Implementation of a Model for Perceptual Completion in $R^2 \times S^1$.

Gonzalo Sanguinetti[1 3], Giovanna Citti[2], and Alessandro Sarti[3]

[1] Universidad de la República, Instituto de Ingeniería Electrica, Montevideo, Uruguay. gsangui@fing.edu.uy
[2] Università di Bologna, Dipartimento di Matematica, Bologna, Italy. citti@unibo.it
[3] Università di Bologna, Dipartimento di Elettronica, Informatica e Sistemistica, Bologna, Italy. asarti@unibo.it

Abstract. In this paper we present an implementation of a perceptual completion model [1] performed in the three dimensional space of position and orientation of level lines of an image. We show that the space is equipped with a natural subriemannian metric. This model allows to perform disocclusion representing both the occluding and occluded objects simultaneously in the space. The completion is accomplished by computing minimal surfaces with respect to the non Euclidean metric of the space. The minimality is achieved via diffusion driven mean curvature flow. Results are presented in a number of cognitive relevant cases.

1 Introduction.

Perceptual completion is performed by the mammalian visual system in a number of phenomenological cases, deeply studied by psychology of Gestalt to understand the underlying structure of visual processing in humans. The most common examples comprehend modal completion, amodal completion, transparency, intersection and self intersection of curves [2]. Modal completion is the process of filling the missing part of an object and building a percept that is phenomenally undistinguishable from real stimuli. It gives rise to the well known phenomenon of illusory boundaries (or subjective contours) and it takes place often to complete occluding objects (in Fig. 1(a) the completed triangle is occluding the 3 circles). Amodal completion (Fig. 1(b)) is a perceptual modality for integrating missing parts of partially occluded objects. Since the occluded figure underlies the occluding one, it is completed without any sensorial counterpart. In case of transparency (Fig. 1(c)) and curve intersection (Fig. 1(d)), both occluding and occluded figures are visible in the scene and the perceptual system is able to disambiguate them and recognize them as different objects. A point made clear by the studies of phenomenology of perception is that in all cases of completion both the occluding and the occluded objects are perceived at the same time in the scene and therefore there are points in the input stimulus corresponding to more than one figure at the perceptual level. Many computer vision techniques have been proposed to model perceptual completion, either heuristically based or

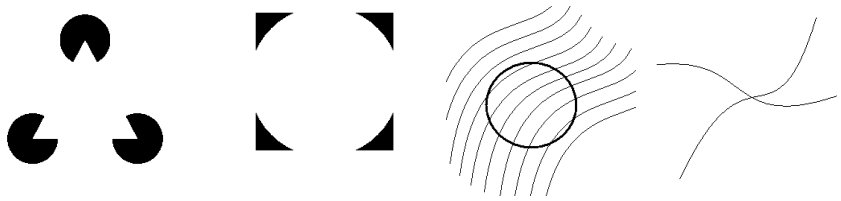

Fig. 1. Some examples of perceptual completion: (a) Modal completion, (b) Amodal completion, (c) Transparency, (d) Curve Intersection.

biologically inspired. Rectilinear and curvilinear subjective contours have been modeled by D.Mumford with Euler elastica as extremality points of curvature functionals [3] and by stochastic fields as solution of the Fokker-Planck equation [4]. In the latter case the stochastic completion field represents the likelihood that a completion joining two contour fragments passes through any given position and orientation in the image. An extension taking into account also the curvature has been proposed in [5]. Amodal completion has been accomplished by a number of techniques. In [6] [7] an extension of the Mumford functional to level lines has been used to fill missing regions. Digital inpainting has been introduced as a technique to diffuse existing information on the boundary toward the interior region [8] [9]. A total variation approach has been proposed in [10]. All these techniques consider the perceptual space in which completion is performed has the same dimensionality of the image. This could be a restriction in case we are interested in the presence of reconstructed occluding and occluded objects in the scene, as in case of mammalian vision. To overcome this restriction, in [1] has been proposed a completion model based on the functional architecture of the visual cortex, where completion is fully performed in the rototranslation group $R^2 \times S^1$, allowing the simultaneous reconstruction of occluding and occluded objects. This model is an improvement of the one proposed in [11] where curves are lifted in the three dimensional Heisenberg group, but using the group of rotations and translations in the plane. In [12] the neural connectivity was modelled as a parallel transport over a tangent fiber bundel. Another model has been introduced in [13] using a tensorial field to complete the missing information. Other high dimensional models are presented in [14] and in [15].

Following [1], a two dimensional image is lifted to a surface in the 3-dimensional sub-Riemannian space, an occlusion is considered as a hole in the surface, and the proposed model complete the missing part of the image with a minimal surface. Computing a minimal surface in the hole and re-projecting it over the image domain, we find the same level lines as Morel and Masnou have found in [6] minimizing an elastica based functional.

In [16] the authors proposed a very fast method for finding the minimal surface explicitly interpolating the level lines represented in the Sub-Riemannian

space even if it is not well suited for simultaneous representation of occluded and occluding objects.

The main objective of this paper is to propose a computational technique for finding minimal surfaces by diffusion driven mean curvature flow. The technique is able to simultaneously construct occluded and occluding objects. The surface is represented as a thin concentrated mass, suitably diffused and concentrated with a two step algorithm adapted to the sub-Riemannian metric. The diffusion driven method was first introduced in the Euclidean settings in [17].

The paper is organized as follows:

- In section 2 we explain the lifting of the image to the 3D position-orientation space and describe the subriemannian structure of the space.
- In section 3 the main model of image completion is proposed and discussed in detail.
- In section 4 we present the numerical scheme for the equations presented in previous sections.
- In section 5 we describe the experiments realized and provide the results obtained.
- Finally, conclusions are presented.

2 Theoretical background.

2.1 Lifting of the image level lines in a 3D space.

An image I can be represented as a bounded function defined on a domain $M \subset \mathbb{R}^2$, $I : M \to \mathbb{R}^+$. The points of M have coordinates (x, y). At every point of the image we detect the tangent direction to the level lines $(I_y, -I_x)$, where I_x and I_y are the components of the image gradient. If θ is the angle between the tangent and the x-axis the tangent can be rewritten as $(\cos(\theta), \sin(\theta))$.

We want to define the orientation independently of the versus of the tangent vector. Therefore, we identify a tangent vector with its opposite one. This means that angles which differ form π will be identified, and

$$\theta(x, y) = -\arctan(I_x/I_y) \ , \ \theta \in S^1,$$

where S^1 is the set of angles $[0, \pi]$.

To every point (x, y) is associated a three dimensional vector (x, y, θ), in a new space homeomorphic to $\mathbb{R}^2 \times S^1$. Since the process is repeated at each point, each level line is lifted to a new curve in the three dimensional space. We will call *admissible curve* a curve in $\mathbb{R}^2 \times S^1$ if it is the lifting of a level line.

2.2 The Tangent Bundle and The Integral Curves.

A tangent vector to the lifted curve has the same two first components as the tangent vector to the level line, i.e. a real multiple of $(\cos(\theta), \sin(\theta))$, and it has the third component in the direction $(0, 0, 1)$. Hence it can be represented as a

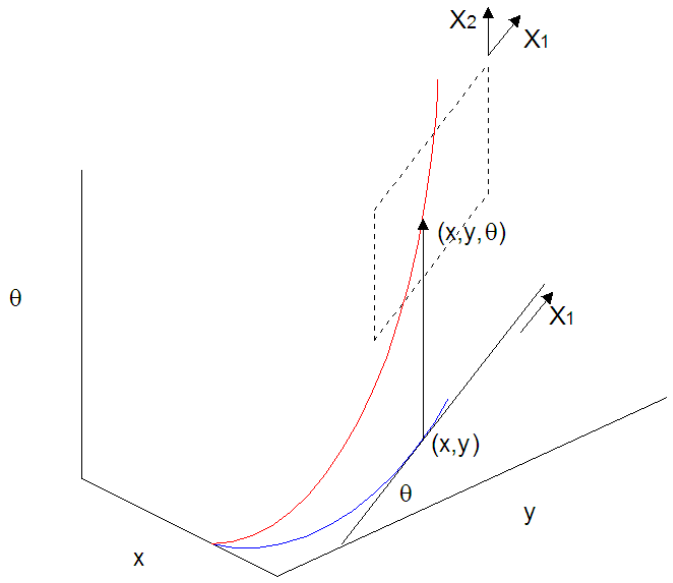

Fig. 2. A lifted level line in $\mathbb{R}^2 \times S^1$ and the tangent space to the point (x, y, θ).

linear combination of the vectors $(\cos(\theta), \sin(\theta), 0)$ and $(0, 0, 1)$ which, from now on, will be called $\boldsymbol{X_1}$ and $\boldsymbol{X_2}$ respectively. The set of vectors $\alpha_1 \boldsymbol{X_1} + \alpha_2 \boldsymbol{X_2}$ defines a plane and every admissible curve is tangent to a vector of the plane. Hence an admissible curve satisfies the differential equation:

$$\gamma'(t) = \alpha_1 \boldsymbol{X_1}(t) + \alpha_2 \boldsymbol{X_2}(t).$$

It is well known that the ratio α_2/α_1 is the the curvature $k(t)$ of its 2D projection, the level line of I.

2.3 Curve length's and metric of the space.

If we equip the tangent planes with an Euclidean metric then the length of an admissible curve can be computed as usual integrating the tangent vector.

$$\lambda(\gamma)(t) = \int_0^t \|\gamma'(s)\| ds = \int_0^t \|\alpha_1 \boldsymbol{X_1} + \alpha_2 \boldsymbol{X_2}\| ds = \int_0^t \alpha_1 \sqrt{1 + k^2} ds. \qquad (1)$$

In order to define a distance in term of the length, we need to answer the following question: *Is it possible to connect every couple of points of* $\mathbb{R}^2 \times S^1$ *using an integral curve?*

This is not a simple question taking into account that in every point we have only directions which are linear combinations of two vectors even if we are immersed in a three dimensional space. However, the answer is yes and it will become clear in the example below. Otherwise, see [1] for a detailed justification.

Consequently, it is possible to define a notion of distance between two points $p_0 = (x_0, y_0, \theta_0)$ and $p_1 = (x_1, y_1, \theta_1)$:

$$d(p_0, p_1) = \inf\{\lambda(\gamma) : \gamma \text{ is an admmisible curve connecting } p_0 \text{ and } p_1\}. \quad (2)$$

In the Euclidean case this infimum is realized by a geodesic that is a segment. Here, the geodesics are locally curvilinear. The metric induced by (2) is clearly Non-Euclidean, moreover it is not even Riemannian. With the chosen metrics on the tangent plane, the space co-metric is given by:

$$g = \begin{pmatrix} \cos(\theta) & 0 \\ \sin(\theta) & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \cos(\theta) & \sin(\theta) & 0 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} \cos^2(\theta) & \cos(\theta)\sin(\theta) & 0 \\ \cos(\theta)\sin(\theta) & \sin^2(\theta) & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Since the matrix g is not invertible, it can not induce a Riemannian metric on the space. Spaces equipped with Sub-Riemannian metrics appears often when one of the dimensions is a state variable depending on the others. In this case the state variable is θ.

2.4 The lifted surface as an implicit function.

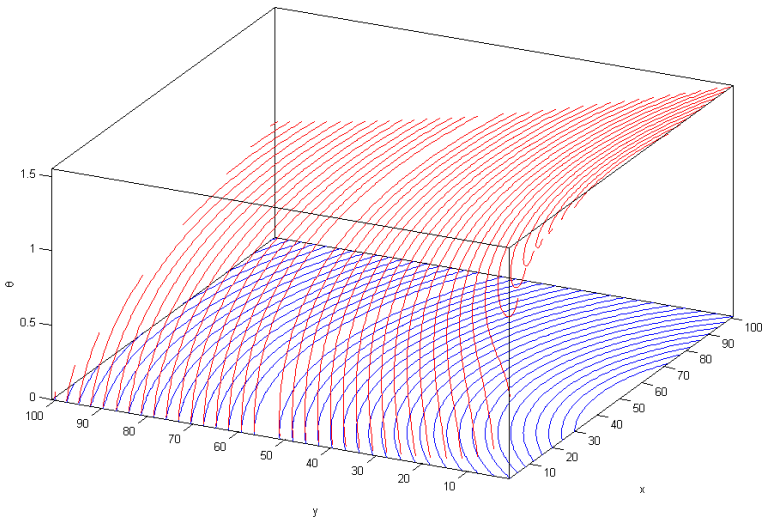

Fig. 3. A lifted image is a surface foliated by the lifted level lines.

When every point of an entire image is lifted up, a three dimensional surface is constructed as:

$$\Sigma = \left\{(x, y, \theta) \in \mathbb{R}^2 \times S^1 : \theta(x, y) = -\arctan(I_x/I_y)\right\}.$$

We can identify the lifting of an image with the lifting of every level line. This point of view allows us to understand a remarkable property of the lifted surface. In fact, since two level lines of an image never cross, also the lifted level lines don't do it. Then we say that the lifted surface is foliated by the lifted curves (see Fig. 3). We will call *rule* an admissible curve foliating a surface.

Let's now represent the surface in terms of the implicit function:

$$u(x, y, \theta) = \left[\cos\left(\theta + \arctan(I_x/I_y)\right)\right]^2 . \tag{3}$$

For every coordinate (x, y) this function attains its maximum in the variable θ in correspondence to a point $(x, y, \bar{\theta})$ of the surface. The cosine function is chosen

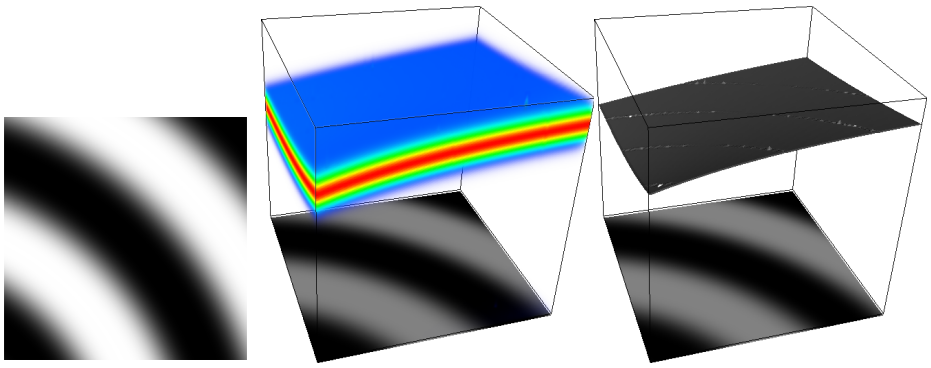

Fig. 4. The lifted image can be viewed as a thick surface and the surface obtained with eq. 4

in order to have periodicity of u in the third coordinate since it is an angle. Note we have imposed that the maximum value of u is 1.

The surface Σ can be represented as the zero level set of the function u_θ:

$$\Sigma = \{(x, y, \theta) \in \mathbb{R}^2 \times S^1 : \partial_\theta u(x, y, \theta) = 0, \partial_{\theta\theta} u(x, y, \theta) < 0\}. \tag{4}$$

The condition over $\partial_{\theta\theta} u$ is imposed in order to avoid minima of u.

2.5 Sub-riemannian differential operators.

We will define differential operators acting over the function u, in terms of the subriemannian structure introduced before on the space $\mathbb{R}^2 \times S^1$, instead of the Euclidean one. We will need to define two differential operators X_1 and X_2 which play the role of the Euclidean partial derivatives, and have the same coefficients as the vector fields $\boldsymbol{X_1}$ and $\boldsymbol{X_2}$. Hence

$$X_1 = \cos(\theta)\partial_x + \sin(\theta)\partial_y, \quad X_2 = \partial_\theta.$$

Accordingly we define the Sub-Riemannian gradient as:

$$\nabla_{\text{SR}} u = (X_1 u, X_2 u).$$

The notation SR (Sub-Riemannian) will be used in order to avoid confusions with the classical operators. We define the so called sub-laplacian operator, which is the analogous of the classical laplacian in this structure:

$$\Delta_{\text{SR}} u = X_1^2 u + X_2^2 u = \cos^2(\theta) u_{xx} + \sin^2(\theta) u_{yy} + 2\cos(\theta)\sin(\theta) u_{xy} + u_{\theta\theta}. \quad (5)$$

and we define the subriemannian diffusion equation as:

$$u_t = \Delta_{\text{SR}} u.$$

Despite of the fact the sublaplacian operator is built just with two directional derivatives in a 3 dimensional space, the diffusion process reaches every point due to the connectivity property of the sub-riemannian geometry.

2.6 Differential geometry of the surface.

Since the surface Σ, is the zero level set of the function $u_\theta = X_2 u$, it is possible to define geometrical properties of Σ, in terms of the function u_θ and its sub-riemannian derivatives. The subriemannian gradient $\nabla_{\text{SR}} u_\theta$ is orthogonal to the surface (w.r. of the subriemannian metric), and an admissible tangent vector is $(-X_2 u_\theta, X_1 u_\theta)$. Correspondingly the rules on the surface have the expression

$$\gamma' = -X_2 u_\theta \boldsymbol{X_1} + X_1 u_\theta \boldsymbol{X_2}. \quad (6)$$

Analogously the diffusion on the surface, which is the diffusion along the rules, is expressed in terms of $\nabla_{SR} u_\theta$.

The foliation feature suggests a natural notion of area in the sub-riemannian structure $\mathbb{R}^2 \times S^1$. Indeed the area of a lifted surface can be defined as the integral of the lengths of every rule. With this definition, a minimal surface with assigned boundary conditions is obtained requiring every rule to have minimal length.

3 The completion model.

3.1 Basic model.

In this section we present our completion model in the rototraslation group, (see also [1]).

Let's consider an image with an occlusion and let us call D the missing part in the two dimensional domain. In order to complete it, we lift the image to a surface in the Sub-Riemannian space. This lifted surface will have a hole, which will be completed with a minimal surface. Indeed, using relation (1), in [1] it has been proved that the subriemannian minimization of the surface area gives

rise to the minimization on the rules on the surfaces, whose projection are the elastica curves. Hence the minimization of the first order area functional on $\mathbb{R}^2 \times S^1$ correspond to the minimisation of a second order curvature functional on the image plane [7] [6].

The method we will use is the following: first we lift the non occluded part of the image with eq. (3) to a function u defined on $(\mathbb{R}^2 \setminus D) \times S^1$. In the occluded region $D \times S^1$ we assign value zero to the function u. Later we built an initial surface in the missing region. Finally we evolve this surface with an approximated diffusion driven mean curvature flow until it becomes minimal. This is a two step algorithm of diffusion and concentration, as shown in [1]:

- Diffusion of existing information in the subriemannian space with the sub-laplacian.
- Concentration of diffused information on the fiber S^1 over every point (x, y).

3.2 Algorithmic implementation.

The image I is lifted to a surface, represented by the maxima over the fiber S^1 of a function u, by using equation (3) The first step is to propagate existing information from the boundary of the missing region $D \times S^1$ with sub-riemannian diffusion:

$$\begin{cases} \partial_t u = \begin{array}{ll} \Delta_{\text{SR}} u & \text{if } (x, y, \theta) \in D \times S^1 \\ \partial_{\theta\theta} u & \text{if } (x, y, \theta) \in (\mathbb{R}^2 \setminus D) \times S^1 \end{array}, t \in [0, h] \\ u(0) = u_0 \end{cases} \tag{7}$$

This first step is necessary to initialize the function u to be a rough solution, which will be refined by diffusion driven mean curvature flow.

In fact after the initial propagation, a mean curvature evolution of the function u is implemented by using a two step iterative algorithm consisting in alternative diffusion and concentration:

- Diffuse with the Sub-Laplacian operator (5) for a short time with fixed boundary conditions in the boundary of $D \times S^1$.
  In the occluded region we diffuse using the sub-Laplacian operator. This operator propagates data in the direction of the vectors X_1 and X_2. The diffusion in the direction of X_1 alone would expand into the occlusion the information taken from the boundary just in a straight line parallel to the (x, y) plane. By adding the diffusion in the X_2 direction, we allow propagation on curvilinear paths on $\mathbb{R}^2 \times S^1$, even if we make thicker the surface represented by u as a side effect. Outside $D \times S^1$ we use the equation $u_t = u_{\theta\theta}$ just to keep the same thickness of the surface as in the interior of $D \times S^1$. Note that if we just use this equation for a short time the maximum of u is not moved and therefore the surface Σ does not change. For the disocclution problem it is only necessary to consider values of u near the boundary of $D \times S^1$. Only this values will be propagated inside $D \times S^1$. Nevertheless, for improving the visualization we will consider a larger domain outside $D \times S^1$.

– Concentrate the function u over the surface, i.e. make thinner the thick version of the surface.

After diffusing u for a period of time h, we perform a concentration over its maximum and denote $\bar{u}$ the new function which implicitly define the concentrated surface:

$$\bar{u}(x, y, \theta) = \left(\frac{u(x, y, \theta)}{u_{\max}(x, y)} \right)^{\gamma}, \ \gamma > 1 \tag{8}$$

where:

$$u_{\max}(x, y) = \max_{\theta \in S^1} \{u(x, y, \theta)\}. \tag{9}$$

This procedure renormalize the function u in such a way that the maximum over each fiber is 1. The concentration, obtained elevating the function u to a suitable power greater than one, preserves the value of the maximum and reduces all the other values of u. Thus this mechanism concentrates the function around its maximum.

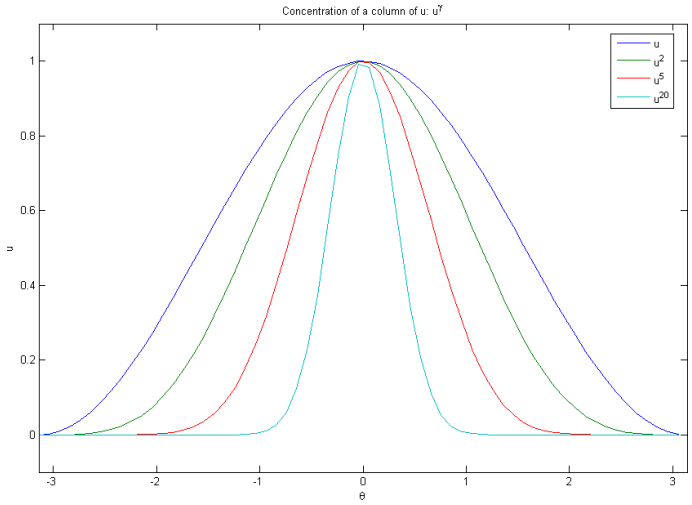

Fig. 5. Example of the concentration proccess of a single fiber.

3.3 Multiple concentration.

The three dimensionality of the space allows the coexistence of occluded and occluding objects at the same time. In terms of the function u it means that we expect to have more than one maximum in each fiber. However, the equation described before (9), allows only one maximum per fiber. The method described above could be slightly modified in order to avoid this limitation. In particular we propose the following renormalization criterion.

We first detect the maxima on a fiber over the point (x, y) as the set $\{\theta \in S^1, \partial_\theta u(x, y, \theta) = 0, \partial_{\theta\theta} u(x, y, \theta) < 0\}$. We call them $\theta_1, \ldots, \theta_n$ with $\theta_i < \theta_{i+1}$. Then we construct a piecewise linear function u_{norm} (Fig. 6) connecting every local maximum detected and periodic in the variable θ:

$$u_{\text{norm}}(x, y, \theta) = u(x, y, \theta_j) + (\theta - \theta_j)\frac{u(x, y, \theta_{j+1}) - u(x, y, \theta_j)}{\theta_{j+1} - \theta_j} \tag{10}$$

with $\theta \in [\theta_j, \theta_{j+1}]$.

We use eq (10) to re-normalize every single column of u as follows:

$$\bar{u}(x, y, \theta) = \left(\frac{u(x, y, \theta)}{u_{\text{norm}}(x, y, \theta)}\right)^\gamma, \gamma > 1.$$

After renormalization the function $\bar{u}$ keep the same points of maximum as the function u and attains value 1 at each of these points.

As we mentioned before, this modification allows more than one maximum on each fiber. Hence applying iteratively this improved concentration technique and

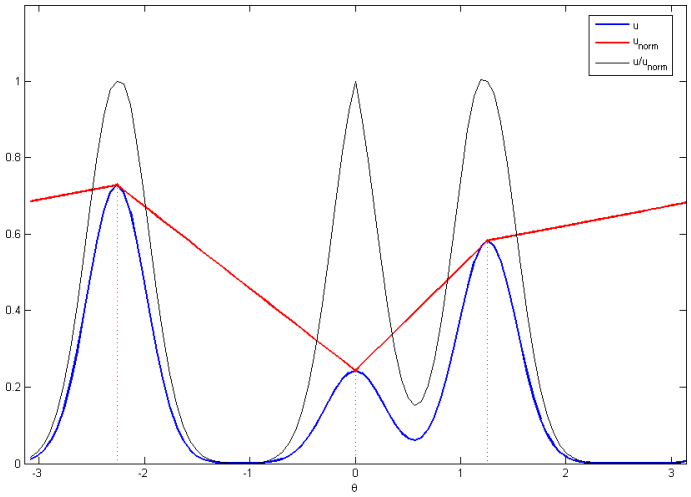

Fig. 6. Example of the improved re-normalization proccess of a single fiber.

the sub-riemannian diffusion, we compute minimal surfaces, in $\mathbb{R}^2 \times S^1$ which are union of graphs of the variable (x, y), which can partially overlap. It corresponds to the completion of both occluding and occluded object.

4 Numerical scheme.

For the diffusion we use a finite difference scheme. Let us consider a rectangular grid in space-time (x, y, θ, t). The grid consist of a set of points $(x_l, y_m, \theta_q, t_n) = (l\Delta x, m\Delta y, q\Delta\theta, n\Delta t)$.

Following the standard notation, we denote by u^n_{lmq} the value of the function u at a grid point. We use forward differences in order to approximate the time derivative:

$$D_t u = \frac{u^{n+1}_{lmq} - u^n_{lmq}}{\Delta t}$$

and center differences for the spatial ones:

$$D_x u^n_{lmq} = \frac{u^n_{(l+1)mq} - u^n_{(l-1)mq}}{2\Delta x},\ D_{xx} u^n_{lmq} = \frac{u^n_{(l+1)mq} - 2u^n_{lmq} + u^n_{(l-1)mq}}{(\Delta x)^2}.$$

The second directional derivatives are approximated with:

$$\begin{aligned} D_{11} u^n_{lmq} &= \cos(\theta_q)^2 D_{xx} u^n_{lmq} + \sin(\theta_q)^2 D_{yy} u^n_{lmq} \\ &\quad +2\cos(\theta_q)\sin(\theta_q) D_{xy} u^n_{lmq} \\ D_{22} u^n_{lmq} &= D_{\theta\theta} u^n_{lmq} \end{aligned}$$

We impose Neumann boundary conditions on x and y and periodic boundary conditions on the third direction θ. The time step Δt is upper bounded by the usual Courant-Friedrich-Levy condition that ensures the stability of the evolution [11].

5 Experiments and results.

5.1 Macula cieca example.

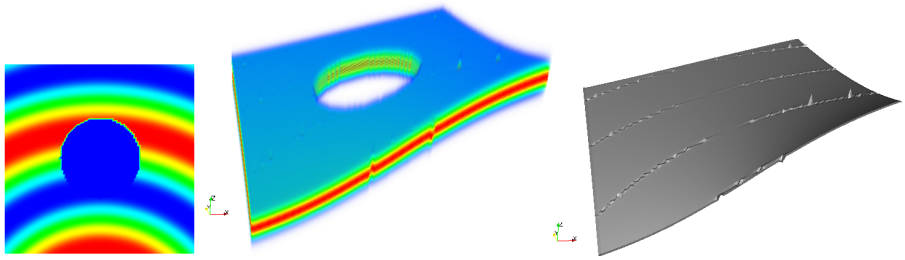

Fig. 7. Macula cieca example: original image, Initially lifted surface and the minimal surface computed.

In this experiment we consider the completion of a figure that has been partially occluded. This example mimics the missing information due to the presence of the macula cieca (blind spot) that is modally completed by the human visual system. As described in the previous section the occluded image is lifted to a surface with a hole in the three dimensional space and an initial surface is defined in the missing part with a classical Euclidean diffusion equation. Then

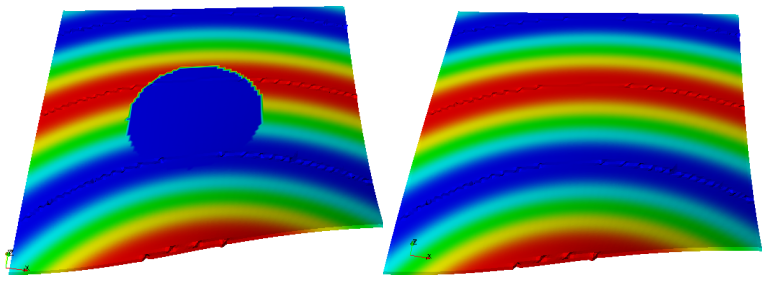

Fig. 8. Gray level diffusion in the macula cieca example.

the surface is evolved applying iteratively equations (7) and (9) until a steady state is achieved.

The image dimensions are 100×100 pixels, and we use 100 values to discretize the variable θ. For the preprocessing step 100 iterations of the Euclidean heat equation were made using a time step of $\Delta t = 0.1$. The steady state was reached after 20 iterations with a concentration power in (8) of $\gamma = 2$ and 20 steps with $\Delta t = 0.1$ of the subriemannian heat equation (7).

At this point we have completed the missing information of the lifted surface with a minimal surface in the Sub-Riemannian space. The lifting and completion processes take into account just the direction of the level lines of the image, as a geometric information. Then the intensity information of the image is completely missed.

Let's define a function v extending the values of the image I on the 3 dimensional space, and constant in the variable θ:

$$v(x,y,\theta) = \begin{cases} I(x,y) & (x,y,\theta) \in (\mathbb{R}^2 \backslash D) \times S^1 \\ 0 & (x,y,\theta) \in D \times S^1 \end{cases}$$

We will use a Laplace Beltrami diffusion algorithm in the sub-riemannian setting to propagate the function v along the rules of the minimal surface. Since the rules of the surface, defined in (6) only depend on $\nabla_{SR} u_\theta$, the Laplace Beltrami operator is a linear operator in the variable v whose coefficients depend on $\nabla_{SR} u_\theta$:

$$v_t = \frac{|X_2 u_\theta|^2 X_1^2 v + |X_1 u_\theta|^2 X_2^2 v}{X_1^2 u_\theta + X_2^2 u_\theta} - \frac{X_1 u_\theta X_2 u_\theta X_1 X_2 v - X_1 u_\theta X_2 u_\theta X_2 X_1 v}{X_1^2 u_\theta + X_2^2 u_\theta}.$$

5.2 Occlusion example.

In Figure 9 an occlusion problem is considered. The initial image (top) shows an underlying object partially occluded by a vertical stripe. The human visual system simultaneously segments the occluding object and amodally completes the occluded one, taking both at the same time as perceived units. In the numerical experiment first the image is lifted in the sub-riemannian space and the missing

information is completed. The result shows that the partially occluded object has been completed and the occluding one has been segmented. Both objects are present at the same time in the three dimensional space.

For this example the dimensions were again $100 \times 100 \times 100$ pixels. Non preprocessing step is needed. The steady state was reached after 10 iterations with a concentration power of $\gamma = 2$ in equation 8 and 10 steps with $\Delta t = 0.1$ of the subriemannian diffusion step.

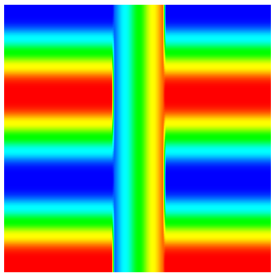

(a) Original Image.

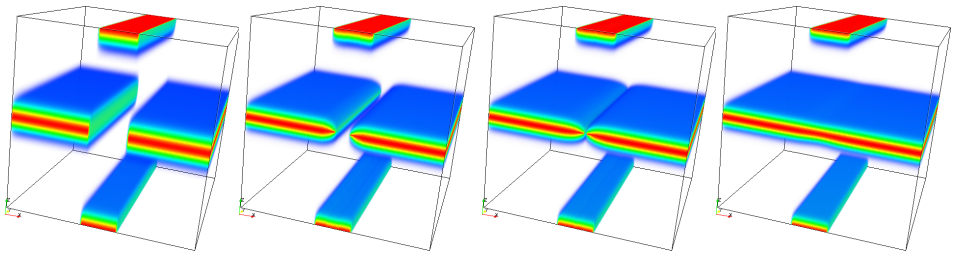

(b) Some steps of the evolution.

Fig. 9. Occlusion example: Mean Curvature Evolution with 2 simultaneous surfaces.

6 Conclusion.

In this paper we utilized a model of perceptual completion inspired from the visual cortex to perform completion of occluding and occluded objects in images. In particular we achieved the task by computing minimal surfaces in subriemannian space via diffusion driven mean curvature flow. The implementation has been performed with a two steps iterative algorithm of diffusion and concentration. A new concentration technique allowing more than one maximum over the fibers has been proposed. This allows to compute a set of graphs partially overlapped representing the occluding and the occluded objects. Computational results on cognitive images have been achieved.

Acknowledgements.

This work was partially supported by ALFA project II-0366-FA and NEST project GALA (Sub-Riemannian geometric analysis in Lie groups) number 028766.

References

1. G. Citti and A. Sarti. A cortical based model of perceptual completion in the roto-translation space. *Journal of Mathematical Imaging and Vision*, 24(3):307–326, May 2006.
2. G. Kanisza. *Organization in Vision: Essays on Gestalt Perception*. Praeger, New York, NY, 1979.
3. M. Nitzberg and D. Mumford. The 2.1-D sketch. In *International Conference on Computer Vision*, pages 138–144, 1990.
4. Lance R. Williams and David W. Jacobs. Stochastic completion fields: A neural model of illusory contour shape and salience. In *ICCV*, pages 408–415, 1995.
5. Jonas August and Steven W. Zucker. Sketches with curvature: The curve indicator random field and markov processes. *IEEE Trans. Pattern Anal. Mach. Intell*, 25(4):387–400, 2003.
6. Simon Masnou and Jean-Michel Morel. Level lines based disocclusion. In *ICIP (3)*, pages 259–263, 1998.
7. L. Ambrosio and S. Masnou. On a variational problem arising in image reconstruction. 2005.
8. Marcelo Bertalmío, Guillermo Sapiro, Vicent Caselles, and Coloma Ballester. Image inpainting. In *SIGGRAPH*, pages 417–424, 2000.
9. Coloma Ballester, Marcelo Bertalmío, Vicent Caselles, Guillermo Sapiro, and Joan Verdera. Filling-in by joint interpolation of vector fields and gray levels. *IEEE Transactions on Image Processing*, 10(8):1200–1211, 2001.
10. T. F. Chan and J. Shen. Mathematical models for local nontexture inpaintings. *Journal of Applied Mathematics*, 62(3):1019–1043, 2001.
11. Jean Petitot and Yannick Tondut. Vers une neurogéométrie. fibrations corticales, structures de contact et contours subjectifs modaux, 1999.
12. Ohad Ben-Shahar and Steven W. Zucker. Geometrical computations explain projection patterns of long-range horizontal connections in visual cortex. *Neural Computation*, 16(3):445–476, 2004.
13. G. Medioni. Tensor voting: Theory and applications, 2000.
14. Alessandro Sarti, Giovanna Citti, and Jean Petitot. The symplectic structure of the primary visual cortex. *Biological Cybernetics*, 98(1):33–48, 2008.
15. Erik Franken, Remco Duits, and Bart M. ter Haar Romeny. Nonlinear diffusion on the 2d euclidean motion group. In Fiorella Sgallari, Almerico Murli, and Nikos Paragios, editors, *SSVM*, volume 4485 of *Lecture Notes in Computer Science*, pages 461–472. Springer, 2007.
16. Robert K. Hladky and Scott D. Pauls. Minimal surfaces in the roto-translation group with applications to a neuro-biological image completion model, September 27 2005. Comment: 35 pages, 15 figures.
17. B. Merriman, J. K. Bence, and S. J. Osher. Diffusion generated motion by mean curvature. *Computational Crystal Growers Workshop, J. Taylor Sel. Taylor (Ed).*, 1998.