypes is imposed using this graph structure, and the ranks are transferred to the whole set of pixels. Finally, these ranks are used as a semantically defined color code to provide an intelligible display of the spatial distribution of the OB response variation. The outlined algorithmic procedure not only serves as an efficient visual data mining scheme but also bridges the gap between the temporal and spatial domain, providing us the possibility to detect more composite encoding scenarios. In addition, it can easily incorporate responses from different recording conditions and, hence, serve as the framework for performing comparisons and even attempting information enhancement.

The article is organized as follows: we briefly describe the experimental setup and then introduce the basic algorithmic elements of the methodology. We describe the two main

variants that emerge by selecting different complexity levels for dimensionality reduction. We exemplify the employed steps using data with high signal to noise ratio (SNR) from experiments in mice. We then demonstrate the potential of the methodology by applying it to data with lower SNR (which is due to experimental conditions) from rats.

Learning the OB Response Manifold

Experimental Setup

Experiments were performed on C57/BL6 mice and Long-Evans rats. Olfactory receptor neurons were loaded with Calcium Green-1 dextran as described in [2]. During surgery and recording, animals were anesthetized with pentobarbital sodium. The bone over one dorsal OB was removed or simply thinned for optical recording access. In mouse experiments, a double tracheotomy was performed to allow the controlled flow of air through the nasal cavity, and the animals breathed freely through the lower canula. In rat experiments, air flow through the nasal cavities occurred spontaneously with breathing. All procedures were approved by the Yale University and the Marine Biological Laboratory animal care committees. Odors were delivered in short pulses (1–2 s) through a custom-made olfactometer described in [4]. In rat experiments, pulses were triggered by the animals' breathing signal. The concentration of the odorants is given as percentage of saturated vapor. We selected odorants that activate the visually accessible dorsal portion of the OB. The dorsal OB was imaged through Wild 10 × 0.4 or 7 × 0.2 numerical aperture objectives mounted onto a Leitz Ortholux II upright microscope. The actual magnifications of the two lenses were 15× and 10×, respectively. In our 128 × 128 pixel charge-coupled device camera, these magnifications resulted in 10 × 10 and 15 × 15 μm/pixel spatial resolution, both adequate for glomerular detection (glomerular diameter is between 50 and 85 μm). Excitation light from a 150-W xenon arc lamp was passed through a 480/50 nm band-pass filter and reflected by a 515-nm dichroic mirror. The fluorescence above 530 nm was recorded with a NeuroCCD-SM or NeuroCCD-SM256 camera using NeuroPlex software. The original sampling frequencies were between 100–200 Hz. In all cases, data were properly resampled at 33 Hz. To enhance the SNR of recordings, multiple (four to eight) trials were collected and subsequently averaged. Here, only averaged data have been considered.

Feature Extraction

Standard preprocessing includes temporal band-pass filtering (within 0.1–2 Hz) and a simple algebraic transformation that associates each pixel (n, m) with a signal, expressing the relative increase in fluorescence $\Delta F/F$ due to odorant presentation

$$x_{n,m}(t) = \frac{f_{n,m}(t) - \frac{1}{(t_{st} - t_{pr} + 1)} \sum_{k=t_{pr}}^{t_{st}} f_{n,m}(k)}{\text{RLI}_{n,m}},$$

$$n, m = 1, 2, \dots, 128$$

where t runs over positive integers to denote latencies, t_{st} is the latency of stimulus onset, t_{pr} is a prestimulus latency (such that $t_{st} - t_{pr} = 3$), and RLI denotes resting light intensity. The fluorescence was averaged over a few initial latencies well before t_{st} .

The ensemble of these signals constitutes a set of temporal patterns $\{x_{n,m}(t)\}$, and our methodology starts in a classic

pattern-analytic mode by a feature extraction step. Latencies of interest (LOIs) are defined, and the corresponding sets of signal values play the role of feature vectors. In this way, the OB response is represented by a point swarm in a space with dimensionality, depending on the number p of selected LOIs. The coordinates of these $N = 128 \times 128$ points are tabulated in a $(N \times p)$ sized matrix \mathbf{X}^{data} with rows of the form

$$X_i = [X_{n,m}(t_1), x_{n,m}(t_2), \dots, x_{n,m}(t_p)], \quad i = 1, 2, \dots, N.$$

Figure 1(b) demonstrates the embedding of the OB response dynamics in a 3-D feature space. For the selection of LOIs, the time course in Figure 1(a) was built, which is the OB response averaged across all pixels (grand average), and a triplet of latencies was defined so as to be located within the time range of the prominent positive deflection. The selected latencies are denoted using dots of different colors [Figure 1(a)]. To show the association between LOI selection and response representation, the same colors have been used for drawing the axes in Figure 1(b). Needless to say, a different triplet would have resulted in a point swarm of different geometry, while a more complete coverage of the deflection using additional LOIs would have resulted in a higher dimensional feature space. Hereafter, as a simple, data-driven strategy for defining LOIs, we select all the latencies corresponding to a continuous time range, which starts around the response onset and covers well the maximal peak seen in the grand average.

Clustering, Vector Quantization, and Prototyping

The partition of all N temporal patterns into groups of homogeneous time courses is the most direct way to summarize the temporal variations in the OB response. The task is accomplished by applying cluster analysis in feature space and, then, transferring the results to the original pattern space. The end product would be a set of cluster centers and a set of cluster-membership lists. The cluster centers serve as reference points in the feature space and can be used in a vector quantization (VQ) scheme for the coarse representation of the entire point swarm. VQ encodes the response manifold in the feature space by utilizing only a finite set of reference vectors (i.e., prototypes), the code vectors. It actually performs a parcellation of the feature space known as Voronoi tessellation [5], in which a Voronoi region is defined around each code vector. This section in the feature space is composed of all the points closer to a specific code vector than to any other. The feature vectors falling within a Voronoi region are represented by the corresponding code vector. The number of code vectors k controls the resolution of the representation (i.e., the level of information abstraction). It is a user-defined parameter that can be selected as the best compromise between the coding error and the computational burden. The red circles in Figure 1(b) denote those ten reference vectors that summarize the point swarm with the minimal distortion.

In our approach, a codebook of k code vectors is designed by applying the neural-gas algorithm to the data matrix $\mathbf{X}^{\text{data}} = [X_1|X_2|\dots|X_N]$. This algorithm is an artificial neural network model, which converges efficiently to a small, user-defined number $k < N$ of codebook vectors, using a stochastic gradient descent procedure with a soft-max adaptation rule that minimizes the average distortion error [7]. The network ability to respect the intrinsic data dimensionality is known [8], and this

makes it the best candidate for our manifold learning purposes. The computed code vectors $O_j \in \mathbf{R}^p, j = 1, 2, \dots, k$ are used in a simple encoding scheme. The nearest code vector is assigned to each X_i in \mathbf{X}^{data} . Strictly speaking, this procedure divides the response manifold $\mathbf{V} \subset \mathbf{R}^p$ into k Voronoi regions.

$$\mathbf{V}_j = \{X \in \mathbf{V} : \|X - O_j\| \leq \|X - O_i\|, \forall i, i = 1, 2, \dots, k\}.$$

From a more practical point of view, the bulk of information contained in the data matrix is represented, in a parsimonious way, by a $(N \times k)$ partition matrix \mathbf{U} , with elements u_{ij} such that

$$u_{ij} = \begin{cases} 1 & \text{if } X_i \in \mathbf{V}_j \\ 0 & \text{if } X_i \notin \mathbf{V}_j \end{cases}, \quad \sum_j \sum_{i=1}^N u_{ij} = \sum_j N_j = N. \quad (1)$$

The fidelity of the encoding is quantified (inversely) via the following index, which is the total distortion error divided by the total dispersion of the data:

$$n_{\text{Distortion}} = \frac{\sum_{i=1}^N u_{ij} \|X_i - O_j\|^2}{\sum_{i=1}^N \|X_i - \bar{X}\|^2}, \quad \bar{X} = \frac{1}{N} \sum_{i=1}^N X_i. \quad (2)$$

The smaller the $n_{\text{Distortion}}$, the better the encoding. This index gets smaller with the increase of k and approaches zero when k becomes comparable to N .

In the second step, the computed k partition is applied to the set of temporal patterns $\{x_{n,m}(t)\}$ to produce prototypical responses. The exact correspondence between the pixel coordinates (n, m) and rows in the data matrix \mathbf{X}^{data} is preserved and the temporal patterns are split into k groups according to the partition of the feature vectors. Representative time courses are then estimated via within-group averaging. The k computed curves reside in pattern space, have full time resolution, and therefore portray the response dynamics in an interpretable format.

Charting the Response Manifold Using MDS

The structure of response manifold can be described using the geometrical relationships between the computed code vectors. The reliability of this estimation depends on the encoding error (i.e., the value of k) and the way in which relationships in the feature space are captured. The computation of all pairwise Euclidean distances, $D(i, j) = \|O_i - O_j\|_{L_2}, i, j = 1, \dots, k$ and their tabulation in a $(k \times k)$ distance matrix \mathbf{D} , is the standard initialization to many structure analysis methods. In the following, we describe the steps to visualize this structure and relate it with spatiotemporal characteristics of the OB response.

A distance-preserving technique, known as metric multidimensional scaling (MDS), is utilized to produce a low-dimensional point diagram, reflecting as close as possible the Euclidean distances mentioned previously. Based on simple algebraic transformations and eigen analysis of \mathbf{D} , the code vectors are placed in a coordinated space of dimensionality $r < p$ (usually 2 or 3) [9]. We denote the overall mapping scheme as $\mathbf{P} = \text{MDS}(\mathbf{D}, r)$, with \mathbf{P} a $(k \times r)$ matrix such that the i th row defines the image P_i of the code vector O_i in a reduced feature space. The new point diagram condenses the structural information estimated in the original feature space.

The success of this representation is measured with indices such as the normalized total discrepancy

$$\text{Stress} = \frac{\sum_{i < j} (D(i, j) - \|P_i - P_j\|_{L_2})^2}{\sum_{i < j} D(i, j)^2}. \quad (3)$$

The smaller the stress, the more reliable the mapping. This index depends strongly on the dimensionality of the mapping. It becomes smaller as r gets bigger. Here, we adopt the value 2 for r , i.e., $X_i \mapsto P_i = [p_1(i) p_2(i)] \in \mathbb{R}^2$, since our goal is to sketch the (approximate) response manifold in a plane and gain direct insight into the temporal variations. We accompany the produced point diagram with the stress index as a measure of its creditability. Figure 1(c) shows the two-dimensional (2-D) point diagram that corresponds to the ten code vectors shown in Figure 1(b). Hence, the dimensionality is reduced by one.

Two nearby points in such a 2-D point diagram correspond to similar code vectors, while two distant ones correspond to deviant code vectors. A compact point diagram is hence indicative of a uniform OB response, whereas a fragmented one reflects significant temporal variations. This sort of structural information is exploited further with the use of a minimal spanning tree (MST) graph [10]. This is a popular graph-theoretic tool that can provide both the nearest neighbor information about each point and the shortest linkage information about subsets of points. A node is associated with the endpoint of each vector P_i , and an MST is the connected graph formed by connecting all these nodes with the unique set of line segments having the minimal total length. As it can be seen in Figure 1(c), there is a unique path connecting each pair of nodes that can be followed to navigate in the (reduced) feature space. Associated with this property of the MST, there is a graph-theoretic procedure for transforming this graph into an ordered list, in which similar ranks are assigned to nearby nodes [5], [11]. It is based on a standard technique for traversing a tree graph and is accomplished by selecting one node of the MST as the root and following a breadth first search for the remaining ones. The rank of each node is the order in which it is visited in this search [see Figure 1(c)]. These ranks replace the original (random) labels of the code vectors and are always implied henceforth. The so-called MST ordering provides a content-dependent reorganization of the list of code vectors that contributes significantly to their efficient parameterization and handling and, consequently, all the temporal patterns [5].

The charting of the response manifold, apart from the MST-ranked point diagram, includes the ordered list of prototypical responses [following the order of the corresponding code vectors, Figure 1(e)], a pictorial representation of the portions $\alpha_j = N_j/N$ of temporal patterns assigned to each of the k groups [Figure 1(d)], and a spatial display in which pixels are colored in accordance with the order of their group [Figure 1(f)]. With the latter step, the spatial distribution of the temporal code is examined. It is the most critical, since the coupling of the temporal with the spatial domain is attempted. The interface is the MST-ordering scheme, and the success relies on how well the assigned ranks reflect the variation in the codebook. This in turn depends on the selection of the rooting code vector. The user can try different rankings and decide on the most suitable. The seriation of corresponding prototypical responses [e.g., Figure 1(e)] can help in specifying the

reliability of a particular ranking. In practice, however, there is a more convenient way of selecting the root. It includes the embedding of the MST graph in a simple graphical user interface (GUI) environment that enables the user interaction with the graphics [5]. Therein, root selection is simplified to the task of identifying points at the ends of graph diameters [11].

Figure 2 provides a realistic example of the MDS-based charting. It is based on the same data used for the didactic example of Figure 1, although there are two important differences. The first is that multiple LOIs have been selected so as to fully cover the peak seen in the grand average [shown in Figure 2(a) as an overplotted, thin black curve]. The horizontal bar denotes the selected latency range, which implies an embedding of the OB response in a feature space of $p = 61$ dimensions. The second difference is the selection of $k = 50$ code vectors, which provides a more reliable summary of response manifold ($n_{\text{Distortion}}$ index was 0.2601, approximately three times smaller than for $k = 10$). The MDS-based point diagram of Figure 2(b) is characterized by the stress value of 0.092, which shows that a satisfying 2-D sketch of manifold structure can be obtained. Such a high-information compression (from $p = 61$ to $r = 2$ dimensions) indicates strong temporal correlations between consecutive latencies and high-spatial covariations between temporal patterns from adjacent pixels. Due to the MST-graph elongation, rooting was an easy task (the only other reasonable alternatives were the nodes currently labeled as 45, 50, and 49, which would have led to opposite ranking but without any qualitative difference). The ranks have been distributed as shown and used to order the list of prototypical responses accordingly. The ordered list can be thought of as the spectrum of the OB's temporal patterning. By sampling some of these responses [Figure 2(a)] and denoting the placement of the corresponding code vectors on the MST graph [Figure 2(b)], we can interpret the whole structure seen in the graph. For instance, it is clear that the rightmost branches of the MST correspond to increased OB activity. Finally, the spatial distribution of ranks in Figure 2(c), using a color code fully compatible with the presentation in previous panels, concludes our exploration of the sensory code. The way to read this display can be summarized as follows: two pixels of comparable hue correspond to similar temporal patterning. A neighbor of similarly colored pixels indicates a coherently responding OB patch and, therefore, a possible glomerulus. The exact form of patterning can be inferred using the accompanying MST graph [Figure 2(b)] and the ordered list in Figure 2(a).

Unfolding the Response Manifold Using ISOMAP

The previously described MDS procedure results in a visualization of structure that provides a global overview of the response manifold. Being akin to the principal component analysis (PCA) preserves the total sum of pairwise distances and, consequently, can recover nonlinear manifolds only partially. To capture intricate geometries, we have to resort to ISOMAP [12] (or any other manifold learning technique from the recently introduced ones [13]), a method that advances MDS by incorporating geodesic distances as the way to express geometrical relationships.

In our case, ISOMAP is applied as follows [9]: a ν nearest-neighbor graph is built over the code vectors O_j by connecting each vector with its ν nearest neighbors. Using the formed graph, the length of the shortest path between every pair of code vectors is computed as an approximation to their geodesic distance on the manifold and tabulated in a $(k \times k)$ matrix \mathbf{GD} .

The application of classic MDS, i.e., $\mathbf{P} = \text{MDS}(\mathbf{GD}, r)$, provides the coordinates of code vectors in a space of lower dimensionality $r < p$. The recovery of the manifold becomes more complete as r is increased. For visualization purposes, however, we always select $r = 2$ and, thus, attempt to represent the intrinsic nonlinear degrees of freedom via a 2-D reduced point diagram. The residual variance quantifies the error in this approximation [12]. Keeping the analogy with (3), we express it as

$$G_{\text{Stress}} = \frac{\sum_{i < j} (\text{GD}(i, j) - \|P_i - P_j\|_{L_2})^2}{\sum_{i < j} \text{GD}(i, j)^2}.$$

A second important parameter is the number v of neighbors used in building the graph. With small values, the details of local geometry are emphasized but at the expense of graph cohesiveness. Disconnected components in the graph require separate treatment, and this might add unaffordable complication in the study. On the other hand, for very high v , ISOMAP degenerates into MDS. In the presented results, v was selected throughout as the minimum value that resulted in a single component graph.

To characterize and interpret the structure reflected in the reduced point diagram of code vectors, it is necessary to understand the association between this diagram and the pattern space. The MST graph is drawn over this point set and utilized to parameterize the set. Unlike in MDS charting, however, here the MST-based ranks may not always be suitable for imposing a smooth seriation of the code vectors (and, consequently, the prototypical responses). The scope of the reduced point diagram, now, is to unfold response dynamics. Clustering tendencies have thus been intensified, and the diagram is better suited for visual data mining. The understanding of its structure can be promoted significantly via the interactive study of the reduced point diagram within a GUI setup [5]. Finally, it is important to note that the increase in the number k of code vectors advances the unfolding, since path distances approach true geodesic distances. It should be kept below ~ 500 , though, for admissible computational burden.

Figure 3 demonstrates, with the data used previously for Figures 1 and 2, the unfolding of the response manifold. While the embedding dimension was kept the same (i.e., $p = 61$), the number k of code vectors was increased to 200 ($n_{\text{Distortion}} = 0.1842$). ISOMAP provided a graph with G_{Stress} of only 0.023 (0.096, 0.023, and 0.022 were the consecutive values for $r = 1, 2$, and 3,

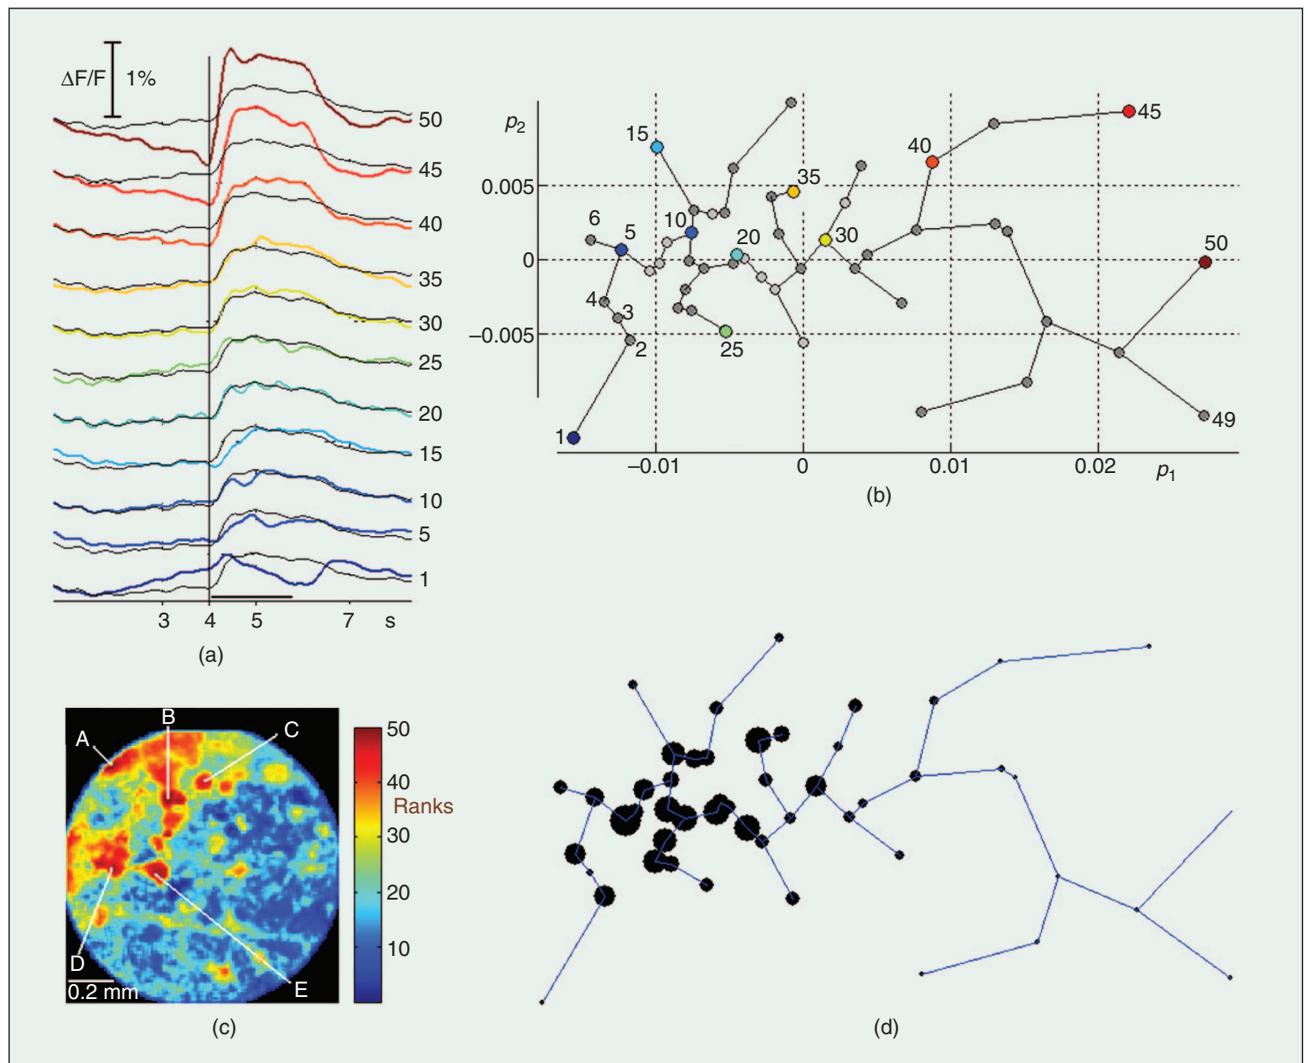


Fig. 2. MDS-based characterization of response manifold.

respectively). This was a clear indication that response dynamics can be faithfully represented using a low-dimensional graph. The shown 2-D point diagram has a main component on its right. This becomes more evident in the inset, in which dark disks denote the relative populations corresponding to the code

vectors. The overplotted MST presents a small tip at one end and two main branches at the other. Ranks have been distributed by rooting at the tip. The complexity of branching is such that the MST ordering cannot provide a single list with a succession that faithfully reflects the shown structure. However,

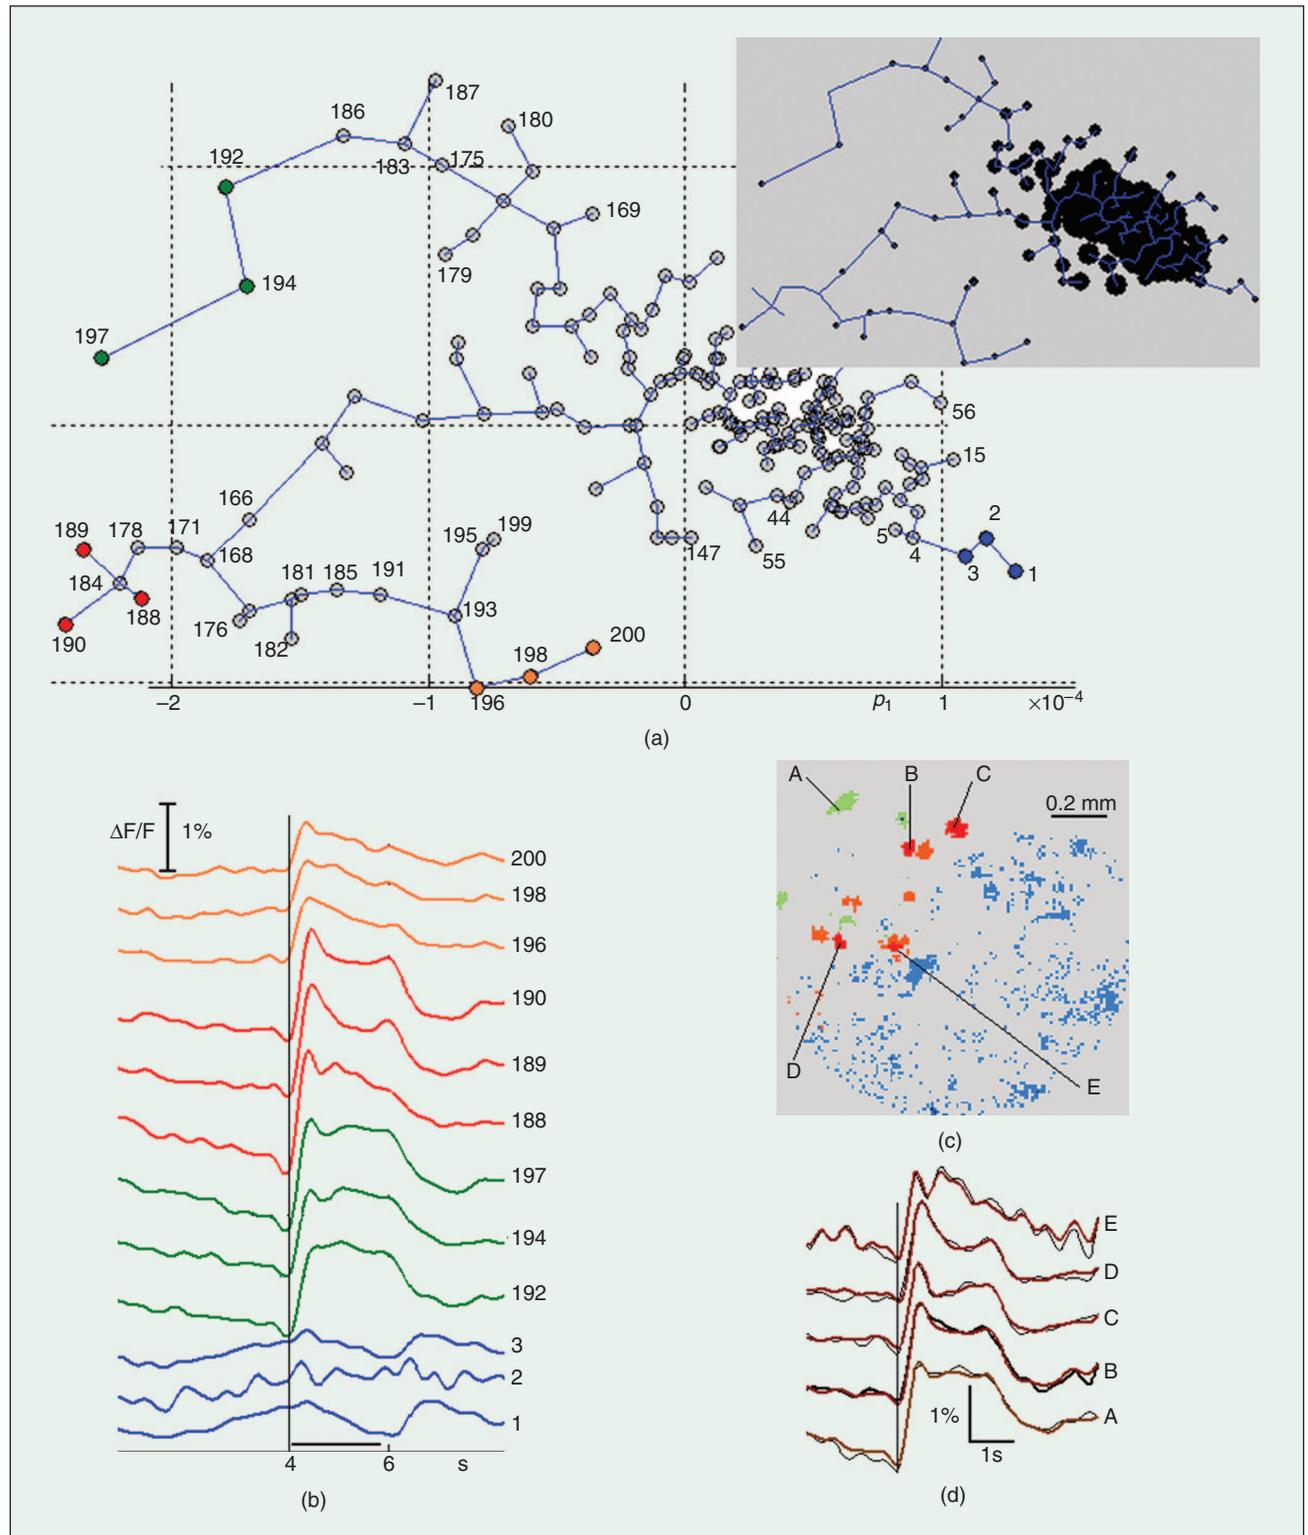


Fig. 3. ISOMAP-based characterization of the response manifold. (d) The five traces from conventional analysis (Figure 1(g)) are compared with the time courses (shown in brown) produced by averaging the activity from all pixels, around each of the A–E pixels, which are connected using the same color.

the true (nonlinear) degrees of freedom can be identified from the different continuous paths emerging in the reduced space. We can, therefore, select a few nodes from the most distant branches and present the corresponding prototypical responses. In Figure 3(a), the selected code vectors are denoted using a different color, one for each distinct path. Keeping the same colors, Figure 3(b) presents the corresponding prototypical responses in an orderly fashion, and Figure 3(c) shows the spatial distribution of the corresponding pixels. The pixels A–E, detected in Figure 1(g) as loci of putative glomeruli, were among the colored ones in the spatial display of Figure 3(c). To further show how our method relates with the standard approach, Figure 3(d) includes the corresponding five traces produced by simply averaging the time courses of each spatially connected and identically colored set of pixels. There is a great resemblance of these traces with the manually produced ones (shown as brown and black curves, respectively). Additionally, there is a remarkable resemblance between these traces and the prototypical time courses (A → 197, B → 190, C and D → 189, and E → 188).

Comparing Responses

Response Submanifolds' Contrast

Comparing responses recorded from the same OB can promote our understanding of sensory encoding. Responses to different odors might reveal a code with innate discriminatory power. Responses to repeated odor presentations will make the systematics clear and help quantify trial to trial variability. Finally, responses from different conditions (e.g., awake versus anesthetized animal) may show the influence of the brain's state.

All the comparisons discussed previously can be mediated through the suggested graphical representation of a point diagram that is accompanied by the semantics of prototypes and the spatial labeling of pixels. However, instead of just creating one single 2-D point diagram for every response and trying to measure the match (a versatile correspondence problem), we suggest a modified strategy that is conceptually straightforward and results in a convenient way of contrasting responses. It starts, as before, by embedding all the responses in a common feature space, by selecting the same or homologous LOIs, and

proceeds with encoding the overall manifold. A 2-D representation of the global codebook is produced using either MDS or ISOMAP. This step is followed by ordering or mining code vectors accordingly. There are only two modifications that are introduced for comparison purposes (see Figures 4–6). The first is in the way we represent the number of patterns populating the different groups. Depending on the situation, either the relative proportions are shown or multiple population graphs are provided. The second modification is in the way the spatial distribution of group labels is shown. Now, there is still a single spatial display for every response, but underlying it there is a common codebook that allows direct visual comparison. Note that by using a common codebook, we may sacrifice a more detailed description of the individual responses. This would be the case whenever the responses form clearly distinct submanifolds.

Specific examples of the suggested visual comparisons follow. The data are obtained from rat OB-imaging experiments. The difficulty in analyzing these data stems from the fact that the animals'

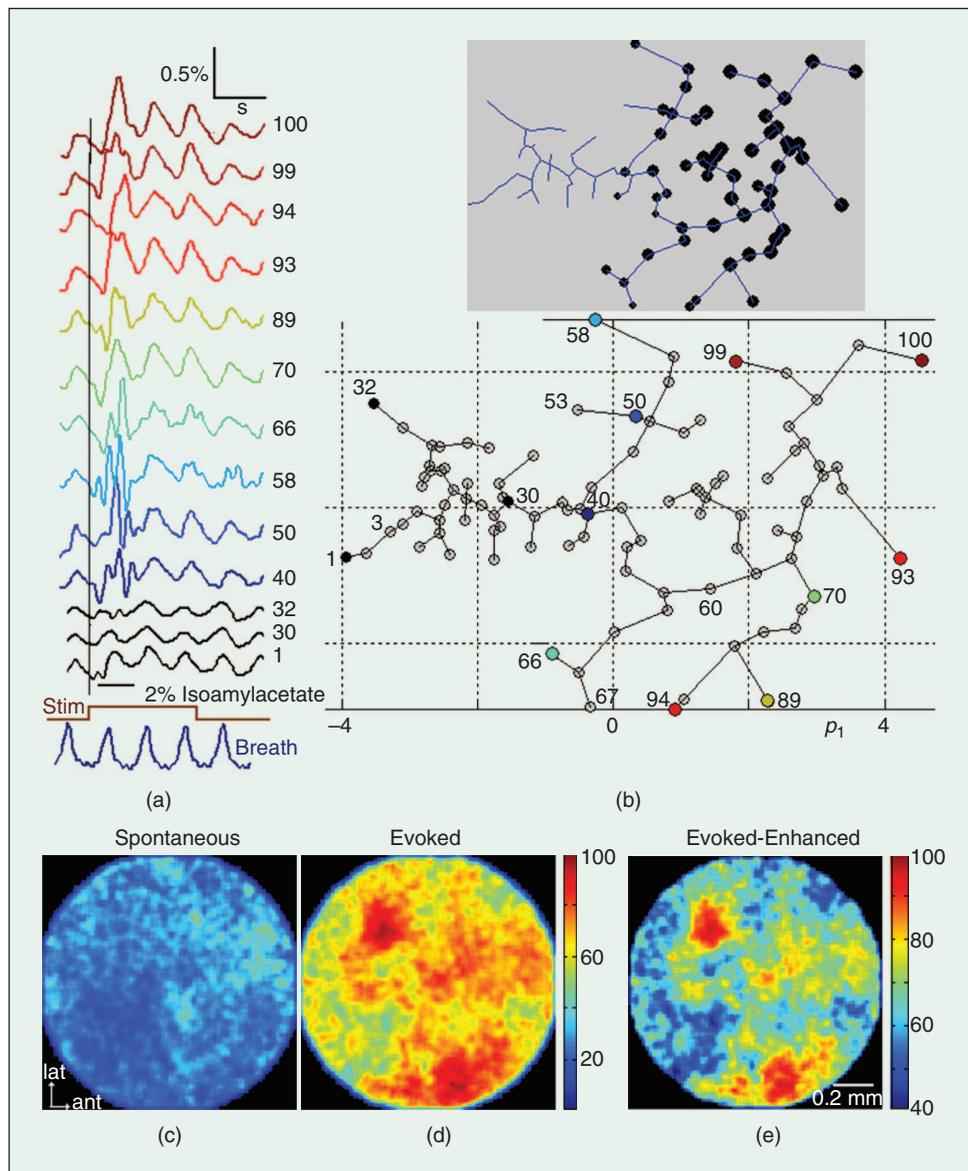


Fig. 4. An MDS-based contrast of spontaneous and evoked activity.

nasal airflow is not controlled, and the dura matter above the OB is more opaque than in mice. As a result, the glomerular pattern of activation is rather diffuse in these experiments.

Spontaneous Versus Evoked Activity

An estimation of true evoked signal, unbiased by the omnipresent ongoing activity, was attempted through the comparison of actual responses with the OB activity recorded via an identical experimental procedure but with no odorant being delivered. LOIs were defined from the evoked response data as the whole latency range containing the first peak in the poststimulus interval. Using the same LOIs, evoked and spontaneous activity patterns were jointly embedded and encoded. We then tried to disentangle them as a means of attributing the observed structure to noise and signal. Two typical cases are discussed below. Apart from referring to different animals, different manifold representation techniques have been employed.

MDS charting was employed in the example of Figure 4 ($n_{\text{Distortion}} = 0.32$; stress = 0.110). The inset portrays the ratio of evoked activity patterns to spontaneous activity ones for each code vector. This visualization makes apparent that the leftmost part of the MST graph is dedicated to spontaneous activity while the rightmost part to evoked activity. Rooting as shown in Figure 4(b) provides ranks that correlate with signal content. Low-ranked (<40) prototypes are indicative of spontaneous activity and, thus, can be excluded from the ordered list in Figure 4(a). In contrast, the top-ranked prototypical responses conceptualize the true evoked signal. In the temporal domain, the response appears to be an amplitude modulation that overrides oscillations related to the (uncontrolled) breathing of the animal [14]. The spatial distribution of ranks for spontaneous and evoked activity is included in Figure 4(c) and (d) accordingly. Their visual contrast clearly shows the existence of a spatial code, in the form of two main extended blobs of activity. The localized character of the spatial code is emphasized in Figure 4(e) by using a color code in which the labels or ranks corresponding exclusively to spontaneous activity have been suppressed.

ISOMAP was employed to represent the global codebook in Figure 5 ($n_{\text{Distortion}} = 17$, $G_{\text{Stress}} = 0.045$). The MST graph of Figure 5(a) clearly indicates the presence of two distinct submanifolds and, in conjunction with Figure 5(b), one of them is related exclusively with patterns of evoked activity. This implies a very strong response or, speaking in engineering terms, very high SNR. The ranks above 45 correlate with signal content, and since they follow smoothly, the elongation of response submanifold can be used for the seriation of prototypical responses [Figure 5(c)] and color coding [Figure 5(d)]. The revealed temporal code is similar with the one shown in Figure 4 (for a different animal). However, a unified,

spatially extended region in the OB appears to be significantly activated here. This is probably a consequence of the response strength. Another observation worth mentioning, and in accordance with recently reported results [15], is that pattern dynamics within the first respiration cycle recur in a similar manner during consecutive cycles. Although the grouping had been based on the latencies of the first response peak, the differences between the prototypical responses remained well thereafter [Figure 5(c)]. This was further verified by repeating the whole procedure (not shown here) for the LOIs corresponding to the second-response peak.

Comparing Odor Representations

Our understanding of the animal's mechanisms for discriminating odors can be promoted significantly via the comparative study of different odor representations in the same OB. Two different odors, 2-hexanone and isoamylacetate (denoted as A and B, respectively), were first presented separately at equal concentrations (~2% of saturated vapor) and then in two mixtures. In the first mixture, odor A had a fivefold larger concentration than B (A, 2%; B, 0.4%; A:B = 5:1) while in the second, odor B was the major component (A, 0.4%; B, 2%; A:B = 1:5). The four different types of response, evoked in the same animal, were recorded in a single experiment and could therefore be contrasted. The LOIs were defined for each response separately, and their union defined the dimensionality for the total embedding. In the reduced point diagram obtained, we tried to understand how manifold structure reflects the two different odor representations (A odor versus B odor) and if the blending of the two odors appears as a geometrical amalgamation.

Figure 6 contains an example of the sketched structure based on 100 code vectors and MDS ($n_{\text{Distortion}} = 0.35$ and stress = 0.19). The absolute proportions, which are depicted for each response separately in Figure 6(d), clearly indicate differences in the temporal patterning between all four responses. The subgraphs of two mixtures lie in between the two single-odor graphs, which are totally separated. The relative spatial position reflects the natural order of the four concentrations (from right to left, A:B is 1:0, 5:1, 1:5, 0:1).

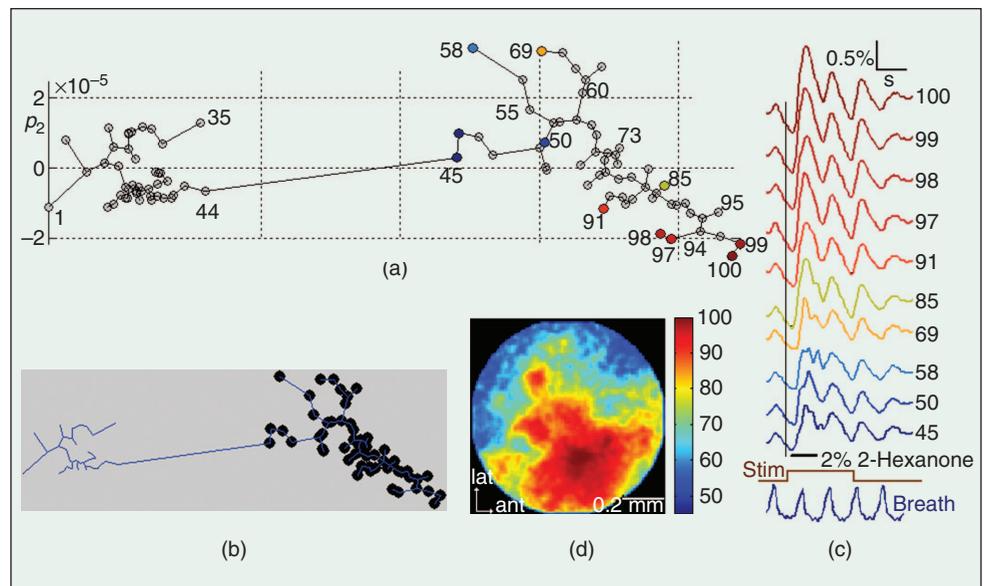


Fig. 5. An ISOMAP-based contrast of spontaneous and evoked activity.

Interestingly, in the mixture graphs, density points have moved away from the regions that were occupied by odors A and B. However, in information space, mixtures stay closer to their major component odor. The prototypical responses with ranks 100 and 40 can easily be recognized as representatives for odor A and B, respectively [Figure 6(a)], provided that stimulus presentations were temporarily aligned (triggered by the breathing cycle). The observed delay in amplitude modulation is indicative of a code carrying enough discriminatory information in the phase of the signal [3]. Figure 6(c) is the spatial counterpart of the OB's discriminatory function.

Discussion

Manifold learning is an emerging field in empirical data analysis. It includes algorithmic techniques that can recover the smooth or regular structure from naturally occurring signals of high-extrinsic dimension. We suggest that the manifold formed by the response spatiotemporal patterning in the imaged OB is actually low dimensional. Hence, its nonlinear degrees of freedom can be identified as a means of fully characterizing the evoked activity.

Our proposal uses one efficient self-organizing network (neural-gas network) for down-sampling the original manifold, two alternative visualization techniques (MDS and ISOMAP) for representing the manifold through 2-D point diagrams, and one graph-theoretic technique (MST) for handling the emerged structure. MDS can provide a fast, global view of the response manifold and therefore mediate a smooth mapping (from 2-D point diagram to pixels) for coloring the imaged OB region. ISOMAP, being a manifold learning technique itself, can offer a detailed representation of manifold geometry at increased computational cost, which is, however, kept at the affordable level, thanks to preceding neural-gas-based prototyping.

The experiments showed that a low-dimensional information space can indeed be recovered and parameterized in a meaningful way. The low dimensionality, a hallmark of self-organization in neural systems, suggests that the present manifold learning approach is very promising for gaining insight into response generation mechanisms. The included examples demonstrate vividly the diversity of sensory encoding schemes in the OB. In this article, we have placed emphasis on the visualization of results and have refrained from quantifications. These are readily

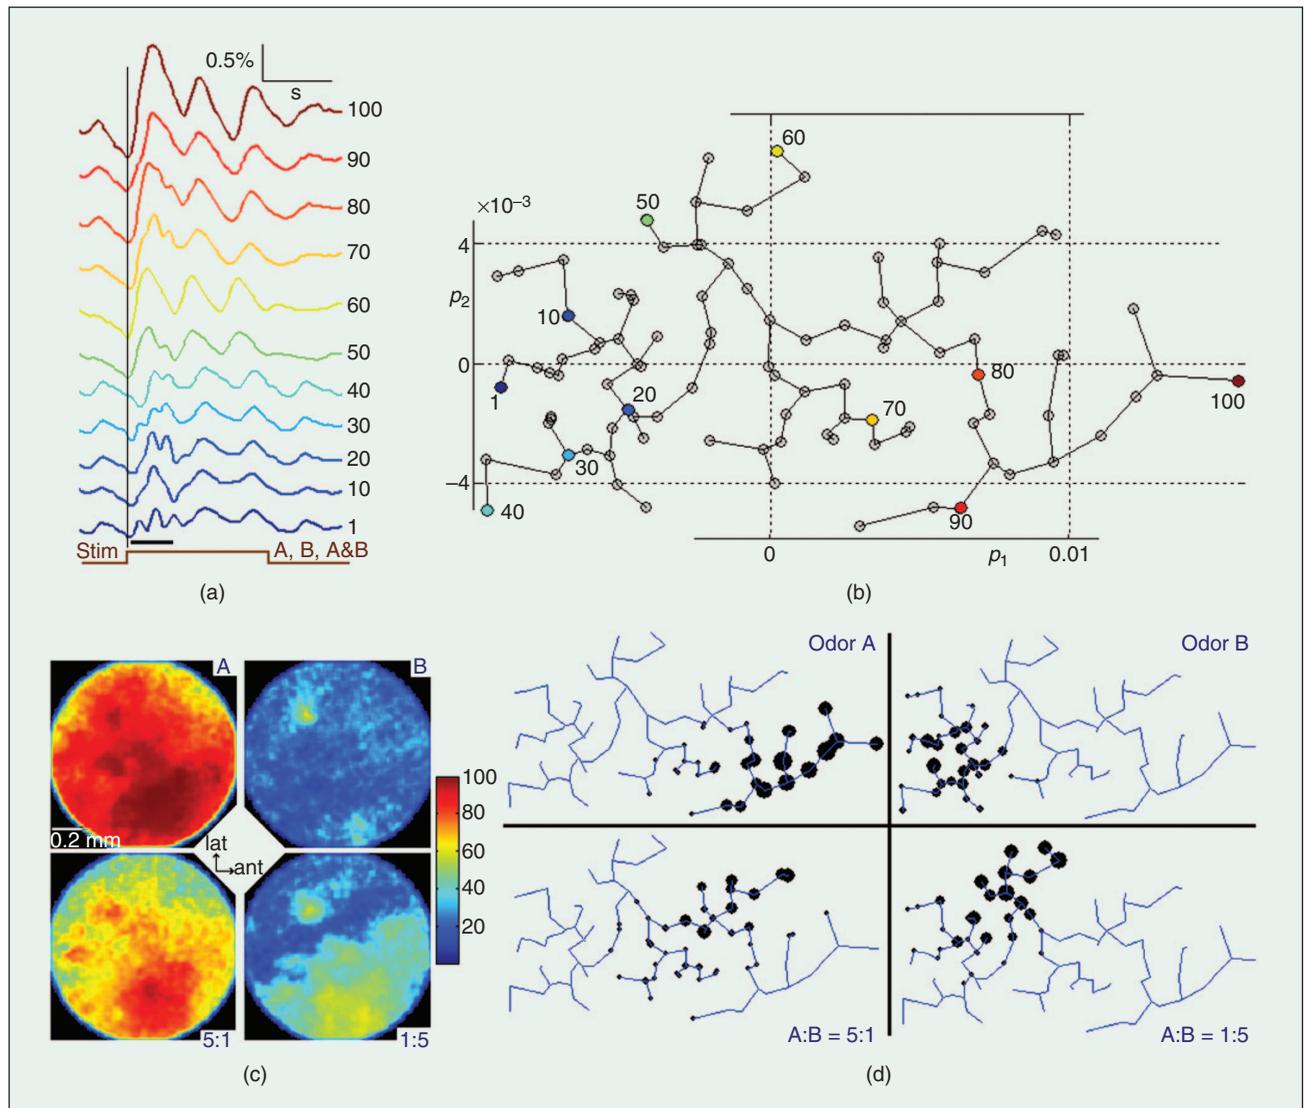


Fig. 6. Comparing odor representations via MDS charting.

available through topological descriptors [5], [10], [17] that can be used for a more thorough analysis.

In a recent article [16], the combination of the PCA with spatial independent component analysis (SICA) was suggested as a powerful technique for unbiased analysis and dissection of related imaging data. Different spatial modes, tightly associated with specific time-dependent amplitudes, are recognized with some of them corresponding to artifacts and others to distinct functional domains. Both PCA-SICA and our approach share an exploratory character, making (almost) no assumptions regarding the different signal sources contributing to the data. The common goal is, by fully exploiting the high-spatiotemporal resolution, to recover the stimulus-related information. However, the two approaches differ in many ways. An important theoretical difference is that a linear additive model relies on independent component analysis (and PCA), while our approach can capture nonlinearities. From a practical point of view, the data are decomposed in different spatial modes with independent component analysis, while, in our case, information is condensed in a single spatial mode, which facilitates direct comparison between multiple responses. Finally, we should point out that our approach must follow proper data preprocessing, which in the future can benefit, subsequently, from the artifact-cleaning abilities of independent component analysis to minimize the contributions from sources unrelated to the stimulus.

Manifold learning finds applications to other brain responses as well. The shaping of event-related dynamics manifold due to trial to trial variability was the focus of a recent study based on human neuromagnetic responses to visual stimuli [17].

Acknowledgment

Nikolaos A. Laskaris is deeply indebted to A. A. Ioannides for previous related discussions. This work was partially supported by NIH Grants DC-05259 and NS-07455.



Nikolaos A. Laskaris received his Ph.D. degree in biomedical engineering from Patras University, Greece. He then joined the Laboratory for Human Brain Dynamics in the Brain Science Institute of RIKEN in Japan (1999–2003). Currently, he is a faculty member of the Informatics Department in Aristotle University of Thessaloniki (AUTH), Greece. His interests include computational intelligence, data mining, nonlinear dynamics, and their applications in biomedicine and neuroscience.



Efstratios K. Kosmidis received his B.S. degree in biology from the AUTH, Greece, and his Ph.D. in computational neuroscience from the Paris 6 University, France. He was a research associate at the Yale School of Medicine, New Haven, Connecticut, working on optical recordings. He is currently serving his military duties in the 424 Army General Training Hospital.



Dejan Vućinić received his Ph.D. degree in physics from the Massachusetts Institute of Technology. As a postdoctoral fellow at Yale University, New Haven, Connecticut, he performed optical recordings in vivo of neuronal activity at the input to the mouse olfactory bulb. He is currently improving the techniques for two-photon optical recording of neuronal membrane voltage.



Ryota Homma received his Ph.D. degree in biophysics from the University of Tokyo, Japan. He was a research scientist in the Laboratory of Integrative Neural Systems in RIKEN Brain Science Institute in Japan (1999–2004). Currently, he is a postdoctoral associate at the Yale School of Medicine, New Haven, Connecticut. He has been working on in vivo optical imaging of the brain.

Address for Correspondence: Nikolaos A. Laskaris, Artificial Intelligence and Information Analysis Laboratory, Department of Informatics, Aristotle University Biology Building, GR-54124 Thessaloniki (box 451), Greece. E-mail: laskaris@aiia.csd.auth.gr.

References

- [1] L. Buck and R. Axel, "A novel multigene family may encode odorant receptors: a molecular basis for odor recognition," *Cell*, vol. 65, no. 1, pp. 175–187, 1991.
- [2] M. Wachowiak and L. B. Cohen, "Representation of odorants by receptor neuron input to the mouse olfactory bulb," *Neuron*, vol. 32, no. 4, pp. 723–735, 2001.
- [3] H. Spors and A. Grinvald, "Spatio-temporal dynamics of odor representations in the mammalian olfactory bulb," *Neuron*, vol. 34, no. 22, pp. 301–315, 2002.
- [4] D. Vucinic, L. B. Cohen, and E. K. Kosmidis, "Interglomerular center-surround inhibition shapes odorant-evoked input to mouse olfactory bulb in vivo," *J. Neurophysiol.*, vol. 95, no. 3, pp. 1881–1887, 2006.
- [5] N. Laskaris, S. Fotopoulos, and A. Ioannides, "Mining information from event-related recordings," *IEEE Signal Processing Mag.*, vol. 21, no. 3, pp. 66–77, 2004.
- [6] W. J. Freeman, *Neurodynamics: An Exploration in Mesoscopic Brain Dynamics*. London: Springer, 2000.
- [7] T. Martinez, S. Berkovich, and K. Schulten, "Neural-gas network for vector quantization and its application to time-series prediction," *IEEE Trans. Neural Networks*, vol. 4, no. 4, pp. 558–569, 1993.
- [8] T. Martinez and K. Schulten, "Topology representing networks," *Neural Netw.*, vol. 7, no. 3, pp. 507–522, 1994.
- [9] N. Laskaris and A. A. Ioannides, "Semantic geodesic maps: A unifying geometrical approach for studying the structure and dynamics of single trial evoked responses," *Clin. Neurophysiol.*, vol. 113, no. 8, pp. 1209–1226, 2002.
- [10] N. Laskaris and A. A. Ioannides, "Exploratory data analysis of evoked response single trials based on minimal spanning tree," *Clin. Neurophysiol.*, vol. 112, no. 4, pp. 698–712, 2001.
- [11] J. Friedman and L. Rafsky, "Graphics for the multivariate two-sample problem," *J. Am. Statist. Assoc.*, vol. 76, no. 374, pp. 277–287, 1981.
- [12] J. B. Tenenbaum, V. de Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [13] M. Law, Manifold learning. [Online]. Available: <http://www.cse.msu.edu/~lawhiu/manifold/>
- [14] A. Arieli, A. Sterkin, A. Grinvald, and A. Aertsen, "Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses," *Science*, vol. 273, no. 5283, pp. 1868–1871, 1996.
- [15] H. Spors, M. Wachowiak, L. B. Cohen, and R. W. Friedrich, "Temporal dynamics and latency patterns of receptor neuron input to the olfactory bulb," *J. Neurosci.*, vol. 26, no. 4, pp. 1247–1259, 2006.
- [16] J. Reidl, J. Starke, D. B. Omer, A. Grinvald, and H. Spors, "Independent component analysis of high-resolution imaging data identifies distinct functional domains," *NeuroImage*, vol. 34, no. 1, pp. 94–108, 2007.
- [17] N. Laskaris, L. C. Liu, and A. A. Ioannides, "Single-trial variability in early visual neuromagnetic responses: An explorative study based on the regional activation contributing to the N70-m peak," *NeuroImage*, vol. 20, no. 2, pp. 765–783, 2003.