

CORTICAL-INSPIRED IMAGE RECONSTRUCTION VIA SUB-RIEMANNIAN GEOMETRY AND HYPOELLIPTIC DIFFUSION

UGO BOSCAIN^{1,2}, ROMAN CHERTOVSKIH^{3,4}, JEAN-PAUL GAUTHIER⁵, DARIO PRANDI⁶
AND ALEXEY REMIZOV⁷

Abstract. In this paper we review 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.

1. INTRODUCTION

Since the pioneering works of Petitot, Citti, and Sarti [15, 34], 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 developing and applying these ideas to image processing and computer vision [10, 19, 20, 24, 36]. 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 such as local orientation, presents a somewhat new framework in which many tasks can be strongly simplified. More precisely, a neuron in the primary visual cortex is represented by a point $(x, y, \theta) \in \mathbb{R}^2 \times \mathbb{P}^1$, where $(x, y) \in \mathbb{R}^2$ is its spatial preference, and $\theta \in \mathbb{P}^1$ is its direction preference. Then, an image on the visual plane, given by $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, is represented in the primary visual cortex by a distribution $\mathcal{L}f(x, y, \theta) = f(x, y)\delta_{Sf}(x, y, \theta)$, where δ_{Sf} is the Dirac mass of the surface $Sf \subset \mathbb{R}^2 \times \mathbb{P}^1$. Here, Sf is composed of the points (x, y, θ) such that $\theta \in \mathbb{P}^1$ is the direction of the level line of f passing through $(x, y) \in \mathbb{R}^2$. The structure of neuronal connections in V1 then implies that the initial cortical stimulus $\mathcal{L}f$ evolves according to a heat-type equation $\partial_t \Psi = \Delta \Psi$, $\Psi|_{t=0} = \mathcal{L}f$, where Δ is not the standard Euclidean Laplacian, but an hypoelliptic operator encoding the horizontal and vertical neuronal connections. When the result of such 3D diffusion is projected back to the visual plane, we obtain a sort of anisotropic 2D diffusion that is very strong in the direction of the level lines of the image and very weak in the orthogonal one.

¹ CNRS, LJLL, Université Pierre et Marie Curie, Paris, France

² Équipe INRIA Paris CAGE

³ SYSTEC, FEUP, University of Porto, Portugal

⁴ Samara University, Russia

⁵ LSIS, Université de Toulon, France

⁶ CRNS, L2S, CentraleSupélec, Gif-sur-Yvette, France

⁷ CMAP, École Polytechnique, Palaiseau, France

The combination of the redundant representation and hypoelliptic diffusion is particularly efficient for image processing tasks as inpainting. Nowadays this is a challenging problem treated with many different techniques, see, e.g., [3, 4, 13, 16, 30] and references therein. In this survey, we present the results obtained in a series of works [5–7, 9, 35] in this framework.

It should be mentioned that efficient PDE-methods for image processing based on 2D anisotropic diffusions have been studied in various contexts. In their simplest form, these methods use a 2D heat-like equation whose coefficients depend on the local structure of the image. See, e.g., [23, 37]. We remark that the 2nd order Laplace-like operator appearing in these methods is anisotropic, but still elliptic. There exists also different approaches using non-linear evolutions or higher order PDEs. The method presented in this paper is deeply different, since the use of the cortical-inspired redundant representation allows to work with a 2nd order linear equation on a 3D space, and permits to consider an “infinitely” anisotropic (i.e., positive but not elliptic) equation. However, this equation is still hypoelliptic and can be efficiently solved via non-commutative Fourier analysis.

The paper is divided in three parts, corresponding to different methods and purposes of image reconstruction.

In the first part, Section 2, we introduce the geometry of the problem by describing the Citti-Petitot-Sarti model for reconstruction of curves, with some new ingredients introduced in [7]. 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 3, 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 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 neither needing nor exploiting any information on the location and shape of the corruption.

Finally, in the third and last part, Section 4, 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 contrary, the methods presented in the last part completely transcend such problem, thus suggesting that, although based on neurophysiological ideas, they do not reflect real mechanisms of V1.

While the results obtained with our methods, are comparable with the state of the art in image inpainting, here our purpose is not to provide a comparison with the existing methods, but to present how ideas from neurophysiology can be exploited for image processing tasks. Nevertheless, a brief study of the quality of the reconstructions is presented in Section 5, together with a comparison with some similar methods, as those presented in [12, 23]. An interesting line of investigation would be to use our framework in combination with other methods based on totally different ideas as for instance exemplar-based methods, e.g., [10, 22].

2. 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 [33, 34]. It is based on previous work by Hubel-Wiesel [26] and Hoffman [25], then it was refined by Citti *et al.* [14, 15], Duits *et al.* [18–21]. and by the authors of the present paper in [5–7, 9]. It was also studied by Hladky and Pauls in [24].

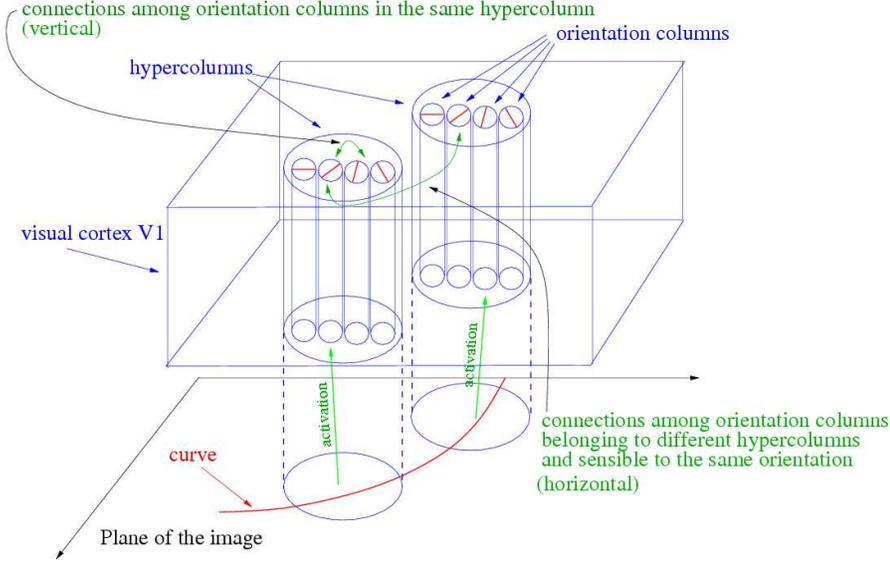


FIGURE 1. A scheme of the primary visual cortex V1.

In a simplified model¹ (see [33, 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 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 $PT\mathbb{R}^2$:

$$X(x, y, \theta) = \begin{pmatrix} \cos \theta \\ \sin \theta \\ 0 \end{pmatrix} \quad \text{and} \quad \Theta(x, y, \theta) = \begin{pmatrix} 0 \\ 0 \\ \beta \end{pmatrix}. \quad (1)$$

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:

$$\frac{d}{dt}\Gamma = uX(\Gamma) + v\Theta(\Gamma), \quad \text{for some } u, v : [0, T] \rightarrow \mathbb{R}. \quad (2)$$

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”.

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 $\Gamma \in W^{1,1}([0, T])$.

Remark 2.1. Observe that the lift of a planar curve $\gamma : [0, T] \rightarrow \mathbb{R}^2$ is not unique in general: for example, if $\dot{\gamma}|_{(\tau_1, \tau_2)} \equiv 0$ for $0 \leq \tau_1 < \tau_2 \leq T$, it is clear that (2) is satisfied on (τ_1, τ_2) as soon as $u|_{(\tau_1, \tau_2)} \equiv 0$, independently of v . Nevertheless, the lift is unique in many relevant cases, e.g., whenever $\dot{\gamma}(t) = 0$ only on a discrete subset of $[0, T]$.

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 (2) holds for some $u, v \in L^1([0, T])$.

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

$$\theta(t) = \arctan \frac{\dot{y}(t)}{\dot{x}(t)}, \quad \text{for a. e. } t \in [0, T]. \quad (3)$$

2.1. 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 (3) 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 [32] 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,

$$J = \int_a^b (u(\tau)^2 + v(\tau)^2) d\tau \rightarrow \min. \quad (4)$$

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

$$\tilde{\gamma}(a) = \gamma_{\text{in}}, \quad \tilde{\gamma}(b) = \gamma_{\text{fin}}, \quad \tilde{\theta}(a) = \theta_{\text{in}}, \quad \tilde{\theta}(b) = \theta_{\text{fin}}. \quad (5)$$

Here $\tilde{\theta}(\cdot)$ is the lift associated with $\tilde{\gamma}(\cdot)$ via (3). Observe that in the cost (4), $u(\tau)^2$ (resp. $v(\tau)^2$) represents the (infinitesimal) energy necessary to activate horizontal (resp. vertical) connections. As pointed out in [7], a solution to the problem (2), (4), (5) always exists.

An example of curve reconstruction obtained via this technique is given in Figure 2.

Remark 2. Minimizers of the cost (4) are minimizers of the reparametrization-invariant cost

$$\mathcal{L} = \int_a^b \sqrt{u(\tau)^2 + v(\tau)^2} d\tau = \int_a^b \|\dot{\gamma}(\tau)\| \sqrt{1 + \frac{\kappa(\tau)^2}{\beta^2}} d\tau, \quad (6)$$

where $\kappa(\cdot)$ is the curvature of the planar curve γ and β is the dimensional parameter appearing in the vector field Θ . For more details about the relations between minimizers of costs (4) and (6), see [31].

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 (2), (4), (5) needs two charts, as explained in detail in [7, Remark 12]. In this paper, the use of two charts is implicit, since it plays no crucial role.

Remark 4. From the theoretical point of view, the weight parameter β is irrelevant: for any $\beta > 0$ there exists a homothety of the (x, y) -plane that maps minimizers of problem (2), (4), (5) to those of the same problem with $\beta = 1$. However its role will be important in our inpainting algorithms (see Section 7).

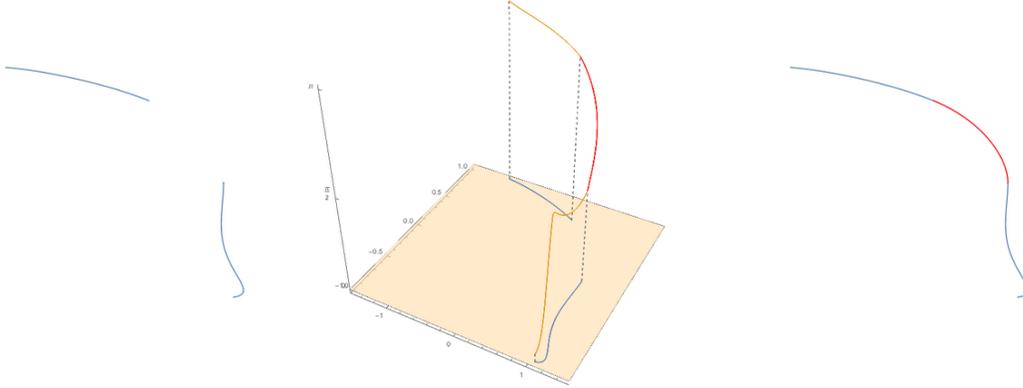


FIGURE 2. Curve reconstruction. *From left to right:* The interrupted curve. The lift of the interrupted curve, in yellow, and the solution to the minimization problem, in red. The final reconstructed curve.

The minimization problem (2), (4), (5) is a particular case of a minimization problem called a sub-Riemannian problem. For more details on this interpretation see the original papers [15, 32, 34], the book [33], or [2, 6, 7] for a language more consistent with the one of this paper.

One of the main interests of the sub-Riemannian problem (2), (4), (5) 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 [7, 15, 19–21], and is the subject of the next section.

3. PURE HYPOELLIPTIC DIFFUSION FOR IMAGE INPAINTING

The sub-Riemannian problem (2), (4), (5) described above was used to reconstruct images whose level sets are smooth curves by Ardentov, Mashtakov and Sachkov [29]. The technique they developed consists of reconstructing 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 2. 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 (2). That is, letting u_t and v_t be two one-dimensional independent Wiener processes, we have

$$dZ_t = u_t X(Z_t) + v_t \Theta(Z_t). \quad (7)$$

Under this assumption, the evolution of a cortical activation $\mathcal{L}f$ is obtained via the diffusion naturally associated with (7), that is,

$$\begin{cases} \frac{\partial \psi}{\partial t} = \frac{1}{2} \Delta \psi, \\ \psi|_{t=0} = \mathcal{L}f, \end{cases} \quad \text{where} \quad \Delta = X^2 + \Theta^2 = \left(\cos(\theta) \frac{\partial}{\partial x} + \sin(\theta) \frac{\partial}{\partial y} \right)^2 + \beta^2 \frac{\partial^2}{\partial \theta^2}. \quad (8)$$

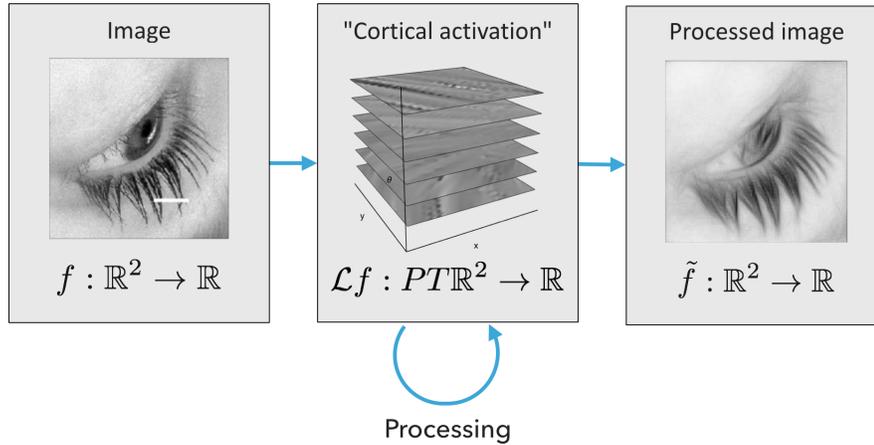


FIGURE 3. A schematic of the pipeline for image reconstruction.

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 (8) to model the spontaneous evolution of cortical activations was first proposed by Citti and Sarti in [15] (although they posed the problem in $SE(2)$, a double covering of $PT\mathbb{R}^2$).

Remark 5. *We have the following relationships between the sub-Riemannian problem control problem (2), (4), (5), and the diffusion equation (8):*

- *By the Feynman – Kac formula, integrating Equation (8), one expects to reconstruct the most probable missing level curves (among admissible ones).*
- *The solution of (8) is strictly related to the solution of the minimization problem (2), (4), (5). Indeed, a result by Léandre [27, 28] shows that, letting p_t be the kernel of (8) and $E(\cdot, \cdot)$ be the value function of problem (2), (5) with cost (6), for all $q, p \in PT\mathbb{R}^2$ it holds*

$$-4t \log(p_t(q, p)) \longrightarrow E(q, p)^2 \quad \text{as } t \rightarrow 0. \quad (9)$$

To be able to use equation (8) 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. (See Figure 3.) 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 (8).

3.1. 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. [17]. This idea is strictly connected with wavelet transforms and orientation scores, which have been implemented, e.g., in [20, 21]. See also [35]. 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. (See Figure 4.) 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

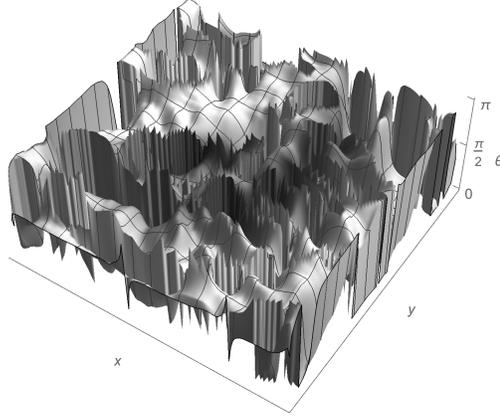


FIGURE 4. Distributional lift of the image on the l.h.s. of Figure 3.

delta measure concentrated on the surface

$$Sf = \{(x, y, \theta) \mid (x, y) \in \mathbb{R}^2, \theta = \theta(x, y)\} \subset PT\mathbb{R}^2,$$

$$\theta(x, y) = \begin{cases} \arctan \frac{\partial_y f(x, y)}{\partial_x f(x, y)} & \text{if } \nabla f(x, y) \neq 0, \\ P^1 & \text{otherwise.} \end{cases} \quad (10)$$

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 [7]. 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 [7] 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 6 ([7]). *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. Variations on the above lifting procedure have been proposed in [10, 15, 36].

3.2. Projection

For our choice of lift, as well as for most of the wavelet transform lifts, solutions to the evolution equation (8) 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

$$\pi\psi(x, y) = \int_{P^1} \psi(x, y, \theta) d\theta. \quad (11)$$

3.3. 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 (8) is covariant w.r.t. the canonical projection $SE(2) \rightarrow PT\mathbb{R}^2$, it is more practical to solve (8) on the latter. Indeed, this allows to exploit the group structure of $SE(2)$. As already remarked, the numerical integration of

(8) is subtle, since multiscale sub-Riemannian effects are hidden inside³ and has been approached in different way. For example, in [15, 36] the authors use a finite difference discretisation of all derivatives while in [19–21] an almost explicit expression for the heat kernel is exploited.

In [5, 7, 9] 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 [5], 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 [5], see also [8, 35], 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

$$Q = (Q^n)_{n \in \mathbb{Z}_N}, \quad \text{where} \quad Q^n(x, y) = \sum_{(\lambda_k, \mu_l) \in \mathcal{K}} c_{k,l}^n e^{i(\lambda_k x + \mu_l y)}, \quad r = 0, \dots, N-1. \quad (12)$$

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 (8) 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

$$\frac{dc_{k,l}}{dt} = -2\pi^2 \text{diag}(\lambda_k \cos \theta_r + \mu_l \sin \theta_r)^2 c_{k,l} + \Lambda_N c_{k,l}, \quad c_{k,l} \in \mathbb{C}^N. \quad (13)$$

Here, $c_{k,l}(t) = (c_{k,l}^0(t), \dots, c_{k,l}^{N-1}(t))^T$. These systems are equipped with initial conditions $c_{k,l}(0) = c_{k,l}$, where the latter are the coefficients in (12) 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 (13) 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 $\mathcal{O}(N)$.

Remark 7. *By [5, Theorem 2], solving (13) for some couple (λ_k, μ_l) is equivalent to solve it for any rotated couple $R_n(\lambda_k, \mu_l)$, associated with $n \in \mathbb{Z}_N$. Thus, if the set \mathcal{K} is invariant with respect to rotations R_n , $n \in \mathbb{Z}_N$, it is indeed sufficient to solve (13) for a slice of \mathcal{K} whose orbit under the rotations covers the whole \mathcal{K} .*

3.4. 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 and type of corruption to be treated (and of course to the desired result).

In Figure 5, we present a sequence of images showing the effect of the pure hypoelliptic diffusion with two different choices for β at different times. In particular, this shows how for large values of β , the effect of the pure hypoelliptic diffusion on $PT\mathbb{R}^2$ becomes, after the projection, close to an isotropic diffusion on \mathbb{R}^2 . In Figure 6, we present the results of the pure hypoelliptic diffusion when the parameter β is chosen to be 0. Notice that,

³For instance, one cannot use the nice method presented in [1] for the Heisenberg group. Indeed, in $SE(2)$ one cannot build a refinable grid having a subgroup properties since such grid has at most 6 angles and hence cannot be refined, see e.g., [11].



FIGURE 5. Hypoelliptic diffusion for increasing times $(\frac{1}{8}, \frac{1}{4}, \frac{1}{2}, 1)$, from left to right. *Upper left:* original image. *Upper right:* $\beta^2 = \frac{1}{4}$. *Bottom:* $\beta^2 = 4$.

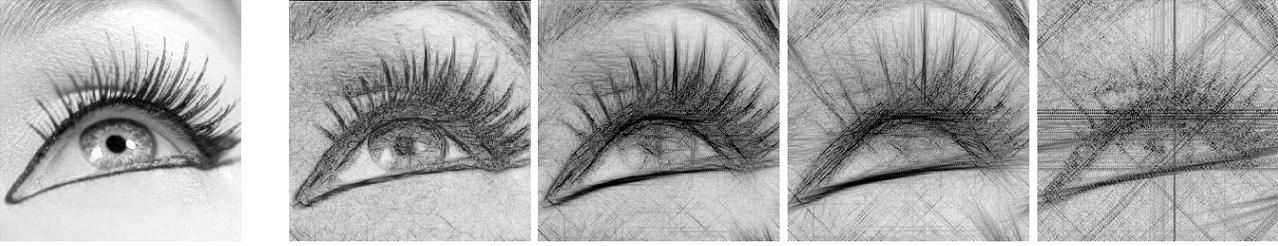


FIGURE 6. Hypoelliptic diffusion of the image on the left with $\beta = 0$ and increasing times $(\frac{1}{8}, \frac{1}{4}, \frac{1}{2}, 1)$. Notice that, due to the absence of diffusion in the angular part, we observe the formation of straight lines tangent to the level lines.

due to the absence of diffusion in the angular part, we observe the formation of straight lines tangent to the level lines.

In order to show the strong anisotropy of the hypoelliptic diffusion, we present also Figures 7, 6 8. In the first one, we show the effects of the evolution when the surface Sf on which the image is usually lifted is replaced with $S_{\theta_0} = \{(x, y, \theta_0) \mid (x, y) \in \mathbb{R}^2\}$, for different choices of $\theta_0 \in P^1$. In accordance with the mathematical formulation, this yields a diffusion restricted to the direction θ_0 . Finally, in Figure 8, we show the hypoelliptic evolution when the lift is replaced with $\tilde{L}_f g = \delta_{Sf} g$ for two different images f, g . Namely, we use the values of the first image only to compute Sf and on this surface we consider the values given by the second image. Here, although the information on the values of the first image is lost, the information on the directions of the level lines alone allows to recover some of the contours.

4. 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 [5] and the varying coefficients hypoelliptic diffusion, and briefly introduce a sort of synthesis of the two, the Averaging and Hypoelliptic Evolution (AHE) algorithm [9]. The results obtained by these methods are



FIGURE 7. Hypocoelliptic diffusion of lift with fixed angle θ_0 . *Left*: original image. *Middle*: $\theta_0 = \frac{\pi}{4}$. *Right*: $\theta_0 = \frac{3\pi}{4}$. In particular, observe that only the white stripe perpendicular to θ_0 is filled, while the other is preserved.



FIGURE 8. Hypocoelliptic diffusion with lift $\tilde{L}_{fg} = \delta_S f g$, where f and g are, respectively, the first and second images. Observe that we recover the contours of the first image only through the surface Sf . Here we chose $\beta^2 = 0.1$ and $T = 1$.

comparable with the current state-of-the-art for PDE based image inpainting algorithms [12, 22]. 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 [12].

4.1. 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 hypocoelliptic diffusion equation on each $[t_r, t_{r+1}]$ with initial condition

$$\psi_{t_r}(k, x, y) = \begin{cases} \psi_{t_r}^-(k, x, y) & \text{if } (x, y) \in B(r) \\ \sigma(x, y, t_k) \psi_{t_r}^-(k, x, y) & \text{if } (x, y) \in G(r). \end{cases} \quad (14)$$

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

$$\sigma(x, y, t_r) = \frac{1}{2} \frac{h(x, y, 0) + h(x, y, t_r)}{h(x, y, t_r)}, \quad h(x, y, t) = \max_k \psi_t(k, x, y). \quad (15)$$

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

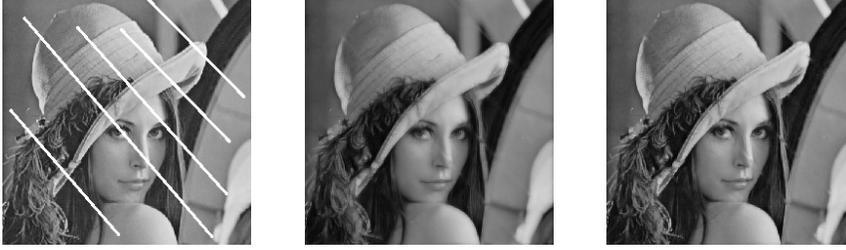


FIGURE 9. Inpaintings via hypoelliptic diffusion with the DR procedure. *Left*: original image. *Middle*: $n = 30$ steps. *Right*: $n = 120$ steps.

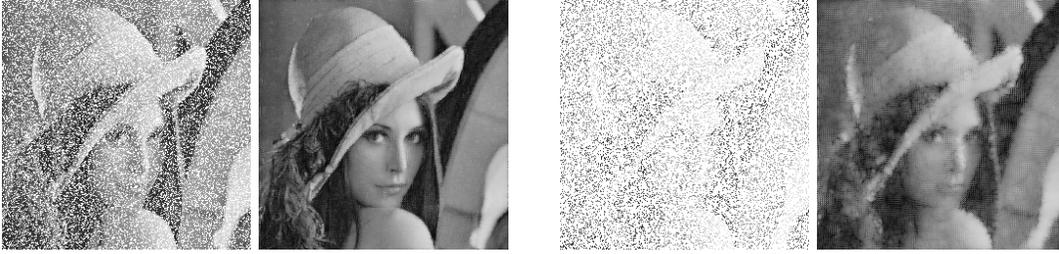


FIGURE 10. Two inpaintings via hypoelliptic diffusion with the DR procedure of images missing, respectively, 30% and 80% of pixels. For both images we chose $T = 1$ and $n = 60$ steps. Notice how in the second image, the big corruption makes it difficult to correctly compute the gradient.

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 Figures 9 and 10. Notice that restorations obtained via this method are much better than those obtained via pure hypoelliptic diffusion. In particular, the DR procedure gives reasonable results even on highly corrupted images, as shown in Figure 10.

4.2. Varying coefficients hypoelliptic diffusion

A different approach w.r.t. the Dynamic Restoration procedure is to directly modify the evolution equation (8), 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 (8) with

$$\Delta_{\mathcal{H}} = a(x, y) \left(\cos(\theta) \frac{\partial}{\partial x} + \sin(\theta) \frac{\partial}{\partial y} \right)^2 + b(x, y) \frac{\partial^2}{\partial \theta^2} \quad (16)$$

for some non-negative continuous coefficients $a, b : \mathbb{R}^2 \rightarrow \mathbb{R}$. By replacing Δ with $\Delta_{\mathcal{H}}$ in equation (8), we try to implement the natural idea of reducing the effect of the diffusion at non-corrupted points. Indeed, when using the operator $\Delta_{\mathcal{H}}$ one can tune the coefficients a, b depending on the point (x, y) , and consequently, weaken the diffusion effect only where it is necessary.

Unfortunately, from the numerical point of view, integration of equation (8) with the operator $\Delta_{\mathcal{H}}$ is much more complicated than with Δ , since the essential decoupling effect that allowed to reduce (8) to (13) does not take place anymore. In order to overcome this problem, at each time step, we approximate the varying



FIGURE 11. Three inpaintings via hypoelliptic diffusion with varying coefficients and the DR procedure: with small diagonal corruptions, and with 30% and 80% of pixels randomly removed. Compare with Figures 5 and 9.

coefficient version of (8) by a similar equation with constant coefficients. Although the constant coefficients operator used in this approximation presents a drift term, this allows to recover the decoupling effect and the Crank-Nicolson method is still pertinent applied to each of the decoupled ODE's. (See the details in [9]).

When choosing the varying coefficients a and b , the idea is to make them larger at bad points and their neighbors. This is especially true for the coefficient a , which encodes the velocity of diffusion along the level lines of the image. Thus, they will be chosen to be a smooth approximation of the indicator function of the set B of bad (corrupted) points and c_0, c_1 are positive constants. In particular, our choice is:

$$a(x, y) = \begin{cases} a_0 + a_1 \epsilon(x, y) & \text{if } \frac{a_0 + a_1 \epsilon(x, y)}{a_0 + a_1} > \epsilon_*, \\ 0 & \text{otherwise,} \end{cases} \quad b(x, y) = \begin{cases} b_0 + b_1 \epsilon(x, y) & \text{if } \frac{b_0 + b_1 \epsilon(x, y)}{b_0 + b_1} > \epsilon_*, \\ 0 & \text{otherwise,} \end{cases}$$

where

$$\epsilon(x, y) = \exp\left(-\frac{f^2(x, y)}{\sigma}\right).$$

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 11. 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 diffusion with 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.

4.3. AHE algorithm

In order to improve on the results for high corruption rates, in [9] 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 12):



FIGURE 12. A depiction of the different steps of the AHE procedure. Notice the increasing of the sharpening from step 2 to step 4.

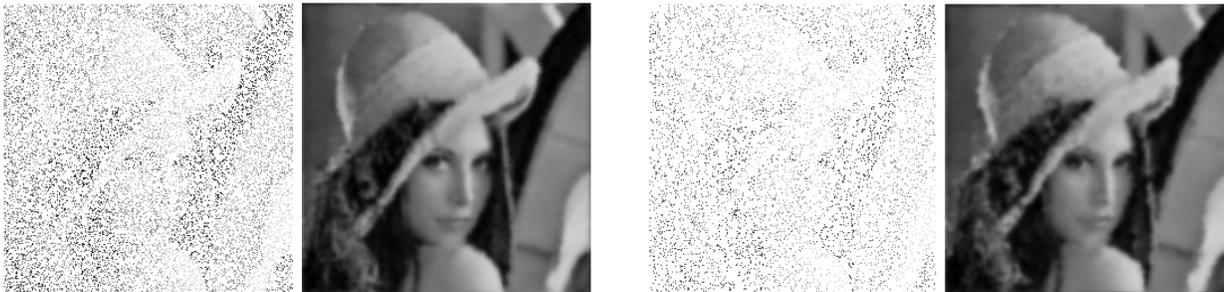


FIGURE 13. Inpaintings via the AHE algorithm of highly corrupted images with, respectively, 80% and 90% of pixels missing. Compare with Figures 9 and 11.

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 4.2.
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 13 we present the results obtained via the AHE algorithm on highly corrupted images. In particular, a comparison with Figure 10 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.

The presented inpainting techniques can be applied independently of the structure and the geometry of the corruption, although it requires the precise knowledge of its location. The quality of reconstruction strongly depends on the accuracy of this information. The quality depends also on the distribution of the corrupted pixels. Indeed, if non-corrupted pixels are well distributed, one can achieve good reconstructions even for highly corrupted images. However, if the image contains large regions entirely composed of corrupted pixels, the reconstructions are no longer satisfactory. This effect is discussed in [9].

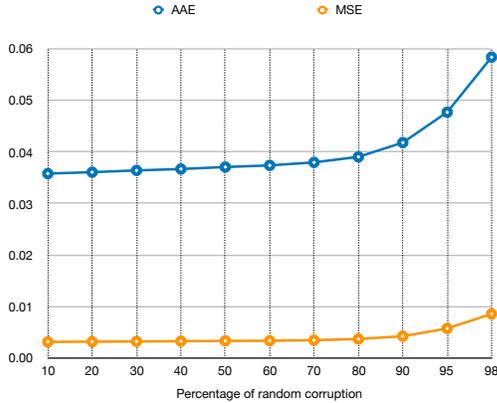


FIGURE 14. Quality of the reconstruction for a cropped version of Lena with increasing percentages of random corruption.

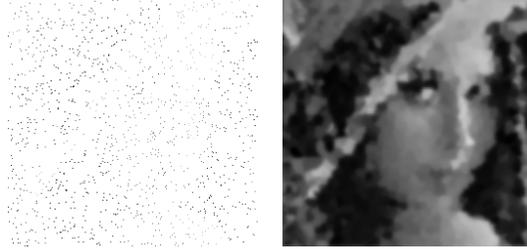


FIGURE 15. Reconstruction of cropped version of Lena with 98% of pixels missing. Compare with [23, Figure 1].

5. QUALITY OF RECONSTRUCTION

An objective comparison of quality of reconstructions is a problem of well-known difficulty, see, for instance, [34, 40, 42]. Moreover, such a measure will forcibly depend on the expected application. Nevertheless, in this section we present some results obtained with a variant of the metric used in [23]. In particular, we consider the following quantities for the quality of a reconstruction f of a corrupted version of an image g :

$$\widetilde{\text{AAE}}(f, g) = \min_{a, b \in \mathbb{R}} \text{AAE}(af + b, g), \quad \widetilde{\text{MSE}}(f, g) = \min_{a, b \in \mathbb{R}} \text{MSE}(af + b, g). \quad (17)$$

Here, we use the *average absolute error* (*AEE*) and the *mean squared error* (*MSE*), as defined in [23]. These essentially amount to a discrete version of the L^1 and L^2 norms, and are given by

$$\text{AAE}(f, g) = \frac{1}{N} \sum_{i,j} |f_{i,j} - g_{i,j}|, \quad \text{MSE}(f, g) = \frac{1}{N} \sum_{i,j} |f_{i,j} - g_{i,j}|^2, \quad (18)$$

where N is the total number of pixels in the images $f = (f_{i,j})$ and $g = (g_{i,j})$. The need for the additional minimizing procedure in (17) is due to the fact that the diffusion of the AHE algorithm generates a scaling of values, including on non-corrupted pixels. Several experiments have shown that the effect of this scaling can be represented as $af + b$ for some $a, b \in \mathbb{R}$.

In Figure 14 we present the results obtained on a 256×256 cropped version of the Lena image (the same used in [23]), with increasing percentages of corruption. We remark that, taking into account the fact that images in [23] are parametrized on the interval $[0, 255]$ while we chose $[0, 1]$, the values obtained for a 98% corrupted image are $\widetilde{\text{AAE}}(f, g) \sim 14.9$ and $\widetilde{\text{MSE}}(f, g) \sim 556.7$. (See Figure 15.) In particular, the first one is slightly worse than the edge-enhancing diffusion presented in [23, Table 1], while the second one is better. One could think this comparison not to be fair, due to the additional minimization procedure in (17), but such minimization does not seem necessary in the case of [23], since the algorithms used in that paper exactly preserve the values of the non-corrupted pixels.

Finally, in Figure 16, we present a comparison of two reconstructions, one obtained with the AHE algorithm, denoted by f_1 , and the other via the methods in [12] (courtesy of S. Masnou), denoted by f_2 . We remark that, in this case, we have $\widetilde{\text{AEE}}(f_1, g) \sim 0.09$ and $\widetilde{\text{MSE}}(f_1, g) \sim 0.05$. These differ from $\widetilde{\text{AEE}}(f_2, g)$ and $\widetilde{\text{MSE}}(f_2, g)$ by a factor of 10^{-9} and 10^{-5} , respectively.



FIGURE 16. Comparison of reconstruction of the image to the left with the AHE algorithm (*middle*) and the methods in [12] (*right*).

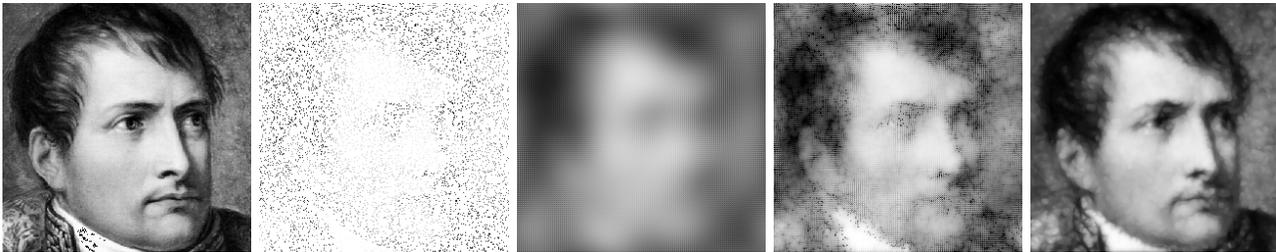


FIGURE 17. A comparison of the different algorithms presented in this paper. *From left to right*: the original image, the corrupted image used as input, the reconstructions obtained via pure hypoelliptic diffusion ($\beta^2 = \frac{1}{4}$, $T = 1$), via hypoelliptic diffusion with the DR procedure ($n = 120$), and via the AHE algorithm.

6. 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. In particular, in Section 5 we showed that the final method we presented, the AHE algorithm, yields reconstructions that are on par with the current state-of-the-art in local PDE-based methods. We also mention that the techniques presented in this paper can be applied to color images, simply by treating separately the three color channels. Some examples are provided in Figure 18.

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 17, 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 Sub-Riemannian Geometry and Interactions, 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. The second author was supported by the project POCI-01-0145-FEDER-006933/SYSTEC financed by ERDF through COMPETE2020 and by FCT (Portugal). The third and fourth authors have been supported by the CNRS INS2I PEPS project CINCIN.

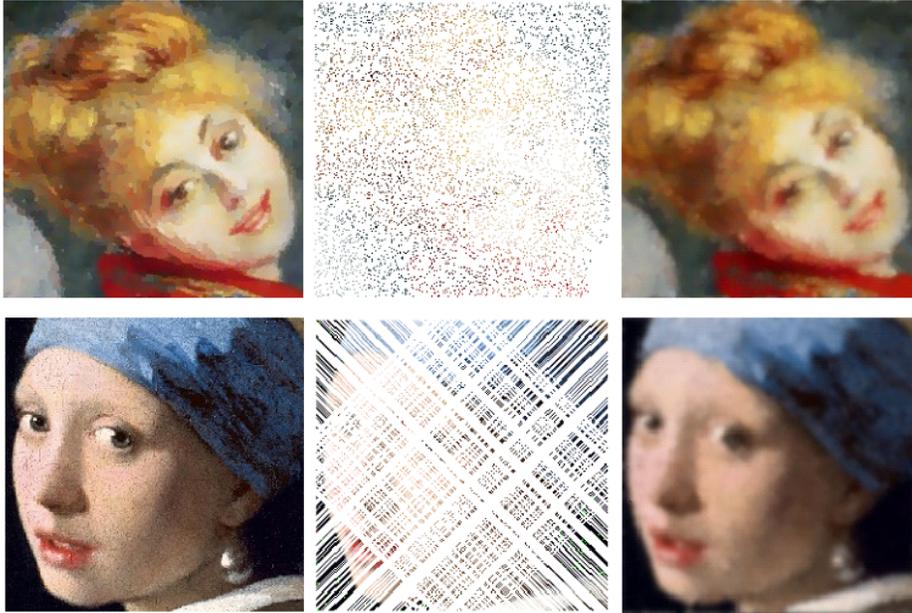


FIGURE 18. Inpaintings of color images. *From left to right:* the original image, the corrupted image used as input, the reconstructions obtained via the AHE algorithm. *Top:* Madame Monet in a Japanese costume, by C. Monet. *Bottom:* Girl with a Pearl Earring, by J. Vermeer.

REFERENCES

- [1] Y. ACHDOU AND N. TCHOU, *A finite difference scheme on a non commutative group*, Numer. Math., 89 (2001), pp. 401–424.
- [2] 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).
- [3] C. BALLESTER, M. BERTALMIÓ, V. CASELLES, G. SAPIRO, AND J. VERDERA, *Filling-in by joint interpolation of vector fields and gray levels*, Image Processing, IEEE Transactions on, 10 (2001), pp. 1200–1211.
- [4] M. BERTALMIÓ, G. SAPIRO, V. CASELLES, AND C. BALLESTER, *Image inpainting*, in Proc. of SIGGRAPH 2000, New Orleans, USA, July, 2000, pp. 417–424.
- [5] 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 (2014), pp. 669–695.
- [6] U. BOSCAIN, R. DUTTS, F. ROSSI, AND Y. SACHKOV, *Curve cusplless reconstruction via sub-Riemannian geometry*, ESAIM Control Optim. Calc. Var., 20 (2014), pp. 748–770.
- [7] U. BOSCAIN, J. DUPLAIX, J.-P. GAUTHIER, AND F. ROSSI, *Anthropomorphic image reconstruction via hypoelliptic diffusion*, SIAM J. Control Optim., 50 (2012), pp. 1–25.
- [8] 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, 2014, pp. 4278–4283.
- [9] U. V. BOSCAIN, R. CHERTOVSKIH, J.-P. GAUTHIER, D. PRANDI, AND A. REMIZOV, *Highly corrupted image inpainting through hypoelliptic diffusion*, Journal of Mathematical Imaging and Vision, (2018).
- [10] K. BREDIES, T. POCK, AND B. WIRTH, *Convex Relaxation of a Class of Vertex Penalizing Functionals*, J. Math. Imaging Vis., 47 (2013), pp. 278–302.
- [11] P. C. BRESSLOFF, J. D. COWAN, M. GOLUBITSKY, P. J. THOMAS, AND M. C. WIENER, *Geometric visual hallucinations, euclidean symmetry and the functional architecture of striate cortex*, Philosophical Transactions of the Royal Society of London B: Biological Sciences, 356 (2001), pp. 299–330.
- [12] F. CAO, Y. GOUSSEAU, S. MASNOU, AND P. PÉREZ, *Geometrically guided exemplar-based inpainting*, SIAM Journal on Imaging Sciences, 4 (2011), pp. 1143–1179.
- [13] T. CHAN, S. KANG, AND J. SHEN, *Euler’s elastica and curvature-based inpainting*, SIAM J. Appl. Math., 63 (2002), pp. 564–592.
- [14] G. CITTI, B. FRANCESCHIello, G. SANGUINETTI, AND A. SARTI, *Sub-riemannian mean curvature flow for image processing*, SIAM J. Imaging Sci., 9 (2016), pp. 212–237.

- [15] G. CITTI AND A. SARTI, *A cortical based model of perceptual completion in the roto-translation space*, J. Math. Imaging Vis., 24 (2006), pp. 307–326.
- [16] A. CRIMINISI, P. PÉREZ, AND K. TOYAMA, *Region filling and object removal by exemplar-based image inpainting*, IEEE Trans. Image Process., 13 (2004).
- [17] 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 (1985), p. 1160.
- [18] R. DUIJS, U. BOSCAIN, F. ROSSI, AND Y. SACHKOV, *Association fields via cusplless sub-Riemannian geodesics in $SE(2)$* , J. Math. Imaging Vision, 49 (2014), pp. 384–417.
- [19] R. DUIJS 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 (2010), pp. 255–292.
- [20] ———, *Left-invariant parabolic evolutions on $SE(2)$ and contour enhancement via invertible orientation scores Part II: non-linear left-invariant diffusions on invertible orientation scores*, Quart. Appl. Math., 68 (2010), pp. 293–331.
- [21] R. DUIJS AND M. A. VAN ALMSICK, *The explicit solutions of linear left-invariant second order stochastic evolution equations on the 2D euclidean motion group*, in Quart. Appl. Math., vol. 66, 2008, pp. 27–67.
- [22] G. FACCIOLO, P. ARIAS, V. CASELLES, AND G. SAPIRO, *Exemplar-based interpolation of sparsely sampled images*, in Energy Minimization Methods in Computer Vision and Pattern Recognition: 7th International Conference EMMCVPR 2009 (Bonn, Germany, August 24–27, 2009) Proceedings, D. Cremers, Y. Boykov, A. Blake, and F. R. Schmidt, eds., Springer, 2009, pp. 331–344.
- [23] I. GALIĆ, J. WEICKERT, M. WELK, A. BRUHN, A. BELYAEV, AND H.-P. SEIDEL, *Image compression with anisotropic diffusion*, Journal of Mathematical Imaging and Vision, 31 (2008), pp. 255–269.
- [24] 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 (2010), pp. 1–27.
- [25] W. C. HOFFMAN, *The visual cortex is a contact bundle*, Applied Mathematics and Computation, 32 (1989), pp. 137–167.
- [26] 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 (1962), pp. 106–154.
- [27] R. LÉANDRE, *Majoration en temps petit de la densité d'une diffusion dégénérée*, Probab. Theory Related Fields, 74 (1987), pp. 289–294.
- [28] ———, *Minoration en temps petit de la densité d'une diffusion dégénérée*, J. Funct. Anal., 74 (1987), pp. 399–414.
- [29] 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 roto-translations*, Numerical Mathematics: Theory, Methods and Applications, 6 (2013), pp. 95–115.
- [30] S. MASNOU, *Disocclusion: a variational approach using level lines*, IEEE Trans. Image Process., 11 (2002), pp. 68–76.
- [31] I. MOISEEV AND Y. L. SACHKOV, *Maxwell strata in sub-Riemannian problem on the group of motions of a plane*, ESAIM Control Optim. Calc. Var., 16 (2010), pp. 380–399.
- [32] J. PETITOT, *The neurogeometry of pinwheels as a sub-riemannian contact structure*, Journal of Physiology-Paris, 97 (2003), pp. 265–309.
- [33] ———, *Neurogéométrie de la vision: modeles mathematiques et physiques des architectures fonctionnelles*, Editions Ecole Polytechnique, 2008.
- [34] J. PETITOT AND Y. TONDUT, *Vers une neurogéométrie. fibrations corticales, structures de contact et contours subjectifs modaux*, Mathématiques informatique et sciences humaines, 145 (1999), pp. 5–102.
- [35] 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, (2017).
- [36] G. SANGUINETTI, G. CITTI, AND A. SARTI, *Image completion using a diffusion driven mean curvature flow in a sub-riemannian space*, in Proceedings of the 3rd International Conference on Computer Vision Theory and Applications (VISAPP 2008), vol. 2, 2008, pp. 46–53.
- [37] D. TSCHUMPERLÉ, *Fast anisotropic smoothing of multi-valued images using curvature-preserving PDE's*, International Journal of Computer Vision, 68 (2006), pp. 65–82.