Image inpainting via a control-theoretical model of human vision

U. Boscain¹, J.P. Gauthier² and D. Prandi³

Abstract—In this paper we consider several algorithms for image inpainting based on the hypoelliptic diffusion naturally associated with a mathematical model of the primary visual cortex. In particular, we present one algorithm that does not exploit the information of where the image is corrupted, and others that do it. While the first algorithm is able to reconstruct only images that our visual system is still capable of recognize, we show that those of the second type completely transcend such limitation providing reconstructions at the state-of-the-art in image inpainting. This can be interpreted as a validation of the fact that our visual cortex actually encodes the first type of algorithm.

Index Terms—mage inpainting, sub-Riemannian geometry, neurogeometry, hypoelliptic diffusion, image inpainting, sub-Riemannian geometry, neurogeometry, hypoelliptic diffusion.

I. INTRODUCTION

Since the founding works of Petitot, Citti and Sarti [25], [12], where the foundations of the mathematical model (known as Citti-Petitot-Sarti model, CPS model in short) of the human primary visual cortex V1 have been posed, many authors worked on develop and apply these ideas to image processing and computer vision [16], [8], [18]. Indeed, the main concept at the basis of the CPS model, i.e., that images in our brains are processed via a redundant representation taking into account local features as local orientation, presents a somewhat new framework in which many tasks can be strongly simplified.

In view of the many well-studied examples of contour completion operated by the human brain, one of the main problems that was attacked from this point of view has been that of image reconstruction, also known as image inpainting [2], [23], [10], [13]. In this paper, we present and extend on the results we obtained in a series of works [6], [4], [5], [3], [26] in this context. The paper is divided in three parts, corresponding to different methods and purposes of image reconstruction.

In the first part, Section II, we introduce the geometry of the problem by describing the Citti-Petitot-Sarti model for reconstruction of curves, with some new ingredients introduced in [6]. Although this simplified model is not very effective to reconstruct images partially corrupted (reconstructing an image level curve by level curve is a hopeless problem), it permits to understand easily the sub-Riemannian structure of V1, which will allow to introduce the hypoelliptic diffusion.

Indeed, in Section III, we present an image inpainting algorithm obtained by building on the curve reconstruction methods of the first part. In particular, we show that, via stochastic considerations, the latter can be extended to image reconstruction by considering the hypoelliptic diffusion associated with the a sub-Riemannian structure of V1. Although not very efficient, this method allows to reconstruct images with small corruptions. The main interest of this technique is to simulate the reconstruction of images by the visual cortex and, as we will show in the next section, to produce very effective image inpainting algorithms when coupled

with other techniques. Moreover, thanks to the exploitation of non-commutative harmonic analysis techniques, the final algorithm that we present has the twofold advantage of, on one side, being very fast and easy to parallelize, and on the other, of not needing nor exploiting any information on the location and shape of the corruption.

Finally, in the third and last part, Section IV, we collect some results detailing how to ameliorate the above hypoelliptic diffusion method by taking into account the location of the corruption. In particular, we show how these methods allow to obtain image reconstructions of highly corrupted images (e.g. with up to 97% of pixel missing), placing themselves at the same level as state-of-the-art inpainting methods.

We conclude this brief introduction by observing that, experimentally, the threshold of the maximal amount of corruption that the pure hypoelliptic diffusion method can manage looks very close to the threshold of corruption above which our brain can recognize the underlying image. This fact could be seen as a validation of the CPS model with pure hypoelliptic diffusion. On the other hand, the methods presented in the last part completely transcend such problem, thus suggesting that, although based on neurophysiological ideas, they do not reflect any real mechanism of V1.

II. THE SUB-RIEMANNIAN MODEL FOR CURVE RECONSTRUCTIONS

In this section, we recall a model describing how the human visual cortex V1 reconstructs curves which are partially hidden or corrupted. The model we present here was initially due to Petitot [25], [24]. It is based on previous work by Hubel-Wiesel [20] and Hoffman [19], then it was refined by Citti *et al.* [12], [11], Duits *et al.* [16], [15], and by the authors of the present paper in [4], [6], [5], [3]. It was also studied by Hladky and Pauls in [18].

In a simplified model¹ (see [24, p. 79]), neurons of V1 are grouped into *orientation columns*, each of them being sensitive to visual stimuli at a given point of the retina and for a given direction² on it. The retina is modeled by the real plane, i.e. each point is represented by $(x, y) \in \mathbb{R}^2$, while the directions at a given point are modeled by the projective line, i.e. $\theta \in P^1$. Hence, the primary visual cortex V1 is modeled by the so called *projective tangent bundle* $PT\mathbb{R}^2 := \mathbb{R}^2 \times P^1$. From a neurological point of view, orientation columns are in turn grouped into *hypercolumns*, each of them being sensitive to stimuli at a given point (x, y) with any direction. In the same hypercolumn, relative to a point (x, y) of the plane, we also find neurons that are sensitive to other stimuli properties, like colors, displacement directions, etc... In this paper, we focus only on directions and therefore each hypercolumn is represented by a fiber P^1 of the bundle $PT\mathbb{R}^2$.

Orientation columns are connected between them in two different ways. The first kind of connections are the *vertical* (inhibitory) ones, which connect orientation columns belonging to the same hypercolumn and sensible to similar directions. The second kind of

¹ CNRS, LJLL, Sorbonne Université, Paris, France and Équipe INRIA Paris CAGE. ugo.boscain@upmc.fr

² LSIS, Université de Toulon, France. gauthier@univ-tln.fr

³ CNRS, L2S, CentraleSupélec, Gif-sur-Yvette, France. dario.prandi@l2s.centralesupelec.fr

¹For example, in this model we do not take into account the fact that the continuous space of stimuli is implemented via a discrete set of neurons.

²Geometers call “directions” (angles modulo π) what neurophysiologists call “orientations”.

Fig. 1: A scheme of the primary visual cortex V1.

connections are the *horizontal* (excitatory) connections, which connect neurons belonging to different (but not too far) hypercolumns and sensible to the same directions. (See Figure 1.) These two types of connections are represented by the following vector fields on $PT\mathbb{R}^2$:

$$X(x, y, \theta) = \begin{pmatrix} \cos \theta \\ \sin \theta \\ 0 \end{pmatrix} \quad \text{and} \quad \Theta(x, y, \theta) = \begin{pmatrix} 0 \\ 0 \\ \beta \end{pmatrix}.$$

The parameter $\beta > 0$ is introduced here to fix the relative weight of the horizontal and vertical connections, which have different physical dimensions.

In other words, when V1 detects a (regular enough) planar curve $\gamma : [0, T] \rightarrow \mathbb{R}^2$, $\gamma(t) = (x(t), y(t))$, it computes a “lifted curve” $\Gamma : [0, T] \rightarrow PT\mathbb{R}^2$, $\Gamma(t) = (x(t), y(t), \theta(t))$, by including a new variable $\theta(\cdot) : [0, T] \rightarrow P^1$ which satisfies:

$$\frac{d}{dt}\Gamma = uX(\Gamma) + v\Theta(\Gamma), \quad \text{for some } u, v : [0, T] \rightarrow \mathbb{R} \quad (1)$$

The new variable $\theta(\cdot)$ plays the role of the direction in P^1 of the tangent vector to the curve. Here it is natural to require that $u, v \in L^1([0, T])$, i.e., that $\Gamma \in W^{1,1}([0, T])$.

Definition 1: A *liftable curve* is a planar curve $\gamma : [0, T] \rightarrow \mathbb{R}^2$ such that $\gamma \in W^{1,1}([0, T], \mathbb{R}^2)$ and there exists a unique $\theta(\cdot) \in W^{1,1}([0, T], P^1)$ such that (1) holds for some $u, v \in L^1([0, T])$.

In particular, for liftable curves $\gamma(\cdot) = (x(\cdot), y(\cdot))$ it holds

$$\theta(t) = \arctan \frac{\dot{y}(t)}{\dot{x}(t)}, \quad \text{for a.e. } t \in [0, T]. \quad (2)$$

A. The curve reconstruction problem

Consider an “interrupted” liftable curve, that is, assume that $\gamma : [0, a] \cup [b, T] \rightarrow \mathbb{R}^2$ where $0 < a < b < T$. Let us call $\gamma_{\text{in}} = \gamma(a)$, $(x_{\text{in}}, y_{\text{in}}) := (x(a), y(a))$, $\gamma_{\text{fin}} = \gamma(b)$, and $(x_{\text{fin}}, y_{\text{fin}}) := (x(b), y(b))$. Notice that, according to (2) the limits $\theta_{\text{in}} := \lim_{\tau \uparrow a} \theta(\tau)$ and $\theta_{\text{fin}} := \lim_{\tau \downarrow b} \theta(\tau)$ are well defined. In the following, we describe a method to reconstruct the missing part, based on the model presented above. That is, we are looking for a curve $\tilde{\gamma} : [a, b] \rightarrow \mathbb{R}^2$ whose support is a “reasonable” completion of the support of γ .

It has been proposed by Petitot [24] that the visual cortex reconstructs the curve by minimizing the energy necessary to activate orientation columns which are not activated by the curve itself. This is modeled by the minimization of the following cost functional,

$$J = \int_a^b (u(\tau)^2 + v(\tau)^2) d\tau \rightarrow \min. \quad (3)$$

Here, a, b are fixed, and the minimum is taken on the set of liftable curves $\tilde{\gamma} : [a, b] \rightarrow \mathbb{R}^2$ which are solution of (1) for some $u, v \in L^1([a, b])$ and satisfying boundary conditions

$$\tilde{\gamma}(a) = \gamma_{\text{in}}, \quad \tilde{\gamma}(b) = \gamma_{\text{fin}}, \quad \tilde{\theta}(a) = \theta_{\text{in}}, \quad \tilde{\theta}(b) = \theta_{\text{fin}}. \quad (4)$$

Here $\tilde{\theta}(\cdot)$ is the lift associated with $\tilde{\gamma}(\cdot)$ via (2). Observe that in the cost (3), $u(\tau)^2$ (resp. $v(\tau)^2$) represents the (infinitesimal) energy necessary to activate horizontal (resp. vertical) connections. As pointed out in [6], a solution to the problem (1), (3), (4) always exists.

Remark 2: Minimizers of the cost (3) are minimizers of the reparametrization-invariant cost

$$\mathcal{L} = \int_a^b \sqrt{u(\tau)^2 + v(\tau)^2} d\tau = \int_a^b \|\dot{\gamma}(\tau)\| \sqrt{1 + \frac{\kappa(\tau)^2}{\beta^2}} d\tau,$$

where $\kappa(\cdot)$ is the curvature of the planar curve γ and β is the dimensional parameter appearing in the vector field Θ .

Remark 3: Observe that here $\theta \in P^1$, i.e., angles are considered modulo π . Notice that the vector field X is not continuous on $PT\mathbb{R}^2$. Indeed, a correct definition of problem (1), (3), (4) needs two charts, as explained in detail in [6, Remark 12]. In this paper, the use of two charts is implicit, since it plays no crucial role.

The minimization problem (1), (3), (4) is a particular case of a minimization problem called a sub-Riemannian problem. For more details on this interpretation see the original papers [12], [25], the book [24], or [6], [5], [1] for a language more consistent with the one of this paper.

One the main interests of the sub-Riemannian problem (1), (3), (4) is the possibility of associating to it a hypoelliptic diffusion equation which can be used to reconstruct images (and not just curves), and for contour enhancement. This point of view was developed in [6], [12], [17], and is the subject of the next section.

III. PURE HYPOELLIPTIC DIFFUSION FOR IMAGE INPAINTING

The sub-Riemannian problem (1), (3), (4) described above was used to reconstruct images whose level sets are smooth curves in [22]. The technique they developed consists of reconstructing as the missing parts of the level sets of the image by applying the above curve-reconstruction algorithm. Obviously, beside the fact that level sets in general need not be smooth nor curves, when applying this method to reconstruct images with large corrupted parts, one is faced with the problem of how to put in correspondence the non-corrupted parts of the same level set.

To avoid this problem, it is natural to model the cortical activation induced by an image $f : \mathbb{R}^2 \rightarrow [0, 1]$ as a function $\mathcal{L}f : PT\mathbb{R}^2 \rightarrow \mathbb{R}$, where $\mathcal{L}f(x, y, \theta)$ represents the strength of activation of the neuron $(x, y, \theta) \in PT\mathbb{R}^2$. It is also known that this cortical activation will evolve with time, due to the presence of the horizontal and vertical cortical connections already introduced in Section II. Then, a natural way to model this evolution is to assume that the activation of a single neuron propagates as a stochastic process Z_t on $PT\mathbb{R}^2$, which solves the stochastic differential equation associated with system (1). That is, letting u_t and v_t be two one-dimensional independent Wiener processes, we have

$$dZ_t = u_t X(Z_t) + v_t \Theta(Z_t). \quad (6)$$

Under this assumption, the evolution of a cortical activation $\mathcal{L}f$ is obtained via the diffusion naturally associated with (6), that is,

$$\begin{cases} \frac{\partial \psi}{\partial t} = \frac{1}{2} \Delta \psi, \\ \psi|_{t=0} = \mathcal{L}f, \end{cases}$$

where,

$$\Delta = X^2 + \Theta^2 = \left(\cos(\theta) \frac{\partial}{\partial x} + \sin(\theta) \frac{\partial}{\partial y} \right)^2 + \beta^2 \frac{\partial^2}{\partial \theta^2}. \quad (7)$$

Observe that the above diffusion is highly anisotropic, since the operator Δ only takes into account two of the three possible directions in $PT\mathbb{R}^2$. Indeed, Δ is not elliptic. However, due to the fact that the family of vector fields $\{X, \Theta\}$ satisfy Hörmander condition, the operator Δ is hypoelliptic. The use of equation (7) to model the spontaneous evolution of cortical activations was first proposed by Citti and Sarti in [12] (although they posed the problem in $SE(2)$, a double covering of $PT\mathbb{R}^2$).

Remark 4: The solution of (7) is strictly related to the solution of the minimization problem (1), (3), (4). Indeed, a result by Léandre [21] shows that, letting p_t be the kernel of (7) and $E(\cdot, \cdot)$ be the value function of problem (1), (4) with cost (5), for all $q, p \in PT\mathbb{R}^2$ it holds $-4t \log(p_t(q, p)) \rightarrow E(q, p)^2$ as $t \rightarrow 0$.

To be able to use equation (7) to reconstruct a corrupted image, one has to specify two things: **i)** the lift operator $\mathcal{L}$, which maps images on $\mathbb{R}^2$ to cortical activations on $PT\mathbb{R}^2$ **ii)** how to project the result of the diffusion on $\mathbb{R}^2$ to get the reconstructed image. We remark that in this pipeline we do not exploit any knowledge on the location and/or shape of the corruption.

In the remainder of this section, we will discuss these two points and present an idea of the efficient numerical scheme we exploit to solve (7).

A. Lifting procedure

Concerning *how* the visual cortex lifts an image to $PT\mathbb{R}^2$, it seems likely that this is done through multiple convolutions with orientation sensitive filters (like Gabor filters), see e.g. [14]. (See also [26].) The main mathematical advantage of this kind of lifts is that, for any $(x, y) \in \mathbb{R}^2$, the vertical components $\mathcal{L}f(x, y, \cdot)$ are essentially 1D gaussians centered around the orientation $\theta(x, y)$ of the gradient $\nabla f(x, y)$.

In this paper, we will consider only a primitive version of this lift. Indeed, we will consider the lift $\mathcal{L}f : PT\mathbb{R}^2 \rightarrow \mathbb{R}$ of a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ to be the distribution $\mathcal{L}f = f \delta_{Sf}$, where δ_{Sf} is the Dirac delta measure concentrated on the surface

$$Sf = \{(x, y, \theta) \mid (x, y) \in \mathbb{R}^2, \theta = \theta(x, y)\} \subset PT\mathbb{R}^2, \\ \theta(x, y) = \begin{cases} \arctan \frac{\partial_y f(x, y)}{\partial_x f(x, y)}, & \text{if } \nabla f(x, y) \neq 0, \\ P^1 & \text{otherwise.} \end{cases}$$

In order to insure that Sf be a well-defined hypersurface of $PT\mathbb{R}^2$, we actually compute it on a smoothed version of f , see [6]. This is in accordance with neurophysiological observations that the retina itself operates a Gaussian smoothing. It is interesting to notice that this smoothing renders Sf particularly regular. More precisely, in [6] it has been shown that the convolution of an $L^2(\mathbb{R}^2)$ function with a Gaussian is generically a Morse function and that the following holds.

Proposition 5 ([6]): If $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a Morse function, then Sf is an embedded $2D$ submanifold of $PT\mathbb{R}^2$.

We remark that this distributional version of the lift can be interpreted as a limit of the wavelet transform lifts, where the variance of the gaussians obtained by the latter for $\mathcal{L}f(x, y, \cdot)$ goes to zero.

B. Projection

For our choice of lift, as well as for most of the wavelet transform lifts, solutions to the evolution equation (7) for any time $t > 0$ do not belong to the range of $\mathcal{L}$. Thus, it is not possible to exploit the fact that the lifting procedure is injective in order to define a projection. The most natural way to project functions $\psi : PT\mathbb{R}^2 \rightarrow \mathbb{R}^2$ to $\pi\psi : \mathbb{R}^2 \rightarrow \mathbb{R}$ is then to simply choose $\pi\psi(x, y) = \int_{P^1} \psi(x, y, \theta) d\theta$.

C. Numerical integration

Observe that, since the double covering of $PT\mathbb{R}^2$ is the group of roto-translations $SE(2) = \mathbb{R}^2 \rtimes S^1$, and (7) is covariant w.r.t. the canonical projection $SE(2) \rightarrow PT\mathbb{R}^2$, it is more practical to solve (7) on the latter. Indeed, this allows to exploit the group structure of $SE(2)$. As already remarked, the numerical integration of (7) is subtle, since multiscale sub-Riemannian effects are hidden inside and has been approached in different way. For example, in [12] the authors use a finite difference discretisation of all derivatives while in [16] an almost explicit expression for the heat kernel is exploited.

In [6], [4], [3] we presented a sophisticated and highly parallelizable numerical scheme, based on the non-commutative Fourier transform on a suitable semidiscretization of the group $SE(2)$, i.e., the semidiscrete group of roto-translations $SE(2, N)$ for $N \in \mathbb{N}$. This is the semi-direct product $SE(2, N) = \mathbb{R}^2 \rtimes \mathbb{Z}_N$, where $\mathbb{Z}_N$ is the cyclic group of order N and the action of $n \in \mathbb{Z}_N$ on $\mathbb{R}^2$ is given by the rotation R_n of angle $2\pi n/N$. As pointed out in [4], considering a discrete number of orientations seems to be in accordance with experimental evidence. Although still an open problem, this is probably connected with topological restrictions given by the fact that neurons in V1 encode a 3D space while V1 itself, physically, is essentially 2D.

A way of interpreting the algorithm presented in [4], see also [7], [26], is the following: For a given finite subset $\mathcal{K}$ of $\mathbb{R}^2$, let $SE(2, N, \mathcal{K})$ be the set of $\mathbb{C}^N$ -valued trigonometric polynomials Q on $SE(2, N)$ that read $Q = (Q^n)_{n \in \mathbb{Z}_N}$, where,

$$Q^n(x, y) = \sum_{(\lambda_k, \mu_l) \in \mathcal{K}} c_{k,l}^n e^{i(\lambda_k x + \mu_l y)}, \quad r = 0, \dots, N-1. \quad (8)$$

Here, $c_{k,l}^r \in \mathbb{C}$. Then, if we assume that (standard) Fourier transforms of images are compactly supported on $\mathcal{K}$, their lifts will belong to $SE(2, N, \mathcal{K})$. Equation (7) can be naturally (semi)discretized on $SE(2, N)$, essentially replacing the operator Θ^2 in Δ with its discretized version $\Lambda_N \in \mathbb{R}^N \times \mathbb{R}^N$, and then restricted to $SE(2, N, \mathcal{K})$. This yields the completely uncoupled systems of linear ordinary differential equations

$$\frac{dc_{k,l}}{dt} = -2\pi^2 \text{diag}(\lambda_k \cos \theta_r + \mu_l \sin \theta_r)^2 c_{k,l} + \Lambda_N c_{k,l}. \quad (9)$$

Here, $c_{k,l}(t) = (c_{k,l}^0(t), \dots, c_{k,l}^{N-1}(t))^T \in \mathbb{C}^N$. These systems are equipped with initial conditions $c_{k,l}(0) = c_{k,l}$, where the latter are the coefficients in (8) corresponding to $\mathcal{L}f$.

These discretized equations can then be solved through any numerical scheme. We chose the Crank-Nicolson method, for its good convergence and stability properties. Let us remark that the operators appearing on the r.h.s. of (9) are periodic tridiagonal matrices, i.e. tridiagonal matrices with non-zero $(1, N)$ and $(N, 1)$ elements. Thus, the linear system appearing at each step of the Crank-Nicolson method can be solved through the Thomas algorithm for periodic tridiagonal matrices, of computational cost $\mathcal{O}(N)$.

D. Numerical experiments

When implementing the pure hypoelliptic diffusion, there are essentially 3 parameters to be tuned: the number of angles N for the semi-discretization of the equation on $SE(2, N)$, the weight parameter β appearing in the operator Δ , and the total time of diffusion T . Clearly, for performance reasons, one would like to have N as small as possible. Indeed, numerical experiments suggests that it suffices to choose $N = 30$, as increasing N beyond this threshold does not affect the resulting image in any visible way. On the other hand, the parameters β and T have to be tuned by hand and the optimal choice seems to be deeply sensitive to the image to be treated (and of course to the desired result). In Figure 2, we present a sequence of images showing the effect of the pure hypoelliptic diffusion at different times.

Fig. 2: Hypoelliptic diffusion for increasing times $(\frac{1}{8}, \frac{1}{4}, \frac{1}{2}, 1)$, from left to right, and $\beta^2 = \frac{1}{4}$.

IV. HEURISTIC COMPLEMENTS: EXPLOITING INFORMATION ON THE LOCATION OF THE CORRUPTION

The method explained in the previous section does not use any information of where the image is corrupted. In this section we show how the suitable use of this information permits to obtain better reconstructions. In particular we will present two extensions to the previous algorithm, the Dynamic Restoration (DR) procedure [4] and the varying coefficients hypoelliptic diffusion, and briefly introduce a sort of synthesis of the two, the Averaging and Hypoelliptic Evolution (AHE) algorithm [3]. The results obtained by these methods are comparable with the current state-of-the-art for PDE based image inpainting algorithms, see, e.g., [9]. Actually, in our opinion, in order to go beyond this state-of-the-art it is necessary to introduce additional external information on the corrupted part as it is done, for instance, in exemplar-based methods [9].

A. Dynamic Restoration

In this section we present a technique to exploit the information on the location of the corruption in the inpainting algorithm. Assume that a partition of the set of pixels of the image is given $I = G \cup B$, where points in G are “good”, i.e., non-corrupted, while those in B are “bad”, i.e., corrupted. The idea is now to periodically “mix” the solution ψ_t of the diffusion on $SE(2, N)$ with the initial function $\mathcal{L}f$ on G , while keeping tabs on the “evolution” of the set of good points.

Namely, fix $n \in \mathbb{N}$ and split the segment $[0, T]$ into n intervals $t_r = r\tau$, $r = 0, \dots, n$, $\tau = T/n$. Let $G(0) = G$, $B(0) = B$ and iteratively solve the hypoelliptic diffusion equation on each $[t_r, t_{r+1}]$ with initial condition

$$\psi_{t_r}(k, x, y) = \begin{cases} \psi_{t_r}^-(k, x, y) & \text{if } (x, y) \in B(r) \\ \sigma(x, y, t_k) \psi_{t_r}^-(k, x, y) & \text{if } (x, y) \in G(r). \end{cases}$$

Here, the function ψ^- is the solution of the diffusion on the previous interval (or the starting lifted function if $r = 0$), and the coefficient σ is given by

$$\sigma(x, y, t_r) = \frac{1}{2} \frac{h(x, y, 0) + h(x, y, t_r)}{h(x, y, t_r)},$$

$$h(x, y, t) = \max_k \psi_t(k, x, y).$$

Moreover, after each step, $G(r+1)$ and $B(r+1)$ are obtained from $G(r)$ and $B(r)$ as follows:

1. Project the solution $\psi_{t_{r+1}}$ to the image f_{r+1} .
2. Define $\text{avg } f_{r+1}(x, y)$ as the average of f_{r+1} on the 9-pixels neighborhood of (x, y) .
3. Define the set $W = \{(x, y) \in \partial B(r) \mid f_{r+1}(x, y) \geq \text{avg } f_{r+1}(x, y)\}$.
4. Let $G(r+1) = B(r) \cup W$, $B(r+1) = B(r) \setminus W$.

Some reconstruction results via the Dynamic Restoration procedure are presented in 3. Notice that restorations obtained via this method are much better than those obtained via pure hypoelliptic

Fig. 3: Two inpainting via the DR method of images missing, respectively, 30% and 80% of pixels. For both images we chose $T = 1$ and $n = 60$ steps. Notice how in the second image, the big corruption makes it difficult to correctly compute the gradient.

diffusion. In particular, the DR procedure gives reasonable results even on highly corrupted images, as shown in Figure 3.

B. Varying coefficients hypoelliptic diffusion

A different approach w.r.t. the Dynamic Restoration procedure is to directly modify the evolution equation (7), in order to enforce a stronger diffusion on the parts of the image that are known to be corrupted. Namely, we can replace the operator Δ in (7) by $\Delta_{\mathcal{H}} := bX^2 + \frac{a}{\beta^2}\Theta^2$ for some non-negative continuous coefficients $a, b : \mathbb{R}^2 \rightarrow \mathbb{R}$, i.e.,

$$\Delta_{\mathcal{H}} = b(x, y) \left(\cos(\theta) \frac{\partial}{\partial x} + \sin(\theta) \frac{\partial}{\partial y} \right)^2 + a(x, y) \frac{\partial^2}{\partial \theta^2}. \quad (10)$$

The above equation with varying coefficients a, b tries to implement the natural idea of reducing the effect of diffusion at non-corrupted points. Indeed, when using (7), this can be done only by decreasing of the total time of diffusion T , but this acts in an homogeneous way on all points, including the corrupted ones. On the other hand, when using equation (10) one can tune the coefficients a, b depending on the point (x, y) , and consequently, weaken the diffusion effect only where it is necessary.

When choosing the varying coefficients a and b , the idea is to make them larger at bad points and their neighbors (especially the

Fig. 4: Three reconstructions via the the varying coefficients restoration procedure: with small diagonal corruptions, and with 30% and 80% of pixels randomly removed.

coefficient b , which has the most influence to the velocity of the diffusion). Thus, they will be chosen to be a smooth approximation of the indicator function of the the set B of bad (corrupted) points and c_0, c_1 are positive constants. In particular, our choice is:

$$a(x, y) = \begin{cases} a_0 + a_1 \epsilon(x, y) & \text{if } \frac{a_0 + a_1 \epsilon(x, y)}{a_0 + a_1} > \epsilon_*, \\ 0 & \text{otherwise,} \end{cases}$$

$$b(x, y) = \begin{cases} b_0 + b_1 \epsilon(x, y) & \text{if } \frac{b_0 + b_1 \epsilon(x, y)}{b_0 + b_1} > \epsilon_*, \\ 0 & \text{otherwise,} \end{cases}$$

where $\epsilon(x, y) = \exp(-f^2(x, y)/\sigma)$. Here, $a_0, b_0 \in \mathbb{R}$, $a_1, b_1, \sigma > 0$, and $\epsilon_* \in (0, 1)$ are constant parameters chosen experimentally.

Some numerical experiments are presented in Figure 4. We observe that for low levels of corruption, reconstruction of images with this method yields results which are comparable to, if not better than, the ones obtained through the DR procedure. However, when the corrupted part becomes larger this method fails. This suggests that, in order to obtain a good inpainting algorithm for highly corrupted images, one has still to use the DR procedure, combining it with the varying coefficients. This is an essential component of the algorithm presented in the next section.

C. AHE algorithm

In order to improve on the results for high corruption rates, in [3] we proposed the Averaging and Hypoelliptic Evolution algorithm. The main idea behind the AHE algorithm is to try to provide the anisotropic diffusion with better initial conditions. More precisely, it is divided in the following 4 steps (see Figure 5):

- 1) **Preprocessing phase (Simple averaging).** We apply a simple iterative procedure that fills in the corrupted area by

assigning to each corrupted pixel the average value of the non-corrupted neighboring pixels.

- 2) **Main diffusion (Strong smoothing).** By using the result of the previous procedure as an input, we apply the varying coefficients hypoelliptic diffusion discussed in Section IV-B.
- 3) **Advanced averaging.** In order to remove the blur introduced by the hypoelliptic diffusion of the previous step, we mix the results of step 1 and step 2.
- 4) **Weak smoothing.** We perform a last hypoelliptic evolution, in order to smooth some of the edges we obtained in step 3.

In Figure 6 we present the results obtained via the AHE algorithm on highly corrupted images. In particular, a comparison with Figure 3 shows that the synthesis between the DR procedure and the varying coefficients, coupled with the averaging steps, yields much better results w.r.t. the simple application of the DR procedure.

V. CONCLUSIONS

In this paper we presented several image inpainting algorithms based on hypoelliptic diffusion. Although pure hypoelliptic diffusion allows to obtain reasonable inpaints of images with small corrupted regions, we showed that coupling such diffusion with heuristic methods exploiting the full knowledge of the location of the corruption allows to obtain very efficient reconstructions. It is interesting to notice that when the image is so corrupted that our visual system is not able to recognize it, as it happens, for instance, for the corrupted image presented in Figure 7, the use of pure hypoelliptic diffusion does not help it. This fact can be interpreted as a validation of the CPS model with pure hypoelliptic diffusion. On the contrary, as already pointed out, both the DR procedure and the AHE algorithm produce reconstructions that go beyond the capabilities of our visual system.

ACKNOWLEDGMENTS

This work was supported the ERC POC project ARTIV1, project number 727283, by the ANR project SRGI, contract number ANR-15-CE40-0018, and by a public grant as part of the ‘‘Investissement d’avenir project’’, reference ANR-11-LABX-0056-LMH, LabEx LMH, in a joint call with the FMJH Program Gaspard Monge in optimization and operation research.

REFERENCES

- [1] A. Agrachev, D. Barilari, and U. Boscain. Introduction to Riemannian and sub-Riemannian geometry (Lecture Notes), http://people.sissa.it/agrachev/agrachev_files/notes.html. 2017.
- [2] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In *Proc. of SIGGRAPH 2000, New Orleans, USA, July*, pages 417–424, 2000.
- [3] U. Boscain, R. Chertovskih, J.-P. Gauthier, D. Prandi, and A. Remizov. Highly corrupted image inpainting through hypoelliptic diffusion. *ArXiv e-prints*, Feb. 2015.
- [4] U. Boscain, R. A. Chertovskih, J. P. Gauthier, and A. O. Remizov. Hypoelliptic diffusion and human vision: a semidiscrete new twist. *SIAM J. Imaging Sci.*, 7(2):669–695, 2014.
- [5] U. Boscain, R. Duits, F. Rossi, and Y. Sachkov. Curve cusplless reconstruction via sub-Riemannian geometry. *ESAIM Control Optim. Calc. Var.*, 20(3):748–770, 2014.
- [6] U. Boscain, J. Duplaix, J.-P. Gauthier, and F. Rossi. Anthropomorphic image reconstruction via hypoelliptic diffusion. *SIAM J. Control Optim.*, 50(3):1–25, 2012.
- [7] U. Boscain, J.-P. Gauthier, D. Prandi, and A. Remizov. Image reconstruction via non-isotropic diffusion in Dubins/Reed-Shepp-like control systems. In *53rd IEEE Conference on Decision and Control*, pages 4278–4283, 2014.
- [8] K. Bredies, T. Pock, and B. Wirth. Convex Relaxation of a Class of Vertex Penalizing Functionals. *J. Math. Imaging Vis.*, 47(3):278–302, 2013.
- [9] F. Cao, Y. Gousseau, S. Masnou, and P. Pérez. Geometrically guided exemplar-based inpainting. *SIAM Journal on Imaging Sciences*, 4(4):1143–1179, 2011.

Fig. 5: A depiction of the different steps of the AHE procedure. Observe the sharpening from step 2 to step 4.

Fig. 6: Image inpainting via the AHE algorithm of highly corrupted images with, respectively, 80% and 90% of pixels missing.

Fig. 7: A comparison of the different algorithms presented in this paper. *From left to right*: the original image, the corrupted image used as input, the reconstructions obtained via pure hypoelliptic diffusion ($\beta^2 = \frac{1}{4}$, $T = 1$), the DR procedure ($n = 120$), and the AHE algorithm.

- [10] T. Chan, S. Kang, and J. Shen. Euler's elastica and curvature-based inpainting. *SIAM J. Appl. Math.*, 63(2):564–592, 2002.
- [11] G. Citti, B. Franceschiello, G. Sanguinetti, and A. Sarti. Sub-riemannian mean curvature flow for image processing. *SIAM J. Imaging Sci.*, 9(1):212–237, 2016.
- [12] G. Citti and A. Sarti. A cortical based model of perceptual completion in the roto-translation space. *J. Math. Imaging Vis.*, 24(3):307–326, 2006.
- [13] A. Criminisi, P. Pérez, and K. Toyama. Region filling and object removal by exemplar-based image inpainting. *IEEE Trans. Image Process.*, 13(9), 2004.
- [14] J. G. Daugman. Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *J. Opt. Soc. Am. A*, 2(7):1160, jul 1985.
- [15] R. Duits, U. Boscaín, F. Rossi, and Y. Sachkov. Association fields via cusplless sub-Riemannian geodesics in SE(2). *J. Math. Imaging Vision*, 49(2):384–417, 2014.
- [16] R. Duits and E. Franken. Left-invariant parabolic evolutions on SE(2) and contour enhancement via invertible orientation scores Part I: linear left-invariant diffusion equations on SE(2). *Quart. Appl. Math.*, 68(2):255–292, 2010.
- [17] R. Duits and E. Franken. Left-invariant parabolic evolutions on SE(2) and contour enhancement via invertible orientation scores Part II: nonlinear left-invariant diffusions on invertible orientation scores. *Quart. Appl. Math.*, 68(2):293–331, 2010.
- [18] R. K. Hladky and S. D. Pauls. Minimal surfaces in the roto-translation group with applications to a neuro-biological image completion model. *J. Math. Imaging Vision*, 36(1):1–27, 2010.
- [19] W. C. Hoffman. The visual cortex is a contact bundle. *Applied Mathematics and Computation*, 32(2-3):137–167, 1989.
- [20] D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, 160(1):106–154, 1962.
- [21] R. Léandre. Majoration en temps petit de la densité d'une diffusion dégénérée. *Probab. Theory Related Fields*, 74(2):289–294, 1987.
- [22] A. P. Mashtakov, A. A. Ardentov, and Y. L. Sachkov. Parallel algorithm and software for image inpainting via sub-riemannian minimizers on the group of rototranslations. *Numerical Mathematics: Theory, Methods and Applications*, 6(1):95–115, 2013.
- [23] S. Masnou. Disocclusion: a variational approach using level lines. *IEEE Trans. Image Process.*, 11(2):68–76, 2002.
- [24] J. Petitot. *Neurogéométrie de la vision: modeles mathématiques et physiques des architectures fonctionnelles*. Editions Ecole Polytechnique, 2008.
- [25] J. Petitot and Y. Tondut. Vers une neurogéométrie. fibrations corticales, structures de contact et contours subjectifs modaux. *Mathématiques informatique et sciences humaines*, 145:5–102, 1999.
- [26] D. Prandi and J.-P. Gauthier. A semidiscrete version of the Petitot model as a plausible model for anthropomorphic image reconstruction and pattern recognition. *ArXiv e-prints*, Apr. 2017.