

RESEARCH ARTICLE

Topographic Independent Component Analysis reveals random scrambling of orientation in visual space

Marina Martinez-Garcia^{1,2}, Luis M. Martinez¹, Jesús Malo^{2*}

1 Instituto de Neurociencias, CSIC, Alicante, Spain, **2** Image Processing Lab., Universitat de València, València, Spain

* jesus.malo@uv.es



OPEN ACCESS

Citation: Martinez-Garcia M, Martinez LM, Malo J (2017) Topographic Independent Component Analysis reveals random scrambling of orientation in visual space. PLoS ONE 12(6): e0178345. <https://doi.org/10.1371/journal.pone.0178345>

Editor: Manabu Sakakibara, Tokai University, JAPAN

Received: November 24, 2016

Accepted: May 11, 2017

Published: June 22, 2017

Copyright: © 2017 Martinez-Garcia et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Original images used in the experiments are publicly available at: <http://research.ics.aalto.fi/ica/imageica/>. Specific code to generate and analyze the results is publicly available at <http://isp.uv.es/code/visioncolor/TICAdomains/TICAdomains.html> and on Github: <https://github.com/neuronavalenciana/OrientationDomainsTICA>.

Funding: The authors acknowledge the financial support received from the Spanish Ministry of Economy and Competitiveness (MINECO) under the grants TIN2015-71130-REDT, BFU2014-

Abstract

Neurons at primary visual cortex (V1) in humans and other species are edge filters organized in *orientation maps*. In these maps, neurons with similar orientation preference are clustered together in iso-orientation domains. These maps have two fundamental properties: (1) *retinotopy*, i.e. correspondence between displacements at the image space and displacements at the cortical surface, and (2) a trade-off between good *coverage* of the visual field with all orientations and *continuity* of iso-orientation domains in the cortical space. There is an active debate on the origin of these locally continuous maps. While most of the existing descriptions take purely geometric/mechanistic approaches which disregard the network function, a clear exception to this trend in the literature is the original approach of Hyvärinen and Hoyer based on infomax and Topographic Independent Component Analysis (TICA). Although TICA successfully addresses a number of other properties of V1 simple and complex cells, in this work we question the validity of the orientation maps obtained from TICA. We argue that the maps predicted by TICA can be analyzed in the retinal space, and when doing so, it is apparent that they lack the required *continuity* and *retinotopy*. Here we show that in the orientation maps reported in the TICA literature it is easy to find examples of violation of the continuity between similarly tuned mechanisms in the retinal space, which suggest a random scrambling incompatible with the maps in primates. The new experiments in the retinal space presented here confirm this guess: TICA basis vectors actually follow a random salt-and-pepper organization back in the image space. Therefore, the interesting clusters found in the TICA topology cannot be interpreted as the actual cortical orientation maps found in cats, primates or humans. In conclusion, Topographic ICA does not reproduce cortical orientation maps.

58776-R, TIN2013-50520-EXP, co-financed by the European Regional Development Fund (ERDF), and the Severo Ochoa Program for Centers of Excellence in R+D (SEV-2013-0317). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

1 Introduction

Retinotopy, continuity and coverage in V1 orientation maps

Neurons at the primary visual cortex in mammals are localized edge filters, i.e. their receptive fields are tuned to certain orientation at certain retinal location [1, 2]. The large scale organization of the V1 surface is *retinotopic*, i.e. the retinal space is mapped onto the cortical space so that neighbor cells in the V1 surface are tuned to close regions in the visual field [3–6]. At a finer scale, the organization is retinotopic as well [7–11]: spatial displacements in V1 (in μm) correspond to spatial displacements in the visual field (in *degrees*). *Retinotopy* refers to the correspondence between the location of the receptive field at the retinal space and the location of the neuron at the cortex.

For some species like cats, primates and humans, on top of retinotopy, V1 cells have a quite special arrangement: the so called *orientation maps*. In these maps, neurons with similar orientation preference are clustered together in the so called iso-orientation domains.

On the one hand, these iso-orientation regions are locally *continuous*, i.e. neighbor cells share the same orientation preference except in the boundaries of the region (the fractures and pinwheels). There is certain level of noise (or receptive field scatter) in this continuous distribution, as reported in [12]. However, high-resolution imaging shows that this noise does not prevent the observation of large continuous regions of neurons clearly segregated by orientation tuning in certain species [10], versus others that follow a discontinuous random arrangement [13]. On the other hand, the iso-orientation domains are interleaved for a good *coverage* of the visual field, i.e. the domains are juxtaposed so that continuous patches of similarly tuned neurons can be found all over the visual field [8]. Note that *continuity* of the domains and good *coverage* of the visual field are opposed properties: strict continuity would imply large regions tuned to a specific orientation and hence a poor coverage of the field by neurons tuned to other orientations, while perfect coverage with all orientations is only possible with infinitesimally small (or discontinuous) regions.

In order to describe the functional maps in the mammalian brain, Wilson and Bednar [14] use the terms *topographic maps* and *topological maps*. While *topographic maps* refer to the spatial layout of sensory surfaces according to location (like the skin or the retina), *topological maps* refer to the spatial organization (or grouping) according to features of the sensory input more abstract than location. In the case of animals with structured orientation maps, these are *topological* because of the local continuity of iso-orientation domains, and they are also *topographic* because the *retinotopy* implies that distances in the image space have a well-defined relationship with distances across the cortical surface [14].

Fig 1 replotted from [8] summarizes the basic experimental properties of orientation maps in the visual cortex of many mammals including cats, primates and humans: (1) *retinotopy*, i.e. correspondence between displacements at the image space and displacements at the cortical surface, and (2) a trade-off between *continuity* of each iso-orientation domain, and good *coverage* of the visual field with the different iso-orientation domains.

The grid and retinotopy illustrated in Fig 1 lead to the following key property: *maps measured at V1 have a retinal counterpart*. Note that the receptive fields of a set of neurons located in a grid in V1, arranged in orientation domains, induce an equivalent grid and equivalent orientation domains in the visual space. This is because the retina-V1 transform is smooth: the slightly distorted black grid at the cortex in subfigure C corresponds to a regular grid back in the retinal space. Therefore, one can imagine equivalent orientation domains, with equivalent continuity and coverage, formed by the centers of the receptive fields in the spatial domain. As a consequence, the models should not only reproduce the arrangement in the cortex (i.e. outputs of the processing units organized in orientation maps), but also in the retinal space




Fig 1. Retinotopy, continuity and coverage in V1 orientation maps (reprinted from [8] with permission of Nature Publishing Group). Subfigures A and C correspond to illustrative V1 surface regions in animals with orientation maps (in this case, tree shrews). Colored neighborhoods imply common orientation tuning properties of the neurons: the iso-orientation domains. The color code corresponding to subfigures A and C is below subfigure C. Black scale bars, 200 μm (A), 1mm (C). Black curves in A represent iso-azimuth contours: the cortical location of the maximum response when stimulating with lines at linearly spaced visual angles in the retinal space. Similarly, the black grid in C represents the cortical location of an iso-azimuth/iso-elevation grid in the retinal space. The orientation domains and the iso-azimuth/elevation grids in A and C were obtained using optical imaging. Moreover, multi-unit electrophysiological recording was used to measure the preferred position (central location of the receptive field in the retinal space) at the cortical locations highlighted with letters in subfigure A. Displacements along the cortex imply equivalent displacements (of the receptive fields) along the visual space given the linear relation in subfigure B. Finally, histograms in subfigure C represent the number of neurons tuned to specific orientations in small and large regions of the cortex (or equivalently in small or large visual regions). Continuity of the domains implies peaks in the histograms of small regions, while good coverage of the visual field by different domains implies more uniform histograms when larger visual/cortical regions are considered.

<https://doi.org/10.1371/journal.pone.0178345.g001>

(i.e. the centers of the receptive fields of the predicted units have to be arranged in locally continuous maps according to their orientation tuning).

A functional approach to orientation maps

There is an active debate on the origin of the local continuity of the orientation maps: *where does it come from?, what is its purpose?* [14, 15]. Note that these questions are different in nature: while the first is mechanistic, the second is a functional question.

Most of the literature follows purely mechanistic approaches, in fact there are interesting competing alternatives to describe how the orientation domains arise: (i) as a result of Moire patterns from the structure of the retinal ganglion cells [16, 17] (questioned in [18, 19]), (ii) as a result of self-organization achieved by Hebbian learning and training with natural images [20–22], and (iii) as a result of the local connections leading to global regularities, earlier formulations in [23, 24] and more recent versions based on reaction diffusion in [25] (revisited in [26, 27]). All these works build a mechanism or model which under certain constraints (wiring, connectivity, continuity, coverage, etc.) gives rise to the maps, but these models are not related to the information processing function carried out by the network, nor clarify the functional advantage of having iso-orientation domains [14].

A remarkable exception is the work of Hyvärinen and Hoyer [28] which is based on a functional argument: the network should maximize the transmitted information. Information maximization (or infomax) implies that the responses of the different neurons to natural images should be as statistically independent as possible [29, 30]. Under this functional constraint, the linear Independent Component Analysis (ICA) is able to explain the shape of the receptive fields [31]. However, the linear nature of ICA prevents the complete achievement of the infomax goal, and residual statistical dependence does remain after linear ICA in natural images [32, 33].

Hyvärinen and Hoyer extended linear ICA by taking into account the residual statistical dependence after linear ICA. This is the core of the so called Topographic-ICA (TICA) [34]. In TICA, the correlations between the energies of first-stage linear filters are used to build a second-stage 2-dimensional cyclic topology (or grid) in which the processing units with stronger relations are closer in the grid. As a result, clusters of units tuned to similar features naturally emerge in this topology.

When applied to natural images, TICA certainly accounts for the basic features of both simple and complex cells in V1 [28], e.g. first-stage units have localized and oriented edge receptive fields (as the simple cells), and second-stage units have broader phase and location tuning (as complex cells). The authors also report the emergence of *orientation maps* derived from TICA, displaying iso-orientation domains and pinwheels, which appear again in [35, 36]. However, the orientation maps resulted in [28, 35, 36] may not be appropriate to explain orientation maps in V1 since they may lack the appropriate continuity and retinotopy.

The continuous orientation maps reported in [28, 35, 36] are found in the topology emerging from the algorithm (assumed to be the cortical space), but this topology can not be automatically transferred into the retinal space. Since the location in the TICA topology may not be automatically interpreted as retinal location, the TICA maps may not be equivalent to the biological maps (which are genuinely retinotopic and continuous). Hyvärinen & Hoyer do recognize that the retinotopy of the TICA topology is only local [28], but the implications of this fact were not discussed. No additional analysis has been done to confirm or refute the biological plausibility of the clusters emerging in the topology. In this work we elaborate on this open question: we argue that a biological interpretation of the predicted clusters would require a smoother retinotopic relation between the retinal space and the TICA topology, because, as

opposed to the animals with map patterns, the TICA orientation maps does not have an appropriate counterpart in the image domain with locally continuous structure.

Specifically, in the examples below we show that the topology reported in [28] (and sequels [35, 36]) violates the continuity of the orientation maps in such a way that, given two close units (artificially predicted neurons) in the TICA topology, it is easy to find another unit tuned to the same spatial position and an orthogonal orientation. Here, spatial position tuning means the center of the receptive field in the retinal space. These examples suggest non-continuous, or randomly scrambled, orientation domains when projecting them back to the visual space. In order to check this intuition, in this work we conduct a new analysis of the units emerging from TICA: we compute the spatial distribution of oriented units in the retinal space. The invariable result (over a wide set of training conditions) is a random salt-and-pepper organization, thus confirming the intuition obtained from the particular examples of continuity violation detected in the results of [28, 35, 36].

This salt-and-pepper organization of the TICA oriented filters in the retinal space implies that the locally continuous TICA topology does not have a locally continuous counterpart in the image space, and hence, it is not biologically plausible.

Examples of continuity and retinotopyviolation in the TICA literature

Fig 2 replots the basic result on the orientation domains from TICA in [28] (specifically Fig 4, and part of Fig 5 of the original paper), along with our annotations to discuss the interpretation of the maps. The left panel shows the TICA basis vectors, or processing units, ordered according to the topology that emerges from energy correlations. The top-right panel shows the location and orientation of the receptive fields represented at the left panel. These 2-d location and orientation colored diagrams were built by fitting a Gabor function to each unit in the position (i, j) in the left panel. Then, each position (i, j) in the topology (each pixel (i, j) in the diagram) is depicted using the appropriate color depending on the central location or fundamental orientation of the fitted Gabor function (and the code relating color-location and color-orientation). Finally, at the bottom right part of Fig 2, there is a cartoon of the retinal space to qualitatively illustrate the spatial tuning (or location) of the considered units in the visual field. The spatial distribution of the receptive field centers in this retinal space is the key in the discussion below. The left panel shows that the information theoretic approach by Hyvärinen & Hoyer shares interesting similarities with biological systems: the emerging units have Gabor-like receptive fields (as the V1 simple cells) and the smooth transition in location and frequency along the topology recalls the continuity of the retinotopic iso-orientation domains in the cortex. In fact, the orientation diagram certainly shows locally continuous orientation maps with distinct iso-orientation domains in the topology. The orientation map also displays pinwheels, marked by white circles in the original paper.

However, this interesting behavior may not be automatically applicable to V1 because the correspondence between the TICA topology and the cortical surface is not obvious. Note that the axes of the topology (rows and columns in the left panel, or row/column axes in the orientation diagram at the top-right panel) do not have a smooth spatial meaning as in V1. Note that one finds the same color, i.e. the same retinal location of the receptive field, at distant positions in the *location diagram*, or in the topology. This breaks the global retinotopy that has been described for V1 [4–11]. This result (local but not global retinotopy) was certainly acknowledged as an eventual limitation of TICA in the original paper [28]. However, the implications of this limited retinotopy were not discussed. And no additional analysis has been done to confirm or refute the biological plausibility of the maps emerging in the topology.




Fig 2. Topology and orientation map in Hyvärinen & Hoyer [28], together with our annotations highlighted with squares. The original source for the left panel (basis vectors in gray scale) and the colored location/orientation diagrams at the top right are Figs 4 and Fig 5b of the paper [28] (reprinted with permission of Elsevier). The highlighted functions with squares and the consequence in the retinal space are the key to question the biological plausibility of the result. The retinal space cartoon in the bottom right illustrates the location (central spatial tuning) of the units in the retina that are highlighted for the discussion (colored squares in the left panel and black and white squares in the diagrams). Here we argue that the interpretation of the blobs in the TICA orientation diagram as *orientation domains* is not possible because, when considering the location of the oriented units back in the retinal space, given two neighboring units tuned to one orientation (e.g. those in green or those in yellow), it is easy to find another unit in between, tuned to an orthogonal orientation (e.g. the one in red or the one in blue). This observation suggests that the TICA orientation domains may be randomly scrambled in the retinal space.

<https://doi.org/10.1371/journal.pone.0178345.g002>

Interestingly, a new look at the reported results suggests that the limited retinotopy of the topology may invalidate the biological plausibility of the reported orientation maps: as shown below, this is particularly relevant when considering the spatial distribution of the predicted sensory units in the retinal space (not only in the topology).

The following examples show that it is easy to find strong violations of the continuity of the orientation domains that suggest a random scrambling of the orientation map in the topology when projecting it back to the visual space. If these violations are frequent enough, the locally continuous orientation domains in the topology would disappear in the retinal space (as opposed to the experimental evidence reviewed in Fig 1).

In Fig 2 we highlight two of such cases: consider the sets of sensory units highlighted in the topology using (a) green and red squares, on the one hand, and (b) yellow and blue squares on the other hand. In the location and orientation diagrams of the top-right panels these units are highlighted using squares in (a) black, and (b) white, respectively.

In each case (a and b), the units are quite far in the topology, probably because their responses are almost independent, and they belong to distinct orientation domains. However, they are tuned to a similar spatial location (in the retinal space). The spatial overlapping of the receptive fields can be seen in the left panel, and also because the cells share the same color in the location diagram. Note that in the location diagram *yellow* corresponds to the lower left part of the considered visual field and *light-blue* corresponds to the upper-center region.

Actually, the center of the receptive field highlighted in red in the topology could be spatially located *in the visual field* between the two units highlighted in green (e.g. as illustrated in the qualitative retinal space cartoon in the bottom-right). The same happens in the other set of units, the case (b): the unit in blue may be surrounded by the two units in yellow. Moreover, the receptive fields of these units, despite being so spatially overlapped, are tuned to nearly orthogonal orientations, as can be seen in the left panel, and also by the very different color in the orientation diagram.

These examples suggest the rupture of the continuity within the orientation domains in the retinal space: (a) we have a unit tuned to 45° (in red) overlapped by two units tuned to 135° (in green); and (b) a unit tuned to 0° surrounded by two units tuned to 90° . If such mixtures are frequent enough, one foresees a salt-and-pepper distribution of the orientation tuning when plotted in the retinal space. This would not be compatible with the continuity of the orientation maps in the spatial domain reviewed above (Fig 1 and associated discussion in the text). We confirm this guess in the following section.

Qualitatively similar units and orientation maps are obtained using different neighborhood sizes when computing the correlations along the topology [36], or using overcomplete versions of TICA [35]. In both cases [35, 36], the same kind of strong continuity violations highlighted in Fig 2 can be found, so the same questions on the validity of the results arise.

2 Results: Topographic ICA does not lead to locally continuous maps in the retinal space

Here we check the distribution of the centers of the receptive fields of the TICA units at the retinal space depending on their orientation tuning, i.e. we study the TICA orientation domains back in the retinal space. This analysis, inspired by the empirical correspondence between a grid in the image space and its counterpart in V1 [8], was not done in [28, 35–37].

Given the straightforward relation between cortical location and spatial tuning in the visual field, Fig 1, a proper theory should lead to locally continuous orientation domains in the retinal space. However, the results in this section show a consistent salt-and-pepper distribution of the oriented TICA units for a variety of visual angles, training sets and sampling frequencies. The whole set of results can be reproduced by using the code and data available at: <http://isp.uv.es/code/visioncolor/TICAdomains/TICAdomains.html>, and at the public repository <https://github.com/neuronaValenciana/OrientationDomainsTICA>.

In order to check the consistency of the results at different visual scales, we considered different training sets spanning an order of magnitude in visual angle: 0.64, 0.80, 1.28, 2.00, 2.56, 4.00, 5.12, and 8.00 degrees. A wide range of visual angles (and hence a wide range of visual structures) is important in the context of the discussion on local but not global retinotopy in the original paper. See Fig 3 for an example of the kind of structures explored.

In these experiments, all the image samples come from the database in the ImageICA Toolbox [38]. However, since we used different sampling frequencies and block sizes to consider scenes subtending different visual angles, in practice, this means using totally separated training sets. First, we assumed an arbitrary baseline sampling frequency in the database ($f_s = 25$ cpd). We considered sets of image patches of different discrete sizes with this baseline rate: 16×16 , 20×20 , 32×32 , and 50×50 ; which implies visual angles of 0.64, 0.80, 1.28 and 2.00, degrees. Then, we undersampled the original images using $f_s = 12.5$ cpd and $f_s = 6.25$ cpd, and, in both cases, we took discrete sizes of 32×32 and 50×50 ; i.e. images with visual angles of 2.56, 4.00, 5.12, and 8.00 degrees. At each of the considered resolutions we gathered 20000 randomly selected patches. In each set we ran the Topographic ICA code in the ImageICA Toolbox [38] with default parameters.




Fig 3. The visual angles considered in our experiments span one order of magnitude: [0.64, 8.00] degrees. In red, the specific resolution considered in [28]. The new experiments largely expand the kind of visual structures considered in a single resolution 32 × 32 set.

<https://doi.org/10.1371/journal.pone.0178345.g003>

Fig 4 shows the central location of the receptive fields of the TICA units for the different datasets back in the retinal space. Each location is colored according to the orientation preferred by the unit. As an additional check, Fig 5 shows the distribution of the central frequencies of the receptive fields of the TICA units in the Fourier domain. Locations, frequencies and orientations were computed from the parameters of the corresponding Gabor functions fitted to the TICA units. While Fig 5 shows that the TICA units reasonably cover all possible orientations in the frequency range allowed by the considered sampling, Fig 4 confirms the random scrambling intuition anticipated by the observations made about Fig 2 in the Introduction. Qualitatively speaking, the salt-and-pepper-like distributions in Fig 4 do not seem compatible with locally continuous iso-orientation domains in the retinal space: results show a good coverage of the visual field with units tuned to the different orientations, but the units tuned to a similar orientation do not seem to be clustered in locally continuous regions. In the Methods

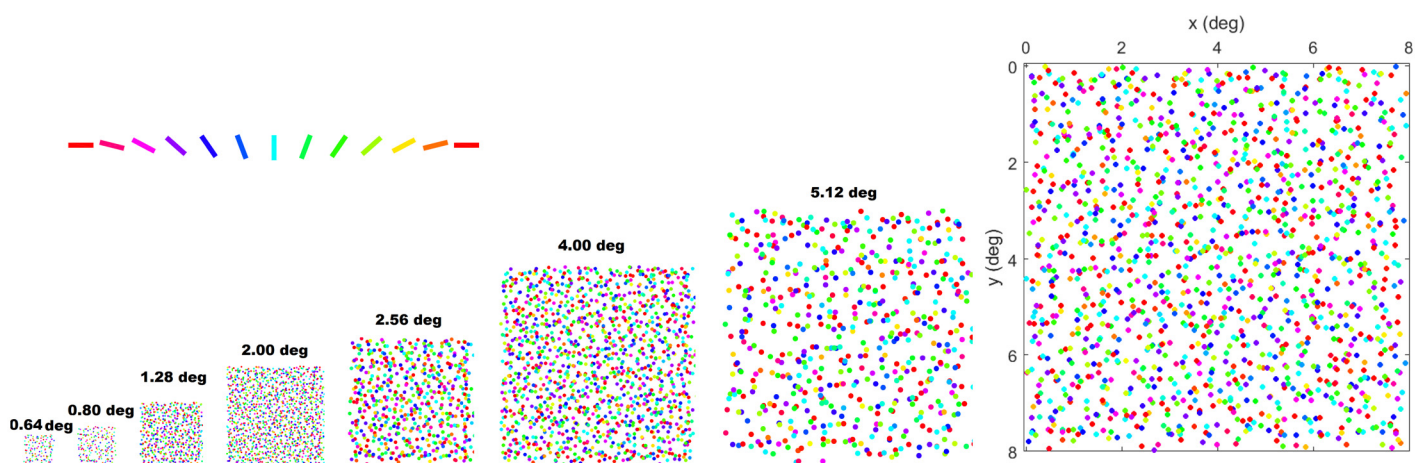


Fig 4. Centers of the receptive fields of TICA units at the retinal space colored depending on their preferred orientation. Results for the different visual angles considered in Fig 3. All the distributions are depicted in the same scale (spanning an order of magnitude in visual angle). Color code for orientation tuning is illustrated by the colored oriented bars at the top. Salt-and-pepper organization is consistently obtained over a range of visual angles, sampling frequencies and training sets.

<https://doi.org/10.1371/journal.pone.0178345.g004>

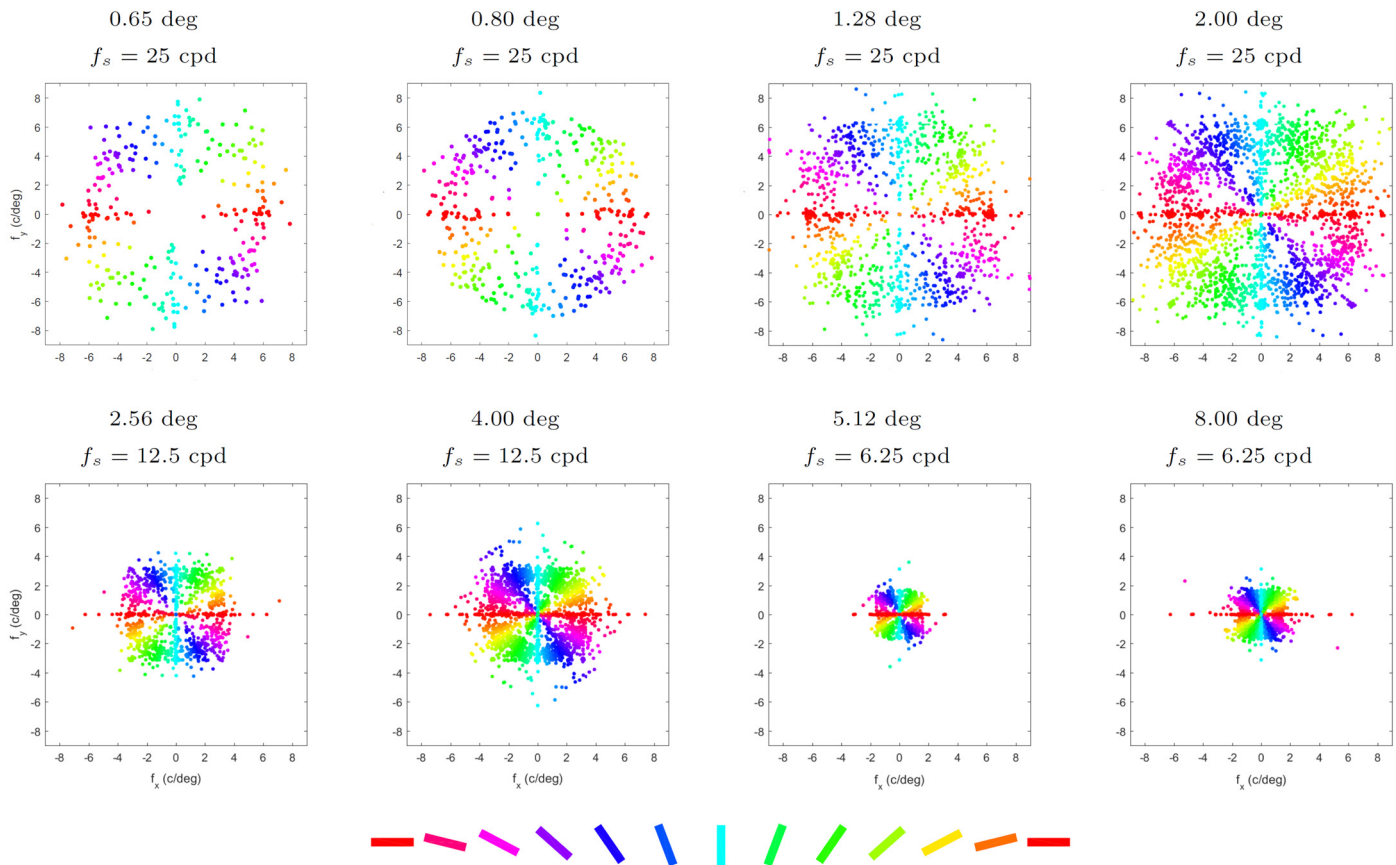


Fig 5. Central frequencies of the TICA units at the Fourier domain colored depending on their orientation tuning for the visual angles considered in Figs 3 and 4. Note that in datasets that have been undersampled the frequency band covered by the mechanisms is smaller as expected by the coarser visual structures due to the antialiasing filter. Color code for orientation is given by the colored bars.

<https://doi.org/10.1371/journal.pone.0178345.g005>

section we present a statistical test that quantitatively rejects the hypothesis of having locally continuous iso-orientation domains in the retinal space.

These salt-and-pepper results are not an artifact of lack of convergence of the iterative algorithm or a poor fit to Gabor functions. In the Methods section we show quantitative and visual illustrations of the proper convergence of the TICA units to Gabor-like functions. The salt-and-pepper distributions do not depend on specific TICA parameter setting either. Free parameters in TICA reduce to [34]: (1) the nonlinearity that is used to enforce the sparsity of the responses during the learning process, and (2) the size of the pooling area after the first stage, which determines the extent of the relations between linear responses. Default settings are [38]: (1) a square-root nonlinearity, and (2) a squared neighborhood of size 5×5 for image blocks of size 32×32 , the dashed square in Fig 2 illustrates this neighborhood. We explored different nonlinearities by varying the exponent in the range $[0.2, 0.8]$, and by considering other saturating functions such as *arc tangent*. We explored the effect of the size of the neighborhood by considering pooling regions of size 7×7 , 9×9 , and 11×11 for the same image size. For every tested nonlinearity and neighborhood we obtained non-continuous salt-and-pepper distributions of the mechanisms similar to those in Fig 4 (results not shown).

All the above suggests that salt-and-pepper arrangement of edge filters in the retinal space is an intrinsic property of the TICA model when trained on natural images.

In summary, the systematic salt-and-pepper results over a wide range of visual angles, data sets, and algorithm settings, show that the examples of rupture of the continuity of the TICA orientation domains back in the retinal space discussed in Fig 2 are extremely frequent. As a result, this prevents the emergence of locally continuous iso-orientation domains in the retinal space no matter the visual scale. The problem of TICA is that it does not have the degree of retinotopy of the biological systems like the one reviewed in Fig 1.

3 Discussion and concluding remarks

Interaction between similarly tuned neurons has functional advantages

Interaction between neurons tuned to similar features (scale, orientation) in nearby locations has a number of functional advantages, both from a lower-level information transmission perspective or from a higher-level pattern recognition perspective.

From the low-level efficient coding point of view, it is well known that natural scenes display characteristic regularities in linear local-frequency representations: statistical relations between linear responses decrease with the departure in preferred frequency and location in natural images [39–42], and the same happens in natural movies analyzed with spatio-temporal filters, where mutual information decreases with the departure in preferred speed [43]. In fact, the observation of these residual relations after linear ICA was one of the qualitative reasons to propose Topographic-ICA [34]. Interaction between mechanisms tuned to similar features (for instance using divisive normalization [44]) can be thought as a sort of predictive coding to remove these residual redundancies. In fact, the statistical interpretation of this interaction between mechanisms sharing features has been proved useful both to derive biological behavior from efficient image coding [32, 33], and the other way around: to derive efficient image coding from biological behavior [45].

From a higher-level, object recognition perspective, interaction between mechanisms tuned to similar orientations in similar locations have two important advantages. On the one hand, this interaction leads to contrast adaptation [44, 46], and hence, to contrast invariant edge detection which improves recognition performance. On the other hand, detection of large-scale contours requires checking the co-activation of neurons tuned to similar orientation at eventually distant points, therefore, it necessarily requires some sort of selective association between potentially collinear mechanisms [47, 48].

As a side comment, the considered cases of interaction (normalization within iso-orientation domains and association of collinear mechanisms), assume that the interaction takes place between mechanisms that already have a sharp orientation tuning. It is important to stress that the emergence of this sharp tuning may also come from feedback intracortical interactions [49, 50], and not only from the classical feedforward model of Hubel and Wiesel [51] based on the alignment of receptive fields of LGN cells afferent to cortical cells. In this context, approaches based on linear ICA, and also TICA, just assume the classical (and eventually limited) feedforward architecture from the retina [28, 31, 36] leading to a straightforward generation of V1 mechanisms as local linear combination of LGN cells [52]. Therefore, their results do not provide extra insight in the discussion between feedback vs feedforward architectures for the emergence of orientation selectivity.

Orientation domains simplify interaction between similarly tuned neurons

Interaction between neurons tuned to similar orientations and locations does happen in animals with [53, 54] and without [55] locally continuous orientation maps. The eventual difference is that in the case of having iso-orientation domains, the interaction would occur between close neighbors in the cortex while in the salt-and-pepper situation the interaction

pattern is necessarily more complex. This difference has been suggested as a possible explanation of the different complexity of the global connectivity between V1 layers in cats/monkeys versus mice [56].

Moreover, keeping the *mechanisms that should interact* together would probably reduce the wiring requirements when implementing the required interactions (either the divisive normalization for redundancy reduction or contrast invariance, or the association between collinear mechanisms for countour detection). This is consistent with the results in [57] which show that imposing connectivity between mechanisms with very similar tuning, the wiring minimization criterion leads to realistic orientation maps.

Are TICA maps related to the above functional advantages?

As opposed to other descriptions of the orientation maps which do not take into account the information processing function of the network [16–27, 57], TICA definitely takes it into account since it is based on information maximization and it considers the redundancies between the responses of the first-stage linear mechanisms. Broadly speaking, improved information transmission with regard to linear ICA and improved recognition performance may be functional arguments in favor of the TICA topology. However, the connection between TICA and these specific functional goals is only conceptual.

First, it is not clear that the reported advantages in object recognition when using TICA [58] actually come from the structured organization in the topology. It could also happen that the improved performance is due to the increased invariance of the mechanisms of the second stage in TICA. Unfortunately, the association pattern between the TICA mechanisms (that could be interesting from the biological point of view [47, 48]) was hidden in the machine learning classification stage in [58], so it is difficult to identify the advantage. Moreover, in any case, classification performance is not a part of the learning algorithm in TICA. Second, taking advantage of the eventually ordered structure for further redundancy reduction (e.g. using divisive normalization after TICA) is not a principle included in the TICA learning, so this *simplify-further-processing* argument could not be an explanation for the emergence of the structured maps either.

As a result, despite the functional origin of TICA through infomax, the emergence of structure in the topology in TICA is not directly explained by further redundancy reduction nor by higher-level object recognition arguments.

After all, as suggested by Wilson and Bednar in their excellent review [14], it may also happen that the emergence of the maps does not have a direct functional reason but it is a by-product of other (structural rather than functional) constraints. Nevertheless, despite the interest of the discussion about the usefulness of iso-orientation domains emerging from TICA, the central point of this paper is questioning the biological interpretation of the structure found in the TICA topology, not questioning its functional advantages.

Conclusion: Orientation maps generated by TICA do not reflect the retinotopy of primate V1 maps

In this work, we saw that a new look at the results reported in [28, 35, 36] suggests a random scrambling of the TICA oriented mechanisms when representing its location in the retinal space. This salt-and-pepper qualitative guess is quantitatively confirmed when systematically checking the organization of the TICA mechanisms in the retinal space (see the [Results](#) section). The problem is that TICA does not have the degree of retinotopy (correspondence between location in the topology and location in the retina) observed in the retina-V1 projection of animals with structured orientation maps. As a consequence, the local continuity of the

orientation domains in the TICA topology is totally broken back in the image space, so that the hypothesis of existence of a clustered structure there can be reliably rejected (see the Methods section). Therefore, structured orientation maps in the V1 cortex of primates and other mammals cannot be interpreted in terms of topologies generated by TICA.

Given the fact that salt-and-pepper organization in rodents is still a puzzle and a challenge for other mechanistic models, one could think that the ability of TICA to generate this kind of organization could be of some interest to understand those other species. However, the negative result for TICA in primates can not have a positive interpretation for rodents because the TICA retinal organization is salt-and-pepper, but the “cortical organization” in TICA is structured and locally continuous, which is not the case in mice [13].

4 Materials and methods

Here we address two methodological issues: (1) the statistical test to show that the set of TICA mechanisms tuned to certain orientation is not clustered in the retinal space as it would be in the case of a structured organization in locally continuous iso-orientation domains; and (2) the appropriate convergence of TICA in our numerical experiments.

Statistical test of the distributions

The test proposed here is focused on pointing out the uniformity of the set of TICA processing units tuned to certain orientation (the set S_{ori}). Specifically, it addresses the following question: is S_{ori} similar to a Cartesian grid? If the answer is yes, we are rejecting the presence of clusters in S_{ori} since a Cartesian grid (S_{cart}) is the less clustered organization one can imagine. In order to use this test in practice, one has to devise a similarity measure and a meaningful reference. As a reference we chose another set lacking of spatial structure: an equivalent finite set drawn from a 2D uniform density function, let's call that S_{rand} . Then, in particular, we check if S_{ori} is more similar to S_{cart} than S_{rand} . This question has three possible outcomes: (1) less similar, (2) equally similar, or (3) more similar. If the answer is more or equally similar, S_{ori} displays no significant clustered structure since its uniformity would be similar or bigger than the uniformity of finite sets drawn from a uniform PDF. In the 1st case we could not extract any conclusion about S_{ori} spatial organization.

The proposed test compares the organization of a set of mechanisms (either the considered S_{ori} or a set of the same size drawn from a 2-D uniform distribution, S_{rand}) against a regular Cartesian grid of the same average density, S_{cart} . Fig 6 shows a realization of these sets, S_{ori} , S_{rand} , and S_{cart} . The similarity measure with S_{cart} could take into account how many samples of the considered set are close to each node of the Cartesian grid. A convenient dissimilarity measure is the *Kullback Leibler* divergence (*KLD*) between the histograms [59].

The example in Fig 6 shows, that in almost every bin in the histogram $h(S_{ori})$ there is a TICA mechanisms, similarly to S_{cart} . In contrast, S_{rand} has many empty bins as well many highly populated bins. Note the black and white deviations in the histogram $h(S_{rand})$.

The *KLD* measures how the histograms ($h(S_{ori})$ and $h(S_{rand})$) depart from the constant value of $h(S_{cart})$. Note that in order to have a lower bound for the divergence, $KLD = 0$, its necessary that the number of samples at the Cartesian grid (and bins of the histograms, $N_1 \times N_2$) its close to the number of samples at the sets S_{rand}/S_{ori} , call it M .

Here it is how the test works. Given certain set S_{ori} of size M , we build $N_1 \times N_2$ Cartesian grid that uniformly tiles the retinal plane, where $N_1 \cdot N_2$ is the best integer approximation to M , ($N_1, N_2 \in \mathbb{N}$). Given the fact that the orientation bandwidth of the mechanisms is in the range $[20^\circ, 45^\circ]$, the literature assumes among 4–6 separate sets of orientated filters. Here we used three different partitions of the continuous $[0, 180^\circ]$ range: $K = 4, 5, 6$, since these




Fig 6. Up left, an example of the distribution of oriented TICA mechanisms in the retinal space, S_{ori} , for certain orientation. Each point corresponds to the center of a mechanism tuned to $135^\circ \pm 22^\circ$. In this example patch size is 32×32 , $f_s = 12.5$, and $K = 4$. Down left is the histogram of the above set. Middle column, an example of a 2-D random sample uniformly distributed, S_{rand} , with the same number of points than S_{ori} , and its corresponding histogram. Right column, an example of Cartesian grid (S_{cart}) and its histogram. The grid of the top figures shows the bins used to build the histograms. In histograms, black corresponds to empty bins while more populated bins are depicted in white.

<https://doi.org/10.1371/journal.pone.0178345.g006>

correspond to bandwidths of 45° , 36° , 30° respectively. Then, we generate 10^4 realizations of S_{rand} of size M drawn from a 2-D uniform distribution defined in the same range, and we measure the divergence between the $h(S_{rand})$ of these realizations and $h(S_{cart})$. This allows us to estimate the PDF of the divergence measure, $P(KLD(S_{rand}, S_{cart}))$. Finally the divergence $KLD(S_{ori}, S_{cart})$ is calculated and compared with $P(KLD(S_{rand}, S_{cart}))$.

The outcomes of the test are as follows: (1) if $KLD(S_{ori}, S_{cart})$ is significantly bigger than the mean of $P(KLD(S_{rand}, S_{cart}))$ no clear conclusion on the structure can be obtained. (2) If $KLD(S_{ori}, S_{cart})$ is similar to the mean of $P(KLD(S_{rand}, S_{cart}))$, S_{ori} would behave similarly as S_{rand} , thus being incompatible with the clustered structure of orientation domains. Finally, (3) if $KLD(S_{ori}, S_{cart})$ falls at the left tail, S_{ori} is clearly more similar to S_{cart} than S_{rand} . In this case, the null hypothesis (uniform distribution) can be rejected, since S_{ori} is actually more uniform than S_{rand} and even more incompatible with orientation domains. Fig 7 are a representative illustration of the results of the test (similar behavior is obtained in the other cases). In these plots $KLD(S_{ori}, S_{cart})$ (a red cross in the Fig 7) falls below the mean $P(KLD(S_{rand}, S_{cart}))$ distribution, so at least we are always in the second case. Even more, in most cases $KLD(S_{ori}, S_{cart})$ falls below the 5% left tail of distribution. This is true regardless of the (arbitrary) number of selected partitions, K , patch size and sampling frequency.

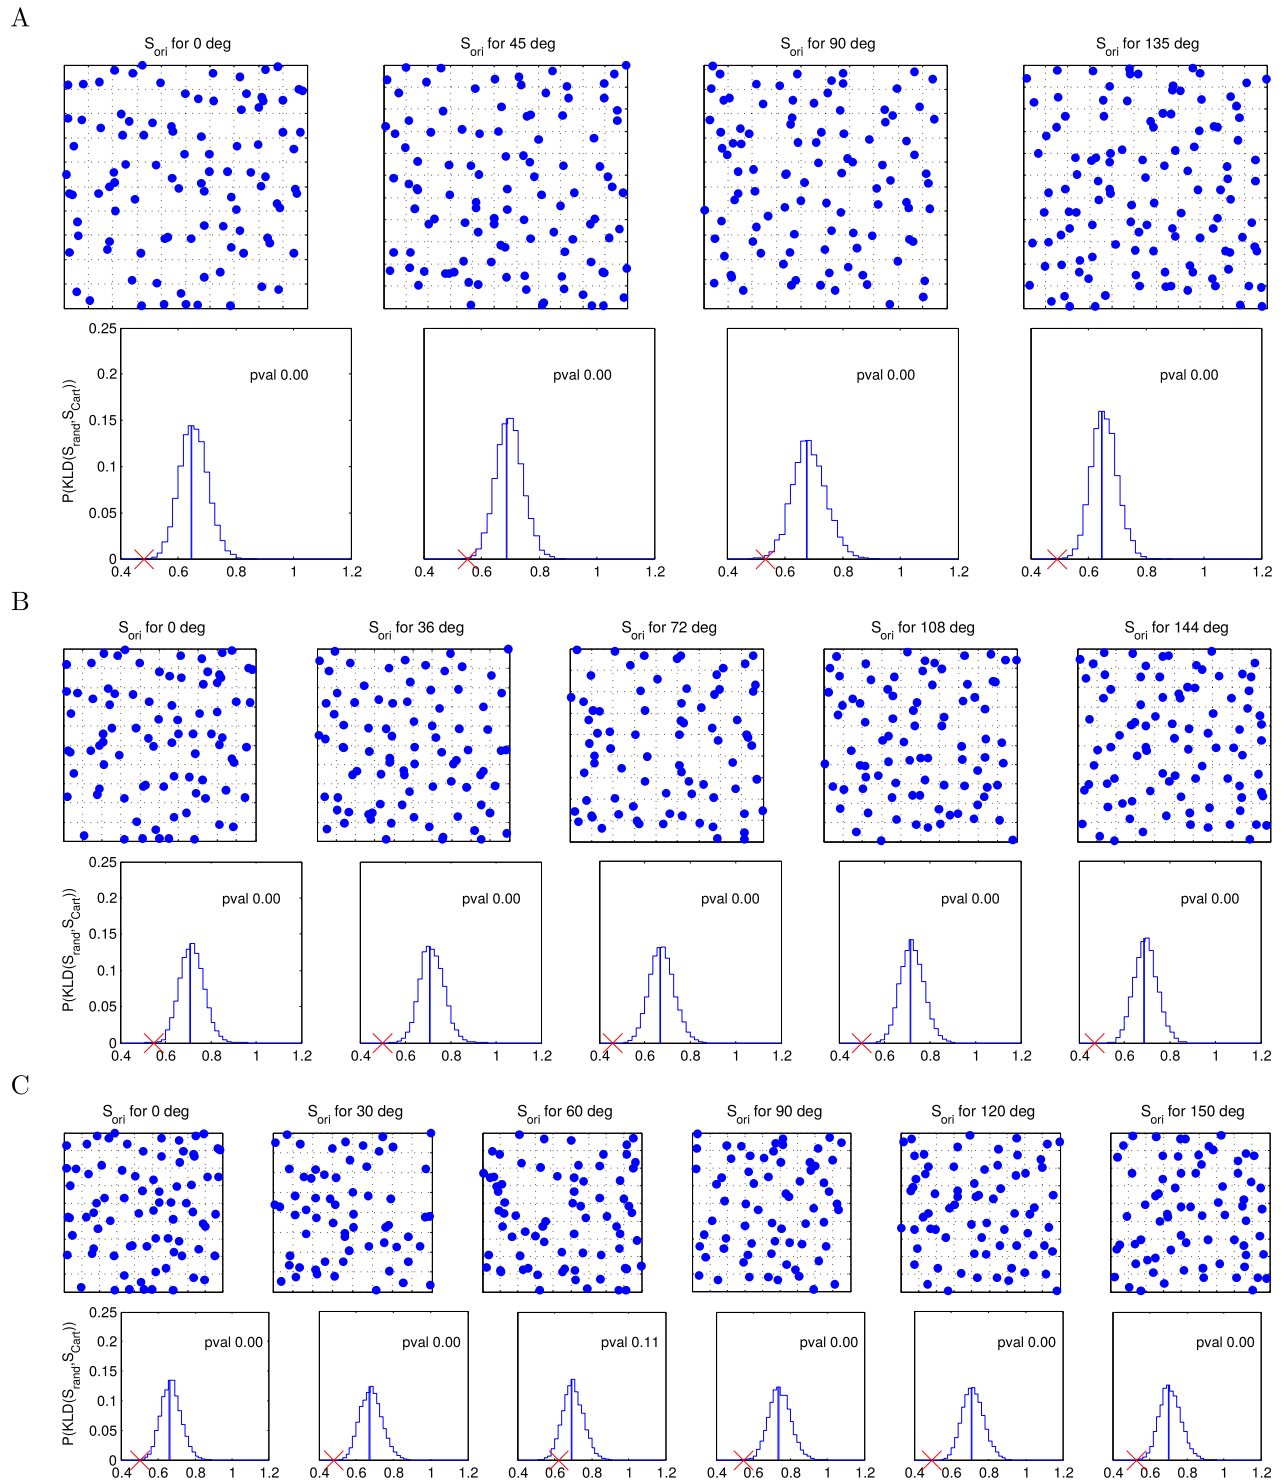


Fig 7. Locations of the TICA sensory units in the retinal space depending on their preferred orientation (sets S_{ori}), and uniformity test results. In order to point out the consistency of the result regardless of the arbitrary partition of the orientation range, here we apply the test splitting the data (the sensory units) according to three different partitions of the $[0, 180^\circ]$ range. We considered different number, K , of preferred orientations: $K = 4$ (top panel A), $K = 5$ (middle panel B), and $K = 6$ (bottom panel C). Image data in this example correspond to scenes subtending 1.28 degrees sampled at 25 cpd. Below each set S_{ori} we represent the corresponding $P(KLD(S_{rand}, S_{cart}))$ distribution, a red cross shows the value of $(KLD(S_{ori}, S_{cart}))$, and the number displays the corresponding p-value. The blue bars at the distributions show the mean interval of confidence $[\mu \pm 2.576 \cdot SEM]$ (standard error of the mean).

<https://doi.org/10.1371/journal.pone.0178345.g007>




Fig 8. Errors of the fits to Gabor functions (left), and evolution of the TICA goal function (right). In the left plot, features were sorted according to the normalized MAE of the fit to a Gabor function. Different visual angles and resolutions lead to topographies of different size (different number of features). The values of the goal functions during the training of TICA [38] were normalized by the initial value in each case. Curves in red and green correspond to full resolution or to spatially undersampled datasets respectively. Note that normalized MAE is small for most of the features in every dataset. For example more than 95% of the features have a normalized error below 12.5%. Note also that the goal curves in red and green achieve a plateau after 2000 iterations indicating a stable result.

<https://doi.org/10.1371/journal.pone.0178345.g008>

This quantitatively confirms the main result (Fig 4): TICA oriented filters display no specific spatial structure in the retinal space, so TICA does not lead to structured orientation maps there.

Convergence and quality of fit to Gabor-functions

We achieved satisfactory convergence of the TICA sensory units in 2000 iterations of the algorithm in all the considered training sets. The quality of the convergence and the resemblance of TICA units to Gabor functions is quantitatively and visually illustrated in Figs 8 and 9 respectively. Each (nonparametric) TICA basis function was fitted to a Gabor function, and Fig 8 (left) shows the errors of these fits for all the considered datasets. The departure of the TICA basis function from the corresponding Gabor function is described here by the Mean Absolute Error normalized by the range of values of the TICA receptive field. We refer to this as *normalized Mean Absolute Error* (normalized MAE). Fig 8 (right) shows the evolution of the goal function optimized in TICA [34, 38] in our specific experiments for the different datasets. Appropriate convergence is indicated by two facts: (1) a plateau in the goal function minimized when training TICA, and (2) the emerging receptive fields resemble Gabor functions.

As a convenient reference, the highlighted threshold in the error in Fig 8 (12.5% in normalized MAE), means pretty good agreement with a Gabor function in practice (see the visual examples in Fig 9, which include only functions below this threshold). In all the experiments (main result, Fig 4, and the corresponding uniformity tests) we did not include TICA units with normalized MAE bigger than 12.5%. Most of the TICA receptive fields were well fitted by Gabor functions, less than 5% were not well fitted, so the location, frequency and orientation values in the main result are reliable.




Fig 9. Illustrative results of the TICA basis functions (left) and the corresponding fitted Gabor functions (right) for the training sets corresponding to the visual angles of 0.64 deg (top), and 2 deg (bottom) sampled at $f_s = 25$ cpd. As can be seen in this example, non-Gabor functions (depicted as flat functions at the right panels) are rare (less than 5%) so fitted parameters are reliable. Qualitatively similar results were obtained for the other datasets.

<https://doi.org/10.1371/journal.pone.0178345.g009>

Acknowledgments

The authors acknowledge the financial support received from the Spanish Ministry of Economy and Competitiveness (MINECO) under the grants TIN2015-71130-REDT, BFU2014-58776-R, TIN2013-50520-EXP, co-financed by the European Regional Development Fund (ERDF), and the Severo Ochoa Program for Centers of Excellence in R+D (SEV-2013-0317).

The authors thank the fruitful discussion with Aapo Hyvärinen before the submission of this work.

Author Contributions

Conceptualization: Luis M. Martinez, Jesús Malo.

Data curation: Marina Martinez-Garcia, Jesús Malo.

Formal analysis: Marina Martinez-Garcia, Jesús Malo.

Funding acquisition: Luis M. Martinez, Jesús Malo.

Investigation: Marina Martinez-Garcia, Jesús Malo.

Methodology: Marina Martinez-Garcia, Jesús Malo.

Project administration: Luis M. Martinez, Jesús Malo.

Resources: Luis M. Martinez, Jesús Malo.

Software: Marina Martinez-Garcia, Jesús Malo.

Validation: Marina Martinez-Garcia, Luis M. Martinez, Jesús Malo.

Visualization: Marina Martinez-Garcia, Jesús Malo.

Writing – original draft: Marina Martinez-Garcia, Luis M. Martinez, Jesús Malo.

Writing – review & editing: Marina Martinez-Garcia, Luis M. Martinez, Jesús Malo.

References

1. Hubel DH, Wiesel TN. Receptive fields and functional architecture of monkey striate cortex. *J Physiol.* 1968; 195:215–243. <https://doi.org/10.1113/jphysiol.1968.sp008455> PMID: 4966457
2. Ringach DL. Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *J Neurophys.* 2002; 88:455–463. PMID: 12091567
3. Schuett S, Bonhoeffer T, Hübener M. Mapping Retinotopic Structure in Mouse Visual Cortex with Optical Imaging. *Journal of Neuroscience.* 2002; 22(15):6549–6559. PMID: 12151534
4. Tootell RBH, Silverman MS, Switkes E, DeValois RL. Deoxyglucose Analysis of Retinotopic Organization in Primate Striate Cortex. *Science.* 1982; 218(4575):902–904. <https://doi.org/10.1126/science.7134981> PMID: 7134981
5. Tootell RBH, Switkes E, Silverman MS, Hamilton SL. Functional Anatomy of Macaque Striate Cortex II: Retinotopic Organization. *J Neurosci.* 1988; 8(5):1531–1568. PMID: 3367210
6. Tootell RBH, Hadjikhani NK, Vanduffel W, Liu AK, Mendola JD, Sereno MI, et al. Functional Analysis of Primary Visual Cortex in Humans. *PNAS.* 1998; 95(2):811–817. <https://doi.org/10.1073/pnas.95.3.811> PMID: 9448245
7. Das A, Gilbert CD. Distortions of visuotopic map match orientation singularities in primary visual cortex. *Nature.* 1997; 387(5):594–598. <https://doi.org/10.1038/42461> PMID: 9177346
8. Bosking W, Crowley J, Fitzpatrick D. Spatial coding of position and orientation in primary visual cortex. *Nature Neurosci.* 2002; 5(9):874–882. <https://doi.org/10.1038/nn908> PMID: 12195429
9. Buzas P, Volgushev M, Eysel UT, Kisvarday ZF. Independence of visuotopic representation and orientation map in the visual cortex of the cat. *Europ J Neurosci.* 2003; 18:957–968. <https://doi.org/10.1046/j.1460-9568.2003.02808.x> PMID: 12925022

10. Ohki K, Chung S, Kara P, Hubener M, Bonhoeffer T, Reid R. Highly ordered arrangement of single neurons in orientation pinwheels. *Nature*. 2006; 442:925–928. <https://doi.org/10.1038/nature05019> PMID: 16906137
11. Frazor RA, Benucci A, Carandini M. Independent Encoding of Position and Orientation by Population Responses in Primary Visual Cortex. *Proc Brain, Vis Artif Intell LNCS*. 2007; 4729:30–41. https://doi.org/10.1007/978-3-540-75555-5_4
12. Albus K. A quantitative study of the projection area of the central and the paracentral visual field in area 17 of the cat. I. The precision of the topography. *Exp Brain Res*. 1975; 24:159–179. <https://doi.org/10.1007/BF00234061> PMID: 1218549
13. Ohki K, Chung S, Chang Y, Kara P, Reid R. Micro-architecture of visual cortex: functional maps with single-cell precision. *Nature*. 2005; 433:597–603.
14. Wilson SP, Bednar JA. What, if anything, are topological maps for? *Developmental Neurobiology*; 75 (6):667–681. <https://doi.org/10.1002/dneu.22281> PMID: 25683193
15. Kaschube M. Neural maps versus salt-and-pepper organization in visual cortex. *Current Opinion in Neurobiology*. 2014; 24:95–102. <https://doi.org/10.1016/j.conb.2013.08.017> PMID: 24492085
16. Ringach DL. On the Origin of the Functional Architecture of the Cortex. *PLoS ONE*. 2007; 2(2):e251. <https://doi.org/10.1371/journal.pone.0000251> PMID: 17330140
17. Paik SB, Ringach DL. Retinal origin of orientation maps in visual cortex. *Nature Neuroscience*. 2011; 14:919–925. <https://doi.org/10.1038/nn.2824> PMID: 21623365
18. Hore VRA, Troy JB, Eglen SJ. Parasol cell mosaics are unlikely to drive the formation of structured orientation maps in primary visual cortex. *Visual Neuroscience*. 2012; 29(6):283–299. <https://doi.org/10.1017/S0952523812000338> PMID: 23110776
19. Schottdorf M, Eglen SJ, Wolf F, Keil W. Can Retinal Ganglion Cell Dipoles Seed Iso-Orientation Domains in the Visual Cortex? *PLoS ONE*. 2014; 9(1):e86139. <https://doi.org/10.1371/journal.pone.0086139> PMID: 24475081
20. Obermayer K, Ritter H, Schulten K. A principle for the formation of the spatial structure of cortical feature maps. *PNAS*. 1990; 87(21):8345–8349. <https://doi.org/10.1073/pnas.87.21.8345> PMID: 2236045
21. Obermayer K, Blasdel G, Schulten K. Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps. *Phys Rev A*. 1992; 45:7568–7589. <https://doi.org/10.1103/PhysRevA.45.7568> PMID: 9906831
22. Bednar JA. Building a mechanistic model of the development and function of the primary visual cortex. *Journal of Physiology-Paris*. 2012; 106(5–6):194–211. <https://doi.org/10.1016/j.jphysparis.2011.12.001> PMID: 22343520
23. Linsker R. From basic network principles to neural architecture: Emergence of orientation columns. *PNAS*. 1986; 83(11):8779–8783. <https://doi.org/10.1073/pnas.83.22.8779> PMID: 3464981
24. Swindale NV. The development of topography in the visual cortex: a review of models. *Network: Comp Neur Syst*. 1996; 7:161–247. <https://doi.org/10.1088/0954-898X/7/2/002> PMID: 16754382
25. Kaschube M, Schnabel M, Löwel S, Coppola D, White L, Wolf F. Universality in the evolution of orientation columns in the visual cortex. *Science*. 2010; 330:1113–1116. <https://doi.org/10.1126/science.1194869> PMID: 21051599
26. Meng Y, Tanaka S, Poon S. Comment on Universality in the evolution of orientation columns in the visual cortex. *Science*. 2012; 336:413–c. <https://doi.org/10.1126/science.1205737> PMID: 22539703
27. Keil W, Kaschube M, Schnabel M, Kisvarday Z, Löwel S, Coppola D, et al. Response to Comment on Universality in the evolution of orientation columns in the visual cortex. *Science*. 2012; 336:413–d.
28. Hyvärinen A, Hoyer PO. A two-layer sparse coding model learns simple and complex cell receptive fields and topography from natural images. *Vis Res*. 2001; 41(18):2413–2423. PMID: 11459597
29. Bell AJ, Sejnowski TJ. An information-maximization approach to blind separation and blind deconvolution. *Neural Comput*. 1995; 7(6):1129–1159. <https://doi.org/10.1162/neco.1995.7.6.1129> PMID: 7584893
30. Cardoso J. Dependence, correlation and Gaussianity in independent component analysis. *J Mach Learn Res*. 2003; 4:1177–1203.
31. Olshausen BA, Field D. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*. 1996; 381:607–609. <https://doi.org/10.1038/381607a0> PMID: 8637596
32. Schwartz O, Simoncelli EP. Natural signal statistics and sensory gain control. *Nature Neurosci*. 2001; 4(8):819–825. <https://doi.org/10.1038/90526> PMID: 11477428
33. Malo J, Gutiérrez J. V1 non-linear properties emerge from local-to-global non-linear ICA. *Network: Computation in Neural Systems*. 2006; 17(1):85–102. <https://doi.org/10.1080/09548980500439602> PMID: 16613796
34. Hyvärinen A, Hoyer PO, Inki M. Topographic ICA. *Neural Comp*. 2001; 13(7):1527–1558.

35. Ma L, Zhang L. Overcomplete topographic independent component analysis. *Neurocomputing*. 2008; 71:2217–2223.
36. Hyvärinen A, Hurri J, Hoyer PO. *Natural Image Statistics: a probabilistic approach to early computational vision*. Springer; 2009.
37. Hyvärinen A, Hoyer PO. Topographic independent component analysis as a model of V1 organization and receptive fields. *Neurocomputing*. 2001; 38–40:1307–1315.
38. Hyvärinen A. *ImageICA Matlab Toolbox*; 2001. <http://research.ics.aalto.fi/ica/imageica/>.
39. Liu J, Moulin P. Information-theoretic analysis of interscale and intrascale dependencies between image wavelet coefficients. *IEEE Trans Im Proc*. 2001; 10:1647–1658. <https://doi.org/10.1109/83.967393> PMID: 18255507
40. Buccigrossi RW, Simoncelli EP. Image compression via joint statistical characterization in the wavelet domain. *IEEE Trans Im Proc*. 1999; 8:1688–1701. <https://doi.org/10.1109/83.806616> PMID: 18267447
41. Malo J, Epifanio I, Navarro R, Simoncelli EP. Nonlinear image representation for efficient perceptual coding. *IEEE Trans Im Proc*. 2006; 15(1):68–80. <https://doi.org/10.1109/TIP.2005.860325> PMID: 16435537
42. Laparra V, Gutiérrez J, Camps-Valls G, Malo J. Image denoising with kernels based on natural image relations. *J Mach Learn Res*. 2010; 11:873–903.
43. Laparra V, Malo J. Visual aftereffects and sensory nonlinearities from a single statistical framework. *Front Human Neurosci*. 2015; 9:557. <https://doi.org/10.3389/fnhum.2015.00557> PMID: 26528165
44. Carandini M, Heeger DJ. Normalization as a canonical neural computation. *Nature Rev Neurosci*. 2012; 13(1):51–62. <https://doi.org/10.1038/nrn3136>
45. Malo J, Laparra V. Psychophysically tuned divisive normalization approximately factorizes the PDF of natural images. *Neural Computation*. 2010; 22(12):3179–3206. https://doi.org/10.1162/NECO_a_00046 PMID: 20858127
46. Watson AB, Solomon JA. A model of visual contrast gain control and pattern masking. *JOSA A*. 1997; 14:2379–2391. <https://doi.org/10.1364/JOSAA.14.002379> PMID: 9291608
47. Field DJ, Hayes A, Hess RF. Contour integration by the human visual system: evidence for a local “association field”. *Vision Res*. 1993; 33:173–193. [https://doi.org/10.1016/0042-6989\(93\)90156-Q](https://doi.org/10.1016/0042-6989(93)90156-Q) PMID: 8447091
48. Polat U, Sagi D. Spatial interactions in human vision: from near to far via experience-dependent connections. *Proc Natl Acad Sci USA*. 1994; 91:12006–12009. <https://doi.org/10.1073/pnas.91.4.1206> PMID: 8108388
49. Ben-Yishai R, Bar-Or L, H S. Theory of orientation tuning in visual cortex. *PNAS*. 1995; 92:3844–3848. <https://doi.org/10.1073/pnas.92.9.3844> PMID: 7731993
50. Lampl I, JF A, D G, Ferster D. Prediction of Orientation Selectivity from Receptive Field Architecture in Simple Cells of Cat Visual Cortex. *Neuron*. 2001; 30(1):263–274. [https://doi.org/10.1016/S0896-6273\(01\)00278-1](https://doi.org/10.1016/S0896-6273(01)00278-1) PMID: 11343660
51. Hubel DH, Wiesel TN. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *J Physiol*. 1962; 160(1):106–154. PMID: 14449617
52. Martinez-Garcia M, Galan B, Martinez LM, Malo J. Towards a functional explanation of the connectivity LGN-V1. In: *Comput. Neurosci. Comm., Proc. MODVIS*; 2016.
53. Gilbert CD, Li W. Adult Visual Cortical Plasticity. *Neuron*. 2012; 75:250–264. <https://doi.org/10.1016/j.neuron.2012.06.030> PMID: 22841310
54. Stettler DD, Das A, Bennett J, Gilbert CD. Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neurons*. 2002; 36:739–750. [https://doi.org/10.1016/S0896-6273\(02\)01029-2](https://doi.org/10.1016/S0896-6273(02)01029-2) PMID: 12441061
55. Cossell L, Iacaruso MF, Muir DR, Houlton R, Sader EN, Ko H, et al. Functional organization of excitatory synaptic strength in primary visual cortex. *Nature*. 2015; 518:399–403. <https://doi.org/10.1038/nature14182> PMID: 25652823
56. Frégnac B Y and Bathellier. Cortical Correlates of Low-Level Perception: From Neural Circuits to Percepts. *Neuron*. 2015; 88:110–126. <https://doi.org/10.1016/j.neuron.2015.09.041> PMID: 26447576
57. Koulakov A, Chklovskii DB. Orientation Preference Patterns in Mammalian Visual Cortex: A Wire Length Minimization Approach. *Neuron*. 2001; 29 (2):519–527. [https://doi.org/10.1016/S0896-6273\(01\)00223-9](https://doi.org/10.1016/S0896-6273(01)00223-9) PMID: 11239440
58. Kavukcuoglu K, Ranzato MA, Fergus R, LeCun Y. Learning invariant features through topographic filter maps. *Proc IEEE Conf Comp Vis Patt Recogn* 2009. 2009; 1:1605–1612. <https://doi.org/10.1109/CVPR.2009.5206545>
59. Cover TM, Thomas JA. *Elements of Information Theory*. 2nd ed. Wiley; 2006. <https://doi.org/10.1002/047174882X>