LETTER Communicated by Steven Zucker

Local and Global Gestalt Laws: A Neurally Based Spectral Approach

Marta Favali
marta.favali@ehess.fr
Center of Mathematics, CNRS - EHESS, Paris, 75244, France

Giovanna Citti
giovanna.citti@unibo.it
Dipartimento di Matematica, Università di Bologna, Bologna, 40126, Italy

Alessandro Sarti
alessandro.sarti@ehess.fr
Center of Mathematics, CNRS - EHESS, Paris, 75244, France

This letter presents a mathematical model of figure-ground articulation that takes into account both local and global gestalt laws and is compatible with the functional architecture of the primary visual cortex (V1). The local gestalt law of good continuation is described by means of suitable connectivity kernels that are derived from Lie group theory and quantitatively compared with long-range connectivity in V1. Global gestalt constraints are then introduced in terms of spectral analysis of a connectivity matrix derived from these kernels. This analysis performs grouping of local features and individuates perceptual units with the highest salience. Numerical simulations are performed, and results are obtained by applying the technique to a number of stimuli.

1 Introduction

Gestalt laws have been proposed to explain several phenomena of visual perception, such as grouping and figure-ground segmentation (Wertheimer, 1938; Kohler, 1929; Koflka, 1935; for a recent review, see Wagemans et al., 2012). In order to individuate perceptual units, gestalt theory introduced local and global laws. Among the local laws, we recall the principle of proximity, similarity, and good continuation. The local law of good continuation plays a central role in perceptual grouping (see Figure 1, top row, left)).

Regarding global laws, in the construction of percepts, the feature of saliency is crucial, yet it is not easy to model quantitatively. In Berliner Gestal theory, the concept of salience denotes the relevance of a form with respect to a contextual frame, that is, the power of an object to be present in the visual field. The role of salience is also pivotal in figure-ground

Neural Computation **29, 394–422** (2017)
doi:10.1162/NECO_a_00921

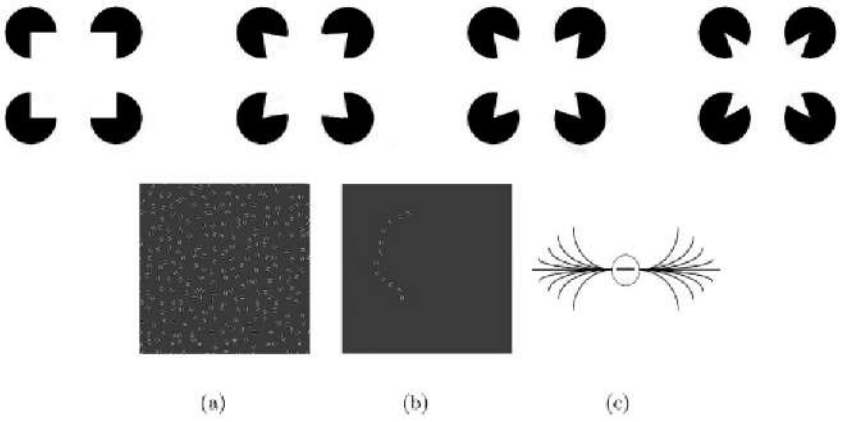

Figure 1: (Top) Deformation of visual stimulus, represented by squares with different angles between the inducers. The angle regularly decreases, and we perceive regular deformations of the subjective Kanizsa square up to a certain value of curvature, when the square suddenly disappears and the inducers are perceived. (Bottom) The stimulus proposed by Field, Hayes, and Hess (1993) (a) and the perceptual unit present in it (b). (c) The field lines of the association field, represent the elements in the path that can be associated with the central point (Field et al., 1993).

articulation. Due to the perceptual grouping process, scenes are perceived as constituted by a finite number of figures, and the salience assigns a discrete value to each of them. The most salient configuration pops up from the ground and becomes a figure (Merleau-Ponty, 1945). Note that in the case of continuous deformation of the visual stimulus, the salient figures can change abruptly from one percept to another (Merleau-Ponty, 1945). This happens, for example, in the top row of Figure 1, where a regular deformation is applied to the Kanizsa square: we progressively perceive a more curved square, until it suddenly disappears and the four inducers are perceived as standing alone (see Lee & Nguyen, 2001; Pillow & Nava, 2002; Petitot, 2008).

A number of results have been provided in order to refine the principles of psychology of form and assess neural correlates of good continuation law. In particular, Grossberg and Mingolla (1985) introduced a "cooperation field" to model illusory contour formation. Similar fields of association and perceptual grouping have been produced by Parent and Zucker (1989). In the 1990s, Kellman and Shipley provided a theory of object perception that adressed the perception of partially occluded objects and illusory contours (Kellman & Shipley, 1991; Shipley & Kellman, 1992, 1994). Heitger and von Der Heydt (Von Der Heydt, Heitger, & Peterhans, 1993) provided a theory of figural completion that can be applied to illusory contour figures (as the

Kanizsa triangle) and real images. Field et al. (1993) introduced through psychophysical experiments the notion of association fields, describing the Gestalt principle of good continuation. They studied how the perceptual unit visualized in Figure 1b pops up from a stimulus of Gabor patches (see Figure 1a). Through a series of similar experiments, they constructed an association field that defines the pattern of position-orientation elements of stimuli that can be associated with the same perceptual unit (see Figure 1c).

Starting from the classical results of Hubel and Wiesel (1977), it has been possible to justify these perceptual phenomena on neurophysiological bases. The results of Bosking, Zhang, Schofield, and Fitzpatrick (1997) and Frégnac and Shulz (1999) confirmed that neurons sensitive to similar orientation are preferentially connected. This suggests that the rules of proximity and good continuation are implemented in the horizontal connectivity of low-level visual cortices. A stochastic model that takes into account the structure of the cortex, with position an orientation feature, was proposed by Mumford (1994), and further exploited by Williams and Jacobs (1997) and August and Zucker (2000). They modeled similar fields with Fokker-Planck equations, taking into account different geometric features, such as orientation and curvature. Petitot and Tondut (1999) introduced a model of the functional architecture of V1, compatible with the association field. Citti and Sarti (2006) proposed the model of functional architecture as a Lie group, showing the relation between geometric integral curves, association fields, and cortical properties. This method has been implemented in Sanguinetti, Citti, and Sarti (2008) and Boscain, Duplaix, Gauthier, and Rossi (2012). An exact solution of the Fokker-Planck equation has been provided by Duits and van Almsick (2008), and their results have been applied by Duits and Franken (2009) to image processing.

The local laws are insufficient to explain the constitution of a percept, since a perceived form is characterized by global consistency. Different authors have qualitatively defined this consistency as pregnancy or global saliency (Merleau-Ponty, 1945), but only a few quantitative models have been proposed (Koch & Ullman, 1985). In particular, a spectral approach for image processing was proposed by Perona and Freeman (1998), Shi and Malik (2000), Weiss (1999), Coifman and Lafon (2006). In Sarti and Citti (2015) showed how this spectral mechanism is implemented in neural morphodynamics in terms of symmetry breaking of mean field neural equations. In that sense, Sarti and Citti (2015) can be considered an extension of Bressloff, Cowan, Golubitsky, Thomas, and Wiener (2002).

In this letter, we further develop the approach introduced in Sarti and Citti (2015) and describe an algorithm for the individuation of perceptual units using both local and global constraints: local constraints are modeled by suitable connectivity kernels, which represent neural connections, and the global percepts are computed by means of spectral analysis. The model is described in the geometric setting of a Lie group equipped with a sub-Riemannian metric introduced in Petitot and Tondut (1999), Citti and

Sarti (2006), Sarti, Citti, and Petitot (2008). Despite the apparent mathematical difficulty, it helps to clarify in a rigorous way the gestalt law of good continuation.

Here we introduce various substantial differences from the techniques in literature. While studying the local properties of the model, we focus on the properties of the connectivity kernels. The Fokker-Planck and the Laplacian kernel in the motion group are already largely used for the description of the connectivity, since they qualitatively fit the experimental data (Sarti & Citti, 2015). Here we perform a quantitative fitting between the computed kernels and the experimental ones in order to validate the model. We show that the cortical architecture is a realization of stochastic sample functions and how, through this realization, we can construct the connectivity kernel. We make a comparison between the fundamental solution of the Fokker-Planck equation with the experimental data of Bosking et al. (1997), Ben-Shahar and Zucker (2004), and Gilbert, Das, Ito, Kapadia, and Westheimer (1996), showing how the stochastic paths are implemented in the neural network. In particular, we consider the distribution of a tracer through lateral connection, modeling each injection with stochastic paths. The bouton distributions are realizations of a stochastic process, in particular, of a random walk in $\mathbb{R}^2 x S^1$ space. We show how the probability density obtained as a combination of Fokker-Planck is an integration of stochastic paths. Moreover, we propose also using the subelliptic Laplacian kernel in order to account for the variability of connectivity patterns. Second we accomplish grouping with a spectral analysis inspired by the work of Sarti and Citti (2015), who proved the neurophysiological plausibility of this process. In the experiments, we manipulate the stimuli to demonstrate the relation between the pop-up of the figure and the eigenvalue analysis. We analyze in particular the swap between one solution and the other while smoothly changing the stimulus in many grouping experiments. Finally, we enrich the model, exploiting the role of the polarity feature, which allows us to work with two competing kernels.

The plan of the letter is the following. Section 2 is divided in two parts, first describing local constraints and then global ones. We first recall the neurogeometry of the visual cortex and see how the cortical connectivity is represented by the fundamental solution of Fokker-Planck, sub-Riemannian Laplacian, and isotropic Laplacian equations. We propose a method for the individuation of perceptual units, first recalling the notions of spectral analysis of connectivity matrices, obtained by the connectivity kernels. We will see how eigenvectors of this matrix represent perceptual units in the image. In section 3, we present numerical approximations of the kernels and will compare kernels with neurophysiological data of horizontal connectivity (Angelucci et al., 2002; Bosking et al., 1997). We will see how the differential equations that we need to solve originate from a stochastic process, estimated with the efficient numerical technique of Markov chain Monte Carlo methods (MCMC). We also perform a quantitative validation of the

kernel considering the experiment of Gilbert et al. (1996), showing the link between the connectivity kernel and the cell's response. Finally in section 4, we present the results of simulations using the implemented connectivity kernels. We identify perceptual units in different Kanizsa figures, highlighting the role of polarity and discussing and comparing the behavior of the different kernels.

2 The Mathematical Model

In this section, we identify a possible neural basis for local gestalt laws in the functional architecture of the primary visual cortex, the first cortical structure that underlies the processing of visual stimuli. We do not claim that the process of grouping has to be attributed exclusively to V1, since several cortical areas are involved in segmentation of a figure. However, neural evidence ensures that it takes place in V1 (see Lee & Nguyen, 2001; Pillow & Nava, 2002). Hence, we focus on this area where the first elaboration is made and is therefore important for the geometrical aspects of the process.

2.1 Local Constraints: The Neurogeometry of V1. In the 1970s, Hubel and Wiesel (1962, 1977) discovered that this cortical area is organized in the so-called hypercolumnar structure. This means that for each retinal point (x, y), there is an entire set of cells, each one sensitive to a specific orientation θ of the stimulus.

The first geometrical models of this structure are due to Hoffman (1989), Koenderink and van Doorn (1987), Williams and Jacobs (1997), and Zucker (2006). They described the cortical space as a fiber bundle, where the retinal plane (x, y) is the basis, while the fiber concides with the hypercolumnar variable θ. More recently, Petitot and Tondut (1999), Citti and Sarti (2006), and Sarti et al. (2008), proposed describing this structure as a Lie group with a sub-Riemannian metric (see also the results of Duits & Franken, 2009). This expresses the fact that each filter can be recovered from a fixed one by translation of the point (x, y) and rotation of an angle θ. In particular, the visual cortex can be described as the subset of points of $\mathbb{R}^2 \times S^1$. Every simple cell is the characterized by its receptive field, classically defined as the domain of the retina to which the neuron is sensitive. The shape of the response of the cell in the presence of a visual input is called the receptive profile (RP) and can be reconstructed by electrophysiological recordings (Ringach, 2002). In particular, simple cells of V1 are sensitive to orientation and are strongly oriented. Hence, their RPs are interpreted as Gabor patches (Daugman, 1985; Jones & Palmer, 1987). They are constituted by two coupled families of cells: an even and an odd-symmetric one.

Via the retinotopy, the retinal plane can be identified with the two-dimensional plane $\mathbb{R}^2$. A visual stimulus at the retinal point (x, y) activates the whole hypercolumnar structure over that point. All cells fire, but the

cell with the same orientation of the stimulus is maximally activated, giving rise to orientation selectivity.

Formally, curves and edges are lifted to new cortical curves, identified by the variables (x, y, θ) where θ is the direction of the boundary at the point (x, y). Citti and Sarti (2006) showed that these curves are always tangent to the planes generated by the vector fields. They modeled these curves as integral curves of suitable vector fields in the $SE(2)$ cortical structure. The vector fields they considered are

$$\vec{X}_1 = (\cos\theta, \sin\theta, 0), \quad \vec{X}_2 = (0, 0, 1). \tag{2.1}$$

All lifted curves are integral curves of these two vector fields such that a curve in the cortical space is

$$c'(s) = (k_1(s)\vec{X}_1 + k_2(s)\vec{X}_2)(c(s)), \quad c(0) = 0. \tag{2.2}$$

Citti and Sarti (2006) noted that these curves, projected on the 2D cortical plane, are a good model of association fields.

2.2 A Model of Cortical Connectivity. From a neurophysiological point of view, there is experimental evidence of the existence of connectivity between simple cells belonging to different hypercolumns—the so-called long-range horizontal connectivity. Combining optical imaging of intrinsic signals with small injections of biocytin in the cortex, Bosking et al. (1997) clarified properties of horizontal connections on V1 of the tree shrew. The propagation of the tracer is strongly directional, and the direction of propagation coincides with the preferential direction of the activated cells. Hence, connectivity can be summarized as preferentially linking neurons with co-circularly aligned receptive fields.

The propagation along the connectivity can be modeled as the stochastic counterpart of the deterministic curves defined in equation 2.2 for the description of the output of simple cells. If we assume a deterministic component in direction X_1 (which describes the long-range connectivity) and stochastic component along X_2 (the direction of intracolumnar connectivity), the equation can be written as

$$(x', y', \theta') = (\cos\theta, \sin\theta, N(0, \sigma^2)) = \vec{X}_1 + N(0, \sigma^2)\vec{X}_2, \tag{2.3}$$

where $N(0, \sigma^2)$ is a normally distributed variable with zero mean and variance equal to σ^2. The probability density of this process, denoted by v, was first used by Williams and Jacobs (1997) to compute the stochastic completion field, August and Zucker (2000, 2003) to define the curve indicator random field, and Duits and Franken (in Duits & van Almsick, 2008; Duits

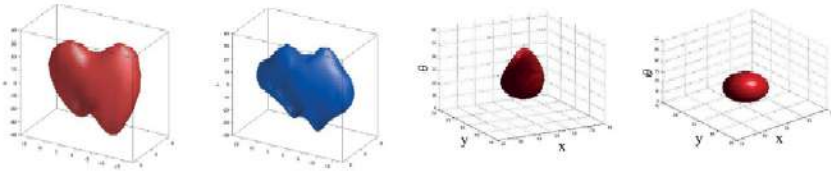

(a) (b) (c) (d)

Figure 2: (a) An isosurface of the connectivity kernel ω_1 obtained by symmetrization of the Fokker-Planck fundamental solution, equation 2.4. (b) The distribution of co-occurrence of edges in natural images (from Sanguinetti et al., 2008). (c) Isosurface of the connectivity kernel ω_2 obtained from the fundamental solution Γ_2 of the sub-Riemannian Laplacian equation, equation 2.7. (d) An isosurface of the fundamental solution of the isotropic Laplacian, equation 2.11.

& Franken, 2009) to perform contour completion, denoising, and contour enhancement.

The kernel obtained integrating in time the density v_1,

$$\Gamma_1(x, y, \theta) = \int_0^{+\infty} v_1(x, y, \theta, t)dt, \tag{2.4}$$

is the fundamental solution of the Fokker-Planck operator $FP = X_1 + \sigma^2 X_{22}$.

The kernel Γ_1 is strongly biased in the direction X_1 and is not symmetric. A new symmetric kernel can be obtained as following

$$\begin{aligned}\omega_1((x, y, \theta), (x', y', \theta')) = \frac{1}{2}(&\Gamma_1((x, y, \theta), (x', y', \theta')) \\ &+ \Gamma_1((x', y', \theta'), (x, y, \theta)).\end{aligned} \tag{2.5}$$

Figure 2a shows an isosurface of the symmetrized kernel ω_1, with its typical twisted butterfly shape. The kernel ω_1 has been proposed in Sanguinetti et al. (2008) as a model of the statistical distribution of edge co-occurrence in natural images. The similarity between the two is proved at both a qualitative and a quantitative level (Sanguinetti et al., 2008; see also Figures 2a and 2b).

If we assume that intracolumnar and long-range connections have comparable strength, the stochastic equation, equation 2.3, reduces to

$$(x', y', \theta') = N(0, \sigma_1^2)\vec{X}_1 + N(0, \sigma_2^2)\vec{X}_2. \tag{2.6}$$

where $N(0, \sigma_i^2)$ are normally distributed variables with zero mean and variance equal to σ_i^2. In this case, the speed of propagation in directions X_1 and X_2 is comparable. The associated probability density is the fundamental

solution of the sub-Riemannian heat equation (Jerison & Sanchez-Calle, 1986). The integral in time of this probability density,

$$\Gamma_2(x, y, \theta) = \int_0^{+\infty} \nu_2(x, y, \theta, t)dt \tag{2.7}$$

is the fundamental solution of the sub-Riemannian Laplacian (SRL): $SRL = \sigma_1^2 X_{11} + \sigma_2^2 X_{22}$.

It is a symmetric kernel, so we do not need to symmetrize it, and we use it as a model of the connectivity kernel:

$$\omega_2((x, y, \theta), (x', y', \theta')) = \Gamma_2((x, y, \theta), (x', y', \theta')). \tag{2.8}$$

Figure 2c shows an isosurface of the connectivity kernel ω_2.

We will see in section 3.2 that a combination of Fokker-Planck and sub-Riemannian Laplacian fits the connectivity map measured by Bosking et al. (1997), where the Fokker-Planck fundamental solution represents well the long distances of the trajectory, while the sub-Riemannian Laplacian represents the short ones. A combination of different Fokker-Planck fundamental solutions can also be used to model the functional architecture of primates experimentally measured by Angelucci et al. (2002).

While validating the model, we will see that a standard Riemannian kernel does not provide the same accurate results. In order to show this, we introduce an isotropic version of the previous model, a standard Riemannian kernel. To constuct it, we complete the family of vector fields in equation 2.1 with an orthonormal one,

$$\vec{X}_3 = (-\sin(\theta), \cos(\theta), 0), \tag{2.9}$$

choosing stochastic propagation in any direction, in such a way that equation 2.3 becomes

$$(x', y', \theta') = N(0, \sigma^2)\vec{X}_1 + N(0, \rho^2)\vec{X}_2 + N(0, \sigma^2)\vec{X}_3, \tag{2.10}$$

where $N(0, \sigma_i^2)$ are normally distributed variables with zero mean and variance equal to σ_i^2. Its probability density will be denoted ν_3, and the associated time-independent kernel,

$$\Gamma_3(x, y, \theta) = \int_0^{+\infty} \nu_3(x, y, \theta, t)dt, \tag{2.11}$$

will be the fundamental solution of the standard Laplacian operator: $L = \sigma^2 X_{11} + \rho^2 X_{22} + \sigma^2 X_{33} = \sigma^2(\partial_{xx} + \partial_{yy}) + \rho^2 \partial_{\theta\theta}$. One of its level sets is represented in Figure 2d.

In section 3.1 we describe a numerical technique to construct the three kernels we have described.

2.3 Global Integration. Since the beginning of the twentieth century, perception has been considered by gestaltists as a global process. Moreover, following Koch and Ullman (1985) and Merleau-Ponty (1945), visual perception is a process of the visual field that individuates figure and background at the same time. Then it continues in segmenting the structures by succeeding differentiations.

A cortical mechanism responsible for this analysis has been outlined by Sarti and Citti (2015), starting from the classical mean field equation of Ermentrout and Cowan (1980) and Bressloff and Cowan (Bressloff et al., 2002; Bressloff & Cowan 2003). This equation describes the evolution of cortical activity and depends on connectivity kernels. The discrete output h of the simple cells selects in the cortical space (x, y, θ) the set of active cells, and the cortical connectivity, restricted on this set, defines a neural affinity matrix. The eigenvectors of this matrix describe the stationary states of the mean field equation—hence, the emergent perceptual units. The system will tend to the eigenvector associated with the highest eigenvalue, which corresponds to the most important object in that scene. Mathematically the approach is strongly linked to spectral analysis techniques for locality-preserving embeddings of large data sets (Coifman & Lafon, 2006; Belkin & Niyogi, 2003; Roweis & Saul, 2000), data segregation and partitioning (Perona & Freeman, 1998; Meila & Shi, 2001; Shi & Malik, 2000), and a grouping process in real images (Weiss, 1999).

2.4 The Cortical Activity Equation. We have seen that in the presence of a visual stimulus, cells aligned to its boundary give the maximal response. We will assume that a discrete number of cells N are maximally activated and denote them (x_i, y_i, θ_i) for $i = 1, \ldots, N$. In Figure 10b, we show as an example the cells responding to a Kanizsa figure, represented with their Gabor-like receptive profiles. Following Sarti and Citti (2015), the cortical connectivity is restricted to this discrete set and reduces to a matrix A:

$$A_{i,j} = \omega((x_i, y_i, \theta_i), (x_j, y_j, \theta_j)). \tag{2.12}$$

In this discrete setting, the mean field equation for the cortical activity reduces to

$$\frac{du}{dt} = -\lambda u(i) + s\left(\sum_{j=1}^{N} A(i, j)u(j)\right), \tag{2.13}$$

where s is a sigmoidal function and λ is a physiological parameter. The solution tends to its stationary states, which are the eigenvectors of the associated linearized equation:

$$\sum_{j=1}^{N} A_{i,j} u_j = \frac{\lambda}{s'(0)} u_i. \tag{2.14}$$

Hence these are the emergent states of the cortical activity that individuates the coherent perceptual unit in the scene and allows segmenting it. This is why we will assign to the eigenvalues of the affinity matrix the meaning of a salience index of the objects. Since we have defined three different kernels, we will define different affinity matrices. However, all kernels are real and symmetric, so matrix A is a real symmetric matrix $A_{i,j} = A_{j,i}$. Their eigenvalues are real, and the highest eigenvalue is defined. The associated principal eigenvectors emerge as symmetry breaking of the stationary solutions of mean field equations, and they pop up abruptly as emergent solutions. The first eigenvalue will correspond to the most salient object in the image.

2.5 Individuation of Perceptual Units. Since the three different kernels assign different roles to different directions of connectivity, the different affinity matrices and their spectrum will reflect these different behaviors. Consequently, the resulting data set partitioning will be stronger in the straight direction using the Fokker-Planck ω_1 kernel or will allow rotation using the ω_2 kernel (see also (Cocci, Barbieri, Citti, & Sarti, 2015) for a deeper analysis). Using the kernel ω_3, we expect an equal grouping capability in the collinear and the ladder directions.

In Figure 3 we show the affinity matrix of the image presented in Figure 10a. It is a square matrix with dimensions $N \times N$, where N is exactly the number of active patches. It represents the affinity of each patch with respect to all the others. The structure of the affinity matrix is composed by blocks, and the principal ones correspond to coherent objects. On the right, we visualize the complete set of eigenvalues in a graph (eigenvalue number, eigenvalues). We explicitly note that the first eigenvector will have the meaning of the emergent perceptual unit. The other eigenvectors do not describe an ordered sequence of figures with different rank. However, their presence is important, especially when two eigenvalues have similar values. In this case, a small deformation of the stimulus can induce a change in the order of the eigenvalues and produce a sudden emergence of the corresponding eigenvector with an abroupt change in the perceived image.

This is in good agreement with the perceptual characteristics of salient figures of temporal and spatial discontinuity, since they pop up abruptly

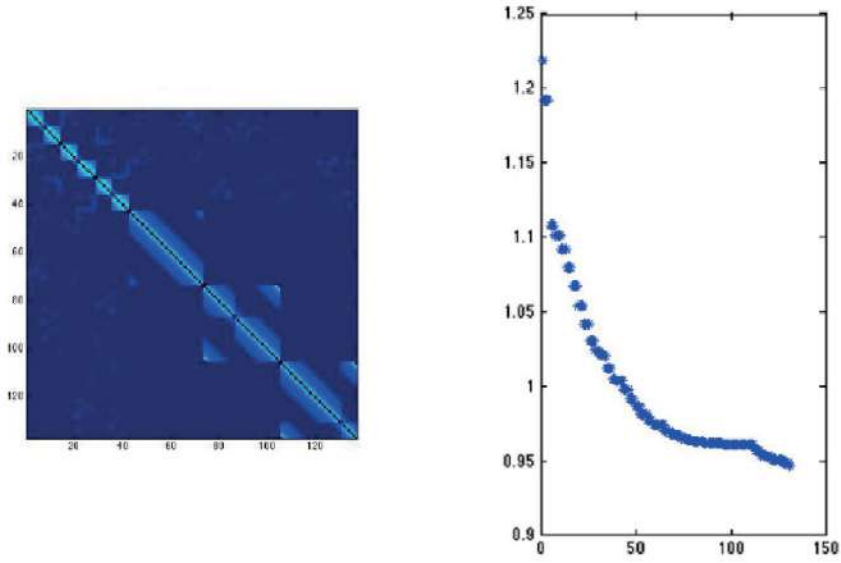

Figure 3: (Left) The affinity matrix that contains information about the affinity of an active patch with respect to all the others. (Right) The set of its sorted eigenvalues.

from the background, while the background is perceived as undifferentiated (Merleau-Ponty, 1945). Spectral approaches give reason to the discontinuous character of figure-ground articulation better than continuous models, which instead introduce a gradual change in the perception of figure and background (Lorenceau & Alais, 2001).

To find the remaining objects in the image, the process is then repeated on the vector space orthogonal to p; the second and the following eigenvectors can be found, until the associated eigenvalue is sufficiently small. In this way, only n eigenvectors are selected; with $n < N$, this procedure reduces the dimensionality of the description. This procedure neurally reinterprets the process introduced by Perona and Freeman (1998).

3 Quantitative Kernel Validations

3.1 Numerical Approximations of the Kernels. In this section, we numerically approximate the connectivity kernels ω_μ defined in section 2.

We obtain the discrete fundamental solution Γ_1 of equation 2.4 by developing random paths from the numerical solution of system 2.3 that can be approximated by

$$\begin{cases} x_{s+\Delta s} - x_s = \Delta s \cos(\theta) \\ y_{s+\Delta s} - y_s = \Delta s \sin(\theta), \quad s \in 0, \ldots, H, \\ \theta_{s+\Delta s} - \theta_s = \Delta s N(\sigma, 0) \end{cases} \tag{3.1}$$

where H is the number of steps of the random path and $N(\sigma, 0)$ is a generator of numbers taken from a normal distribution with mean 0 and variance σ. In that way, the kernel is numerically estimated with MCMC methods (Robert & Casella, 2013). Various realizations n of the stochastic path will be given to solve this finite difference equation n times; the estimated kernel is obtained averaging their passages over discrete volume elements, as described in detail in Higham (2001) and Sarti and Citti (2015). We first fix a discretization step $\Delta s = 1$ without loss of generality; then we simulate several n discrete-times random paths, assigning a value between 0 and 1 corresponding to the number of paths that passed through it, divided by n. This provides a distribution over the cells that, for a large value of n gives a discrete approximation of the connectivity kernel (Cocci et al., 2015). Proceeding with the same methodology, the numerical evaluation of fundamental solution Γ_2 of the hypoelliptic Laplacian, equation 2.7, is obtained, and system 2.6 is discretized:

$$\begin{cases} x_{s+\Delta s} - x_s = \Delta s R \cos(\theta) \\ y_{s+\Delta s} - y_s = \Delta s R \sin(\theta), \quad s \in 0, \ldots, H, \\ \theta_{s+\Delta s} - \theta_s = \Delta s N(\sigma_3, 0) \end{cases} \tag{3.2}$$

where $R = N(\sigma_1, 0)$ and σ_3 is the variance in the θ direction. The kernel represented in Figure 2c is obtained by numerical integration of that system and averaging as before the resulting paths.

Finally, system 2.10, a model for the isotropic diffusion equation (see equation 2.11), is approximated by

$$\begin{cases} x_{s+\Delta s} - x_s = \Delta s N(\sigma, 0) \\ y_{s+\Delta s} - y_s = \Delta s N(\sigma, 0), \quad s \in 0, \ldots, H, \\ \theta_{s+\Delta s} - \theta_s = \Delta s N(\rho, 0) \end{cases} \tag{3.3}$$

where σ, ρ are the variances in the x, y, θ directions. In order to obtain the approximation of the kernel ω_3, visualized in Figure 2d, the system is integrated with the same technique used before. These kernels will be used to construct the affinity matrices in equation 2.12.

3.2 Stochastic Paths and Cortical Connectivity. In this section, we describe the cortical architecture as a realization of stochastic sample functions; in particular, we will see how the connectivity is associated with random

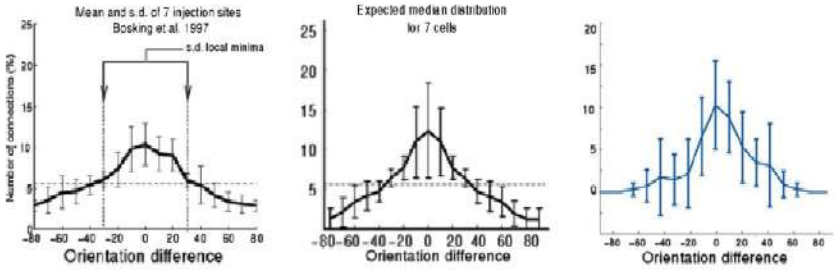

Figure 4: From left to right, the mean connection distribution computed in Ben-Shahar and Zucker (2004) from Bosking et al. (1997), the expected median distribution for seven cells, and the results of our model considering seven different random paths. It is evident that in our model, the standard deviation is nonmonotonic with two local minima at $+30$ and -30 degrees.

paths. We will show that the position of presynaptic boutons in the images of Bosking et al. (1997) can be seen as the realization of stochastic paths via the anatomy. Every random walk starts from the injection site of a tracer and gives the position of a set of boutons, as visualized in Figure 5a. The probability density, which is described as a kernel, is the integration of all the random paths; it is estimated as the density of the boutons. Finally, the probability density, which is described as a kernel, is obtained as the integration of all the random paths. From a neural point of view, this integration, which can be interpreted as the action of a columnar population, provides an estimation of the density of the boutons.

We consider a hypercolumn of the ice cube scheme visual cortex, composed of approximately 100 neurons. In the connectivity map in Figure 5a, we notice the presence of an average of six boutons. In this way, the number of possible connections in the visual cortex is 100^6, and in our model, we use a number of paths compatible with these data. Now we make a comparison between the connectivity kernel previously defined and the experiments of Bosking et al. (1997) and Ben-Shahar and Zucker (2004).

In Figure 4, we can see the results of Ben-Shahar and Zucker (2004). On the left are the mean and standard deviation of the distribution of long-range connections of seven injection sites considering the data of Bosking et al. (1997), and in the middle are the expected median distribution for seven cells from the curve model described in Ben-Shahar and Zucker (2004). They noticed that the standard deviation is nonmonotonic, finding two local minima at approximately $+30$ and -30 degrees. Confirming their results, we show that our model implies a nonmonotonically changing variance as the orientation difference increases. In particular, on the right side of Figure 4, the mean and the standard deviation of seven random paths, at a fixed orientation, are visualized. We notice the presence of the nonmonotonicity of the standard deviation and that the two local minima at almost 30 degrees are preserved.

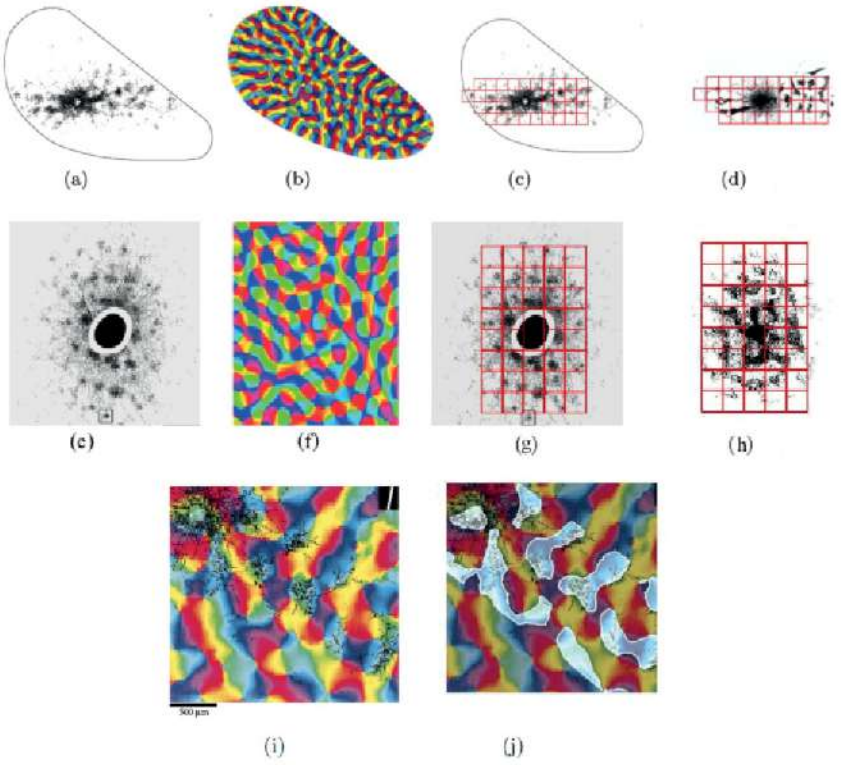

Figure 5: The connectivity map measured by Bosking et al. (1997) (a) and Angelucci et al. (2002) (e), the pinwheel structure used for the estimate (b, f), the tracer partitioned according to rectangles with sides equal to the distance between pinwheels (c, g), and the best-fit results (d, h). (i) The tracer superimposed to the pinwheel structure found by Bosking et al. (1997). (j) The isocontours obtained from a combination of Fokker-Planck.

Moreover, the fact that the mean and the variance of the model are similar to the experimental data suggests that the choice of the normal distribution allows us to find physiological values. For these reasons, the connectivity represents the anatomical implementation of random paths.

We will now examine to what extent kernels ω_i, $i = 1, 2$ are models of connectivity. The kernel ω_3 is used for comparison and to show that a uniform Euclidean kernel does not capture the anysotropic structure of the cortex. The random paths that we compute through MCM are implemented in the functional architectures in terms of the horizontal connectivity of a single cell. However, the connectivity of an entire population of cells corresponds to the set of all single cells' connectivities and then to the Fokker-Planck fundamental solution.

A first qualitative comparison between the kernels ω_1, ω_2 and the connectivity pattern has been provided in Sarti and Citti (2015). Here we follow the same framework, but we propose a more accurate quantitative comparison.

It is well known that the 3D cortical structure is implemented in the 2D cortical layer as a pinwheel structure, which codes for position and orientations (see Figure 5b). The pinwheel structure has a large variability from one subject to the next, but within each species, common statistical properties have been obtained. Cortico-cortical connectivity has been measured by Bosking et al. (1997) by injecting a tracer in a simple cell and recording the trajectory of the tracer. In Figure 5a, the propagation through the lateral connections is represented by black points. Bosking found a large variability of injections, which have common stochastic properties as the direction of propagation, a patchy structure with small blobs at approximately fixed distance and the decay of the density of tracer along the injection site.

We model each injection with stochastic path solutions of equation 2.3. Then we evaluate the stochastic paths on the pinwheel structure.

Due to the stochastic nature of the problem, we do not compare pointwise the image of the tracer and the stochastic paths, but we average them on the pinwheels. We partition the images of both the tracer and the stochastic paths in M regions corresponding to the pinwheels,

$$I = \cup_i R_i, \tag{3.4}$$

and for every R_i, we compute the density of tracer DT_i and the density of the stochastic paths DP_i. The two vectors DT_i and DP_i are then rescaled in such a way to have unitary L^2-norm, and the mean square error is computed:

$$E = \sqrt{\frac{1}{M}\sum_{i=1}^{M}\left(DP_i - DT_i\right)^2}. \tag{3.5}$$

The free parameters of the model are the value of the standard deviation, and the number of paths, the number of steps, appearing in equation 2.3 and in system 3.1. The best fit between the experimental and simulated distributions has been accomplished by minimizing the mean square error by varying these parameters.

Due to the different role of the directions X_1 and X_2 in the definition of these kernels, the sub-Riemannian Laplacian paths and the Fokker-Planck paths have different structures.

The sub-Riemannian Laplacian allows diffusion in direction X_2 and favors the change of the angle; it can be used to describe short-range connectivity as described in section 4.4. Hence, it is responsible for the center blob in a neighborhood of injection points (see Figure 5c). The Fokker-Planck kernel produces an elongated, patchy structure and seems responsible for the long-range connection (see Figure 5d). We apply our quantitative fit

only to the long-range connectivity, discarding the tracer in the neighborhood of the injection. For this reason, the sub-Riemannian Laplacian is not involved in validating the model.

The method is first applied to fit the image of the tracer taken by Bosking et al. (1997; see Figure 5a). They evaluate all the kernels on the pinwheels (see Figure 5b) to obtain a patchy structure. In order to apply formula 3.4, we cover both the image of the tracer and the Fokker-Planck with a regular distribution of rectangles, with edges equal to the mean distances between pinwheels (see Figures 5c, and 5d); clearly we do not cover the central zone, where we cannot fit the Fokker-Planck kernel). The resulting error value is $E < 8\%$, showing that the model accurately represents the experimental distribution.

A similar procedure has been applied to the image of the tracer provided in Angelucci et al. (2002; see Figure 5e). They obtain their result with various injections in the neighborhood of a pinwheel, so that all orientations are present and the tracer propagates in all directions. In this case, we do not have natural pinwheels; hence we use artificial pinwheels, obtained with the algorithm presented in Barbieri, Citti, Sanguinetti, and Sarti (2012); see Figure 5f), with the constraint that the mean distance between the artificial pinwheels is equal to the mean distance between the blobs produced by the tracer. Here we consider Fokker-Planck paths with all directions to obtain the apparent isotropic diffusion. Also in this case, we cover it with rectangles and perform a best fit; the minimum error value is $E < 8\%$ (see Figures 5g, and 5h).

Bosking et al. (1997) showed a famous image with the tracer superimposed to the pinwheel structure (see Figure 5i). In this case, we have the tracer and the pinwheel of the same animal. This allows going below the scale of the pinwheel, and we correctly recover the orientation with the pinwheel (see Figure 5j). The estimated kernel is again a combination of Fokker-Planck. As before, we focus on orientations; hence, we model only the long-range part of the image, discarding the center blob. The evaluation of the error is made with squared regions at a scale smaller that the pinwheel, and the error goes below $E < 9\%$.

3.3 Perceptual Facilitation and Density Kernels. In order to obtain a stable and deterministic estimate of this stochastic model, we used the density kernel, which is a regular deterministic function, coding the main properties of the process. We perform a quantitative validation of these regular kernels, comparing them to an experiment by Gilbert et al. (1996). They study the capability of cells to integrate information out of the single receptive field of the cells. This integration process is due to the long-range horizontal connections; hence, it can be used to validate our model of long-range connectivity. As we recognized in the previous section, it is the Fokker-Planck kernel that can be considered a model for long-range connectivity; hence, we use this kernel.

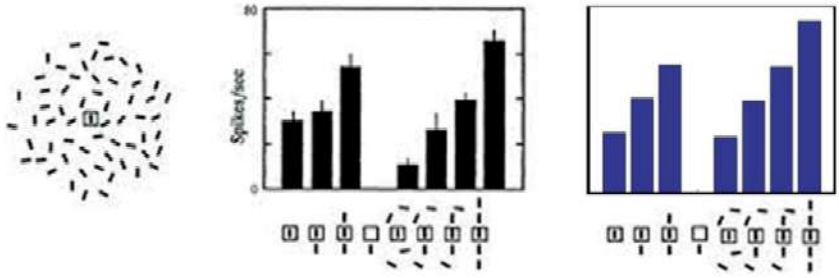

Figure 6: (Left) The experiment of Gilbert et al. (1996), with the stimulus composed by randomly placed and oriented lines and the black histogram of the cell's response. (Right) The histogram evaluated from the probability density in response to the same distribution of lines.

Figure 6 (left) shows the results of Gilbert et al. (1996), visualized by the cell's response to randomly placed and oriented lines in a black histogram. A vertical line is present in the receptive field of a cell selective to this orientation, and the intensity of its response is represented in the first column of the histograms. If the stimulus is surrounded by random elements aligned with the first one, the cell's response increases (resp. the second, third, and last column of the histograms). When the other random elements are not aligned with the fixed one (as in the fifth, sixth, and seventh columns), the cell's response decreases because it reflects an inhibitory effect.

On the right in the blue histogram, we evaluate the probability density modeled by the kernel in equation 2.5 in the presence of the same configuration of elements. The same trend is obtained considering the probability density distribution, as visualized in Figure 6 (right). In order to consider the inhibitory effect, we evaluate the kernel with 0 mean. A quantitative analysis of the differences between them has been evaluated considering the mean square error between the two normalized histograms. The error of 8% underlines how this connectivity kernel well represents neural connections.

4 Emergence of Percepts

In the following experiments, some numerical simulations will be performed in order to test the reliability of the method for performing grouping and detecting perceptual units in an image. The kernel considered here depends only on orientation. Hence, it can be applied to detect the salience of geometrical figures, which can be very well described using this feature.

The purpose is to select the perceptual units in these images, using the following algorithm:

1. Define the affinity matrix $A_{i,j}$ from the connectivity kernel.
2. Solve the eigenvalue problem $A_{i,j}u_i = \lambda_i u_i$, where the order of i is such that λ_i is decreasing.
3. Find and project on the segments the eigenvector u_1 associated with the largest eigenvalue.

The parameters used are 1,000,000 random paths with $\sigma = 0.15$ in system 3.1, $\sigma_1 = 1$ and 2, $\sigma_3 = 0.11$ in system 3.2, and $\sigma, \rho = 0.15$ in system 3.3. The value of H is defined as follows: $H = \frac{1}{3}d_{\max}$, where $d_{\max}$ is the maximum distance between the inducers of the stimulus. Similar parameters have been used for all the experiments.

4.1 The Field, Hayes, and Hess Experiment. In this section we consider some experiments similar to the ones of Field et al. (1993), where a subset of elements organized in a coherent way is presented out of a ground formed by a random distribution of elements. A first stimulus of this type is represented in Figure 7a. The connectivity among these elements is defined as in equations 2.4 and 2.7.

After the affinity matrix and its eigenvalues, the eigenvector corresponding to the highest eigenvalue is in red. The results show that the stimulus is well segmented with the fundamental solutions of Fokker-Planck and sub-Riemannian Laplacian equations (see Figure 7b).

Now we consider a similar experiment proposed in Field et al. (1993), where the orientation of successive elements differs by 15, 30, 45, 60, and 90 degrees and the ability of the observer to detect the path was measured experimentally. It was proved that the path can be identified when the successive elements differ by 60 degrees or less. With our method, we obtain similar results: if the angle between successive elements is less than 60 degrees (Figures 7c–7e), the identification of the unit is correctly performed. With an angle equal to 60 degrees (see Figure 7f) only part of the curve is correctly detected: this can be interpreted as the observer's increasing difficulty to detect the path. Finally, with higher angles (Figure 7g) the first eigenvector of the affinity matrix corresponding to random inducers, confirming the results.

Finally, we present an example where there are two units in the scene with roughly equal salience and roughly equal eigenvalues. In the first and the second rows of Figure 8, the stimuli are composed by a curve and a line in a background of random elements. In stimulus a, represented in the first row, the elements composing the curve are perfectly aligned and very nearby, so this has the highest saliency; it represents the eigenvector associated with the first eigenvalue (as shown in red in Figure 8b). The second eigenvalue in this case is slightly smaller. After computing the first eigenvector, the stimulus is updated (see Figure 8c); the first eigenvector of the new affinity matrix is computed, corresponding to the inducers of the line (see Figure 8d).

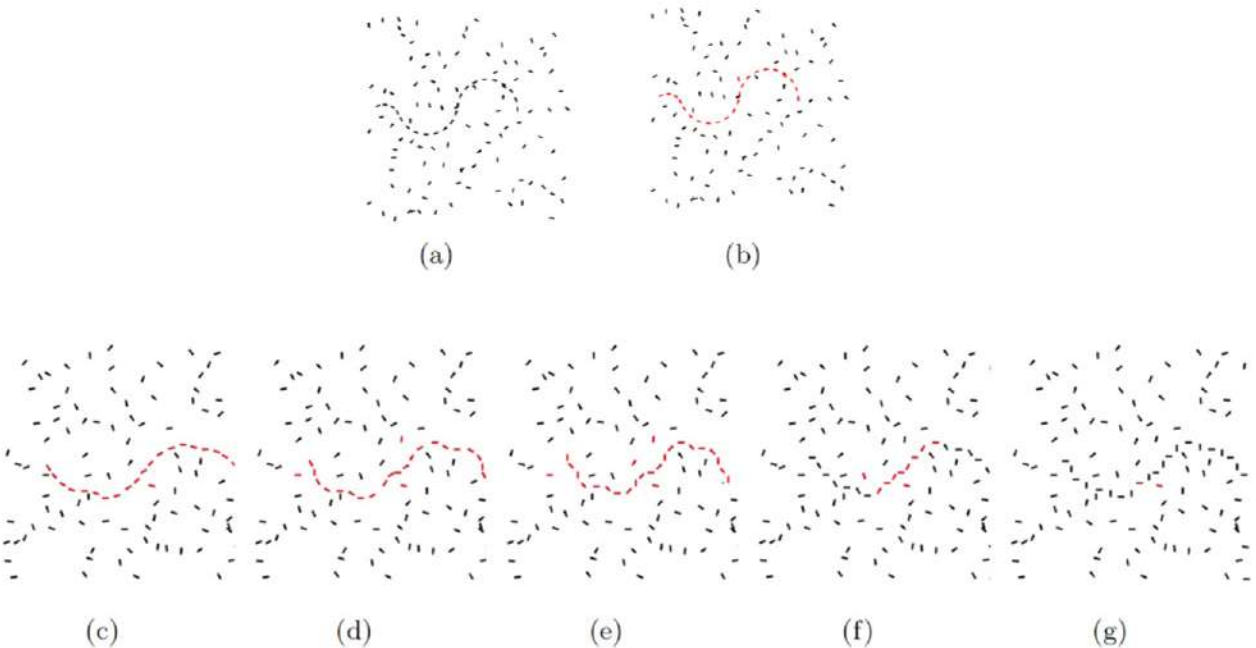

Figure 7: (Top) Example of stimulus a similar to the experiments of Field et al. (1993). The stimulus containing a perceptual unit is segmented with Fokker-Planck and sub-Riemannian Laplacian (b), using the first eigenvector of the affinity matrix. (Bottom) In red, the first eigenvector of the affinity matrix considering images containing paths in which the orientation of successive elements differs by 15 (c), 30 (d), 45 (e), 60 (f), and 90 (g) degrees.

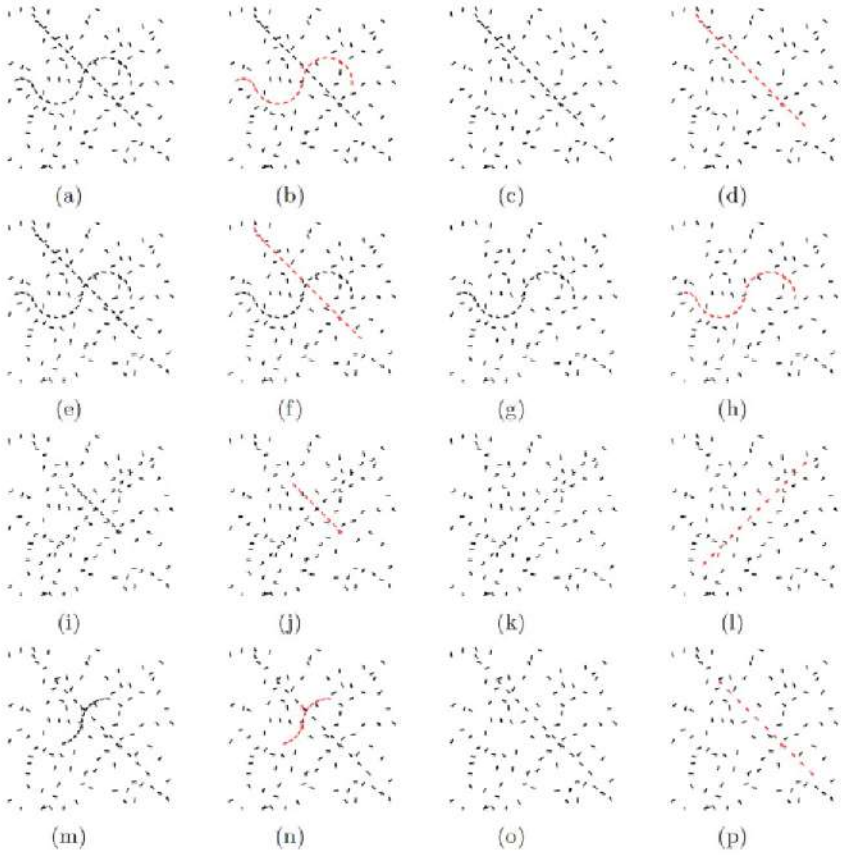

Figure 8: (Rows 1 and 2) Examples (a, e) with two units in the scene, where a change in the angle leads to a change in the order of the eigenvalues (b, f), (c, g), (d, h). (Rows 3 and 4) Examples (i, m) with two units in the scene with different length. (j, n; k, o; l,p) The results of simulation.

In the second row (see Figure 8e), we slightly modify the stimuli, in particular the alignement of the element forming the curve (e.g., an angle of pi/18). As a consequence, the line becomes the most salient perceptual unit and the first eigenvector (see Figure 8f). The stimulus is updated (see Figure 8g), and the first eigenvector of the new affinity matrix corresponds to the inducers of the curve (see Figure 8h). It is notable that in this case, a small change of the eigenvalues corresponds to small change of the eigenvectors, but the first eigenvalue swaps with the second one, and, consequently, we obtain an abrupt change in the perceived object.

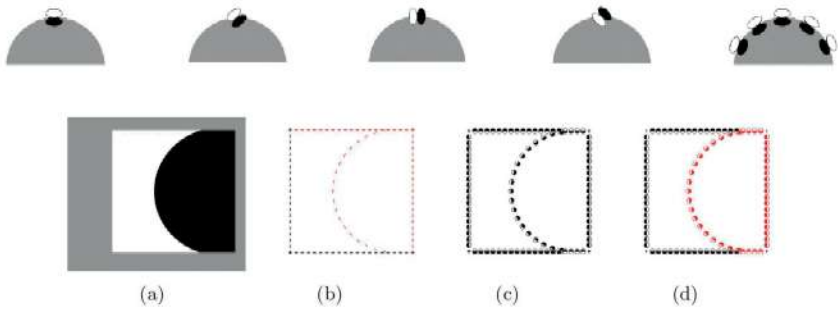

Figure 9: (Top) Schematic description of the hypercolumn of odd simple cells centered in a point (x, y). The maximal activity is observed for the simple cell sensitive to the direction of the boundary of the visual stimulus. The set of maximally firing cells is visualized in the last image. (Bottom) (a) A cartoon image, (b) the first eigenvector of the affinity matrix without polarity, (c) its representation with polarity-dependent Gabor patches, and (d) the corresponding first eigenvector.

In the previous examples, we have considered all contours with almost the same length. We show here that this length does not affect the feature of saliency. The two perceptual units in Figures 8i and 8m have different lengths. The results underline how the proximity of contours is stronger than length: the shortest units with nearer segments are the first perceptual units, associated with the most salient eigenvectors (see Figures 8j, and 8n). Then the stimuli are updated in Figures 8k and 8o), and the second eigenvectors are visualized in Figures 8l and 8p.

In this analysis, different features can be considered. In particular, the distances between the segments also play a central role. Consider, for example, the straight line in Figure 8a. If one or more segments is missing from the contour, we could obtain a less accurate segmentation (a similar effect is noticed in the case of unaligned segments). Favali, Abbasi-Sureshjani, ter Haar Romeny, and Sarti (2016) considered a similar analysis with small or disconnected contours applied to the study of vessel connectivities.

4.2 The Role of Polarity. The feature of polarity leads to inserting in the model the feature of contrast: contours with the same orientation but opposite contrast are referred to opposite angles. For this reason, we assume that the orientation θ takes values in $[0, 2\pi)$ when we consider the odd filters and in $[0, \pi)$ while studying the even ones.

The response of the odd filters in presence of a cartoon image is schematically represented in Figure 9. At every boundary point, the maximally activated cell is the one with the same direction of the boundary. Then the

maximally firing cells are aligned with the boundary (see Figure 9, top right).

In order to clarify the role of polarity, we consider an image in Figure 9a, studied by Kanizsa (1980), in the contest of a study of convexity in perception. In this case, if we consider only the orientation of the boundaries without polarity, we completely lose any contrast information and obtain the grouping in Figure 9b. Here, the upper edge of the square is grouped as a unique perceptual unit. On the other side, while inserting polarity, the Gabor patches on the upper edge boundary of the black or white region have opposite contrast and detect values of θ, which differs from π (see Figure 9c). There is no affinity between these patches; the first eigenvector of the affinity matrix, represented in red, correctly detects the unit present in the image and corresponds to the inducers of the semicircle (see Figure 9d). This underlines the important role of polarity in perceptual individuation and segmentation. We also note that the first perceptual unit detected is the convex one, as predicted by the gestalt law (see Kanizsa, 1980).

4.3 The Kanizsa Illusory Figures. We consider here stimuli formed by Kanizsa figures, represented by oriented segments that simulate the output of simple cells. Lee and Nguyen (2001) describe the completion of Kanizsa figures taking place in V1.

We first consider the stimulus of Figure 10a. The connectivity among its elements will be analyzed with the kernels defined in equations equations 2.4 and 2.7.

The results of simulations with the fundamental solutions of Fokker-Planck and sub-Riemannian Laplacian equations are shown in Figure 10. The first eigenvector, visualized in red, corresponds to the inducers of the Kanizsa triangle (see Figure 10c). In this example, after computing the first eigenvector of the affinity matrix, this matrix is updated by removing the identified perceptual unit and then computing the first eigenvector of the new matrix (see Figure 10d). These simulations show that circles are associated with the less salient eigenvectors. In that way, the first eigenvalue can be considered a quantitative measure of saliency because it allows the segmentation of the most important object in the scene and the results of simulations confirm the visual grouping. When the affinity matrix is formed by different eigenvector with almost the same eigenvalues, as in Figure 10d, it is not possible to recognize the most salient object because they all have the same influence. Here, we show just one inducer in red. The other two have the same eigenvalue. That also happens, for example, when the inducers are not aligned circularly or are rotated.

Now we consider the Kanizsa square as stimulus and change the angle between the inducers, so that the subjective contours become curved (see Figures 10e–h, second row). The fact that illusory figures are perceived depends on a limit curvature. Indeed, we perceive a square in the first three cases but not the last one. The results of simulations with the fundamental

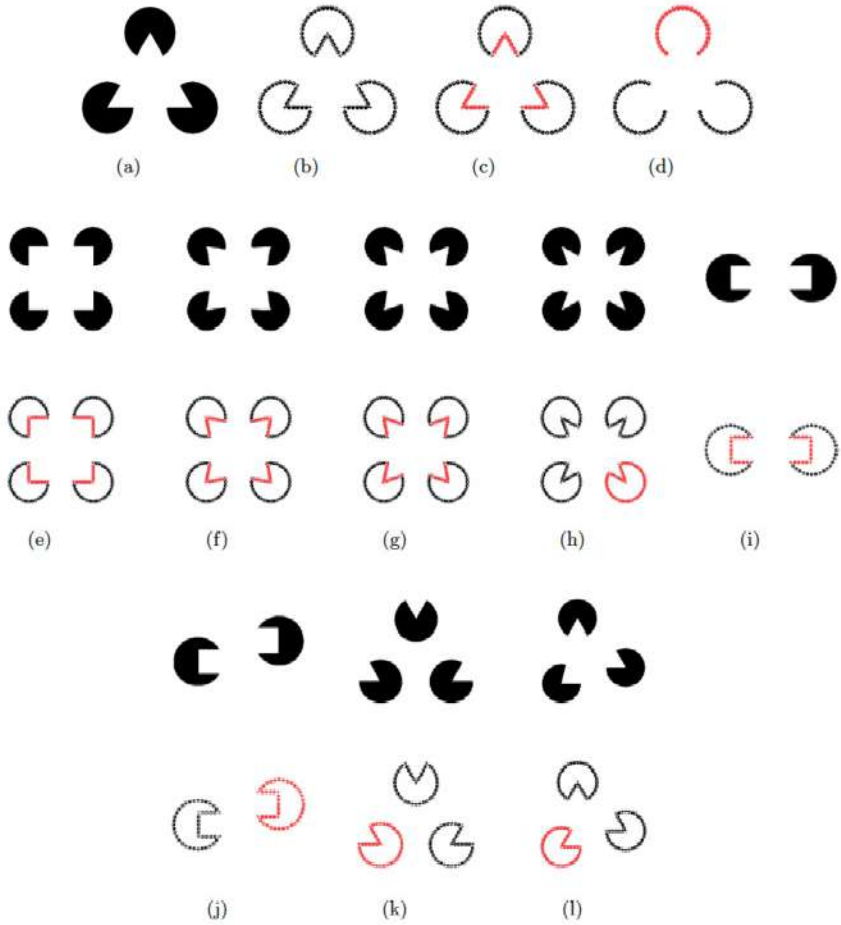

Figure 10: The Kanizsa triangle (a) and the maximally responding odd filters (b). (c) The first eigenvector of the affinity matrix, using the fundamental solutions of Fokker-Planck (see equation 2.4 and 3.1) and sub-Riemannian Laplacian equations 2.7 and 3.2. After this computation, the affinity matrix is updated by removing the detected perceptual unit; the first eigenvector of the new matrix is visualized (d). Examples of stimulus (row 2) with aligned and unaligned inducers. Stimulus with rotated (k) and unaligned (j, l) inducers (row 4). The first eigenvectors of the affinity matrix using the fundamental solutions of Fokker-Planck and sub-Riemannian are visualized in red (rows 3 and 5).

solutions of Fokker-Planck and sub-Riemannian Laplacian equations confirm the visual grouping (see Figures 10e–h, third row). When the angle between the inducers is not too high, Figures 10e to 10g, the first eigenvector corresponds to the inducers that form the square; otherwise, in

Figure 10h, the "Pacman" becomes the most salient object in the image. In this case, we obtain four eigenvectors with almost the same eigenvalue.

Now we consider the Kanizsa bar in Figure 10i, second row) that is perceived only if the inducers are aligned. Also, the result of simulation confirms the visual perception if we use the fundamental solutions of the Fokker-Planck and the sub-Riemannian Laplacian equations (see Figure 10i, third row). When the inducers are not aligned, all the kernels confirm the visual perception, showing two different perceptual units (see Figure 10j).

Considering a stimulus composed of rotated or unaligned inducers, as in Figures 10k and 10l, it is not possible to perceive it, and the results of simulations, using all the connectivity kernels described, confirm the visual grouping. In that case, the affinity matrix is decomposed in three eigenvectors with almost the same eigenvalues, which represent the three perceptual units in the scene.

4.4 Sub-Riemannian Fokker-Planck versus Sub-Riemannian Laplacian. The two kernels we analyze are not mutually exclusive and can be implemented in different cells. The presence of different populations of cells in relation to mathematical models has been also studied in Ben-Shahar and Zucker (2004). We have outlined in sections 2.2 and 3.2 that the Fokker-Planck kernel accounts for long-range connectivity and the sub-Riemannian Laplacian for a short range. In the previous examples, we obtained good results with both kernels, but this difference emerges while we change the parameters. In Figure 11, we compare the action of these two kernels.

In the first row of Figure 11, we see some segments, which form a unique perceptual unit. If they are not too far, the grouping is correctly performed by both the Fokker-Planck and the sub-Riemannian Laplacian (see Figures 11a and 11b). When we separate the inducers, the perceptual unit is correctly detected by the Fokker-Planck kernel (see Figure 11c), while the sub-Riemannian Laplacian is unable to perform the grouping (see Figure 11d). This confirms that the Fokker-Planck kernel is responsible for long-range connectivity. In the second row, we consider an angle. When the angle is sufficiently large, the Fokker-Planck becomes unable to perform the grouping (see Figure 11e), while the sub-Riemannian Laplacian correctly performs the grouping of the perceptual unit (see Figure 11f). This confirms that the sub-Riemannian Laplacian can be used as a model for short-range connectivity.

4.5 Sub-Riemannian versus Riemannian Kernels. In order to further validate the sub-Riemannian model, we show that the model applied with the isotropic Laplacian kernel does not perform correctly. In Figure 12 (top), the visual perception is not correctly modeled: the first eigenvectors coincide with one of the inducers and the squares are not recognized. That also happens for the stimulus of Figure 10a and when the inducers are unaligned circularly or are rotated.

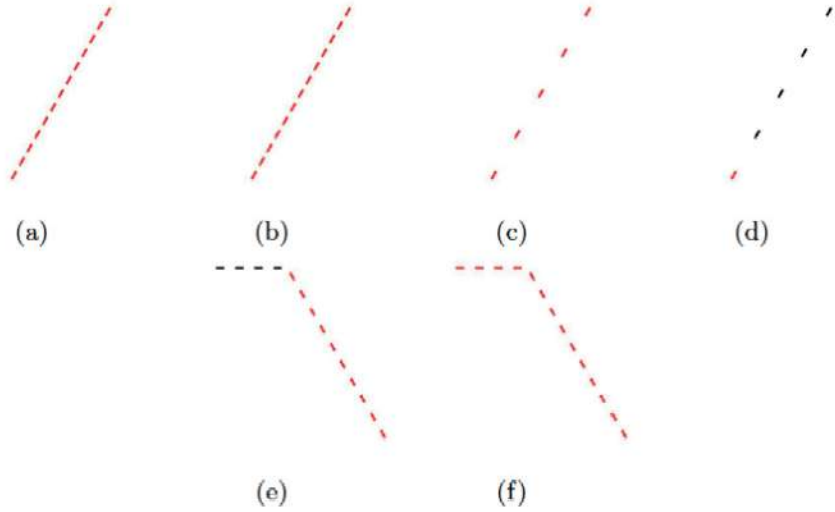

Figure 11: (Top) A few aligned segments that are correctly grouped by the Fokker-Planck and the sub-Riemannian Laplacian (a, b). When we separate the inducers, the perceptual unit is correctly detected using the Fokker-Planck kernel (c), while the sub-Riemannian Laplacian is not able to perform the grouping (d). (Bottom) We consider an angle. In this case, the Fokker-Planck is unable to perform the grouping (e), while the sub-Riemannian Laplacian can correctly perform the grouping (f).

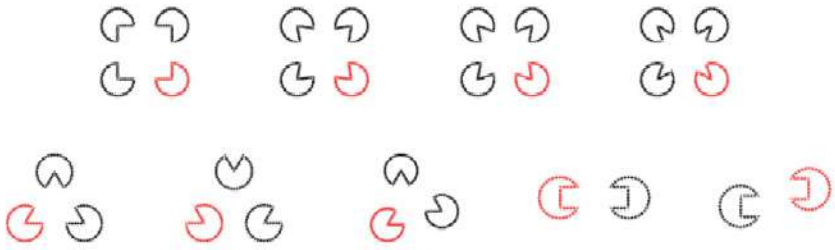

Figure 12: Stimulus of Figure 10. The results do not fit the visual perception if we use the isotropic Laplacian equation (2.11 and 3.3) and confirm the need to use a sub-Riemannian kernel to model the cortical connectivity.

5 Conclusion

In this work, we have presented a neurally based model for figure-ground segmentation using spectral methods, where segmentation has been performed by computing eigenvectors of affinity matrices.

Different connectivity kernels that are compatible with the functional architecture of the primary visual cortex have been used. We have modeled

them as fundamental solutions of Fokker-Planck, sub-Riemannian Laplacian, and isotropic Laplacian equations and compared their properties.

With this model, we have identified perceptual units of different Kanizsa figures, showing that this can be considered a good quantitative model for the constitution of perceptual units equipped by their saliency. We have also shown that the fundamental solutions of Fokker-Planck and sub-Riemannian Laplacian equations are good models for the continuation law, while the isotropic Laplacian equation is less representative for this gestalt law. However, it retrieves information about ladder parallelism, a feature that can be analyzed in the future. All three kernels are able to accomplish boundary, completion, with a preference for the Fokker-Planck and the sub-Riemannian Laplacian operators.

The proposed mathematical model is then able to integrate local and global gestalt laws as a process implemented in the functional architecture of the visual cortex. The kernel considered here depends only on orientation. Hence, it can be applied to detect the salience of geometrical figures, which can be very well described using this feature. The same method can be applied to natural images if their main features are related to orientations, as in retinal images (see Favali et al., 2016). The ideas presented here could be extended to more general kernels able to detect geometrical features different from orientation as curvature (Abbasi et al., 2016), and we are confident that there is a relation between the highest eigenvector and the salient object. However, for general images, we cannot rely on this simple geometric method, since different cortical areas can be involved in the definition of the salience, with a modulatory effect on the connectivity of V1.

Acknowledgments

We are grateful to the anonymous reviewers for their constructive comments, which helped us to improve this letter. The research leading to these results has received funding from the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme FP7/2007-2013/ under REA grant agreement 607643.

References

Abbasi-Sureshjani, S., Favali, M., Citti, G., Sarti, A., & Romeny, B. M. (2016). Cortically-inspired spectral clustering for connectivity analysis in retinal images: Curvature integration. arXiv:1608.08049.

Angelucci, A., Levitt, J. B., Walton, E. J. S., Hupe, J. M., Bullier, J., & Lund, J. S. (2002). Circuits for local and global signal integration in primary visual cortex. *Journal of Neuroscience*, *22*(19), 8633–8646.

August, J., & Zucker S. W. (2000). The curve indicator random field: Curve organization via edge correlation. In K. L. Boyer & Sudeep Sarkar (Eds.), *Perceptual organization for artificial vision systems* (pp. 265–288). New York: Springer.

August, J., & Zucker, S. W. (2003). Sketches with curvature: The curve indicator random field and Markov processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 25*(4), 387–400.

Barbieri, D., Citti, G., Sanguinetti, G., & Sarti, A. (2012). An uncertainty principle underlying the functional architecture of V1. *Journal of Physiology, 106*(5–6), 183–193.

Belkin, M., & Niyogi, P. (2003). Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation, 15*(6), 1373–1396.

Ben-Shahar, O., & Zucker, S. (2004). Geometrical computations explain projection patterns of long-range horizontal connections in visual cortex. *Neural Computation, 16*, 445–476.

Boscain, U., Duplaix, J., Gauthier, J. P., & Rossi, F. (2012). Anthropomorphic image reconstruction via hypoelliptic diffusion. *SIAM Journal on Control and Optimization, 50*(3), 1309–1336.

Bosking, W., Zhang, Y., Schofield, B., & Fitzpatrick D. (1997). Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *Journal of Neuroscience, 17*(6), 2112–2127.

Bressloff, P. C., & Cowan, J. D. (2003). The functional geometry of local and horizontal connections in a model of V1. *Journal of Physiology–Paris, 97*, 221–236.

Bressloff, P. C., Cowan, J. D., Golubitsky, M., Thomas, P. J., & Wiener M. C. (2002). What geometric visual hallucinations tell us about the visual cortex. *Neural Computation, 14*(3), 473–491.

Citti, G., & Sarti, A. (2006). A cortical based model of perceptual completion in the roto-translation space. *Journal of Mathematical Imaging and Vision, 24*(3), 307–326.

Cocci, G., Barbieri, D., Citti, G., & Sarti, A. (2015). Cortical spatio-temporal dimensionality reduction for visual grouping. *Neural Computation, 27*, 1252–1293.

Coifman, R. R., & Lafon, S. (2006). Diffusion maps. *Applied and Computational Harmonic Analysis, 21*(1), 5–30.

Daugman, J. (1985). Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *Journal of the Optical Society of America, A, 2*(7), 1160–1169.

Duits, R., & Franken E. M. (2009). Line enhancement and completion via linear left invariant scale spaces on SE(2). In X.-C. Tai, K. Moerken, M. Lysaker, & K.-A. Lie (Eds.), *Scale space and variational methods in computer vision* (pp. 795–807). New York: Springer.

Duits, R., & van Almsick, M. (2008). The explicit solutions of linear left-invariant second order stochastic evolution equations on the 2D-Euclidean motion group. *Quarterly of Applied Mathematics, 66*(1), 27–67.

Ermentrout, G. B., & Cowan, J. D. (1980). Large scale spatially organized activity in neural nets. *SIAM: SIAM Journal on Applied Mathematics, 38*(1), 1–21.

Favali, M., Abbasi-Sureshjani, S., ter Haar Romeny, B., & Sarti, A. (2016). Analysis of vessel connectivities in retinal images by cortically inspired spectral clustering. *Journal of Mathematical Imaging and Vision, 56* (2016), 158–172.

Field, D., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local association field. *Vision Research, 33*(2), 173–193.

Frégnac, Y., & Shulz, D. E. (1999). Activity-dependent regulation of receptive field properties of cat area 17 by supervised Hebbian learning. *Journal of Neurobiology, 41*(1), 69–82.

Gilbert, C. D., Das, A., Ito, M., Kapadia, M., & Westheimer, G. (1996). Spatial integration and cortical dynamics. *Proceedings of the National Academy of Sciences, 93*(2), 615–622.

Grossberg, S., & Mingolla, E. (1985). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review, 92*(2), 173–211.

Higham, D. J. (2001). An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Review, 43*(3), 525–546.

Hoffman, W. C. (1989). The visual cortex is a contact bundle. *Applied Mathematics and Computation, 32*(2), 137–167.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology, 160*(1), 106–154.

Hubel, D. H., & Wiesel, T. N. (1977). Ferrier lecture: Functional architecture of macaque monkey visual cortex. *Proceedings of the Royal Society of London B: Biological Sciences, 198*(1130), 1–59.

Jerison, D. S., & Sanchez-Calle, A. (1986). Estimates for the heat kernel for a sum of squares of vector fields. *Indiana University Mathematics Journal, 35*, 835–854.

Jones, J. P., & Palmer, L. A. (1987). An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology, 58*(6), 1233–1258.

Kanizsa, G. (1980). *Grammatica del vedere*. Bologna: Il Mulino.

Kellman, P. J., & Shipley, T. F. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology, 23*(2), 141–221.

Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology, 4*, 219–227.

Koenderink, J. J., & van Doorn, A. J. (1987). Representation of local geometry in the visual system. *Biological Cybernetics, 55*(6), 367–375.

Koflka, K. (1935). *Principles of gestalt psychology*. New York: Harcourt.

Kohler, W. (1929). *Gestalt psychology*. New York: Liveright.

Lee, T. S., & Nguyen M. (2001). Dynamics of subjective contour formation in the early visual cortex. *Proceedings of the National Academy of Sciences, 98*, 1907–1911.

Lorenceau, J., & Alais, D. (2001). Form constraints in motion binding. *Nature Neuroscience, 4*(7), 745–751.

Meila, M., & Shi, J. (2001). A random walks view of spectral segmentation. In *Proceedings of the 8th International Workshop on Artificial Intelligence and Statistics*. Delhi: Motilal Banarsidass.

Merleau-Ponty, M. (1945). In *Phenomenology of perception*. Delhi: Motilal Banarsidass.

Mumford, D. (1994). Elastica and computer vision. In C. Bajaj (Ed.), *Algebraic geometry and its applications* (pp. 491–506). New York: Springer.

Parent, P., & Zucker, S. W. (1989). Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 11*(8), 823–839.

Perona, P., & Freeman, W. (1998). A factorization approach to grouping. In *Proceedings of Computer Vision—ECCV'98* (pp. 655–670). New York: Springer.

Petitot, J. (2008). *Neurogeomtrie de la vision: Modeles mathmatiques et physiques des architectures fonctionelles*. Paris: d. cole Polytech.

Petitot, J., & Tondut, Y. (1999). Vers une neurogeometrie: Fibrations corticales, structures de contact et contours subjectifs modaux. *Mathematiques Informatique et Sciences Humaines, 145*, 5–102.

Pillow, J., & Nava, R. (2002). Perceptual completion across the vertical meridian and the role of early visual cortex. *Neuron, 33*, 805–813.

Ringach, D. (2002). Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *Journal of Neurophysiology, 88*(1), 455–463.

Robert, C., & Casella, G. (2013). *Monte Carlo statistical methods*. New York: Springer Science & Business Media.

Roweis, S. T., & Saul, L. K. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science, 290*(5500), 2323–2326.

Sanguinetti, G., Citti, G., & Sarti, A. (2008). Image completion using a diffusion driven mean curvature flow in a sub-Riemannian space. In *Proceedings of the Int. Conf. on Computer Vision Theory and Applications* (pp. 46–53). New York: Springer.

Sarti, A., & Citti, G. (2015). The constitution of visual perceptual units in the functional architecture of V1. *Journal of Computational Neuroscience, 38*(2), 285–300.

Sarti, A., Citti, G., & Petitot, J. (2008). The symplectic structure of the primary visual cortex. *Biological Cybernetics, 98*(1), 33–48.

Shi, J., & Malik, J. (2000). Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 22*(8), 888–905.

Shipley, T. F., & Kellman, P. J. (1992). Perception of partly occluded objects and illusory figures: Evidence for an identity hypothesis. *Journal of Experimental Psychology: Human Perception and Performance, 18*(1), 106.

Shipley, T. F., & Kellman, P. J. (1994). Spatiotemporal boundary formation: Boundary, form, and motion perception from transformations of surface elements. *Journal of Experimental Psychology: General, 123*(1), 3–20.

Von Der Heydt, R., Heitger, F., & Peterhans, E. (1993). Perception of occluding contours: Neural mechanisms and a computational model. *Biomedical Research, 14*, 1–6.

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & von der Heydt, R. (2012). Century of gestalt psychology in visual perception: I. Perceptual grouping and figure-ground organization. *Psychological Bulletin, 138*(6), 1172–1217.

Weiss, Y. (1999). Segmentation using eigenvectors: A unifying view. In *Proceedings of the Seventh IEEE International Conference on Computer Vision* (vol. 2, pp. 975–982). Piscataway, NJ: IEEE.

Wertheimer, M. (1938). *Laws of organization in perceptual forms*. London: Harcourt.

Williams, L. R., & Jacobs, D. W. (1997). Stochastic completion fields. *Neural Computation, 9*(4), 837–858.

Zucker, S. (2006). Differential geometry from the Frenet point of view: Boundary detection, stereo, texture and color. In N. Paraglos, Y. Chen, & O. Faugeras (Eds.), *Handbook of mathematical models in computer vision* (pp. 357–373). New York: Springer.

Received October 15, 2015; accepted October 1, 2016.