PDE & Analysis Seminar

Geometry of the Visual Cortex with Applications to Image Inpainting and Enhancement

Francesco Ballerin

UNIVERSITY OF BERGEN

Objective: image processing

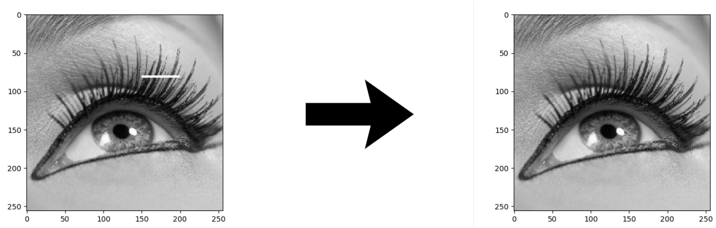

Figure: Image restoration

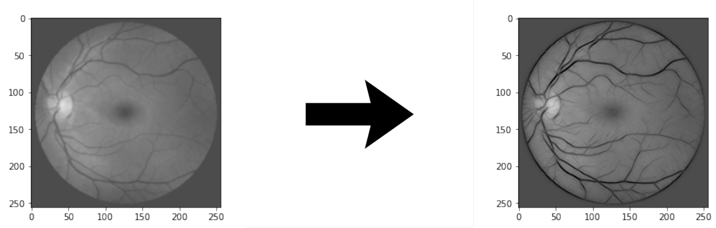

Figure: Image enhancement

Curve completion in 2D

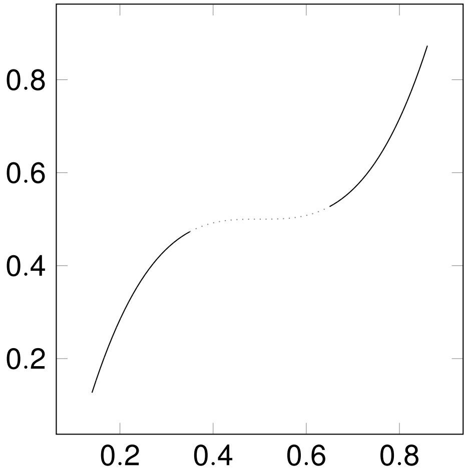

Curve completion in 2D

- $\gamma_0 : [a, b] \cup [c, d] \to \mathbb{R}^2$ a smooth curve that is partially hidden in the interval $t \in (b, c)$.
- We want to find a curve $\gamma : [b, c] \to \mathbb{R}^2$ that completes γ_0 while minimizing a cost $J[\gamma]$.
- Constraints on position: $\gamma(b) = \gamma_0(b)$, $\gamma(c) = \gamma_0(c)$
- Constraints on orientation:
  - $\dot{\gamma}(b) \sim \dot{\gamma}_0(b)$, $\dot{\gamma}(c) \sim \dot{\gamma}_0(c)$
    or
  - $\dot{\gamma}(b) \approx \dot{\gamma}_0(b)$, $\dot{\gamma}(c) \approx \dot{\gamma}_0(c)$
- $J_\beta[\gamma] = \int_b^c \sqrt{\|\dot{\gamma}(t)\|^2 + \beta\|\dot{\gamma}(t)\|^2 K_\gamma^2(t)} dt$

The visual cortex V1

- The **visual cortex V1** is the main region of a mammal's brain for **processing vision.**
- Composed of simple cells that are sensitive to **position** and **orientation.**
- Simple cells that receive enough stimulus from an external source **spike.**
- Simple cells that spike stimulate other simple cells with **same position but different orientation**, and simple cells with **same orientation and close position.**

The visual cortex V1

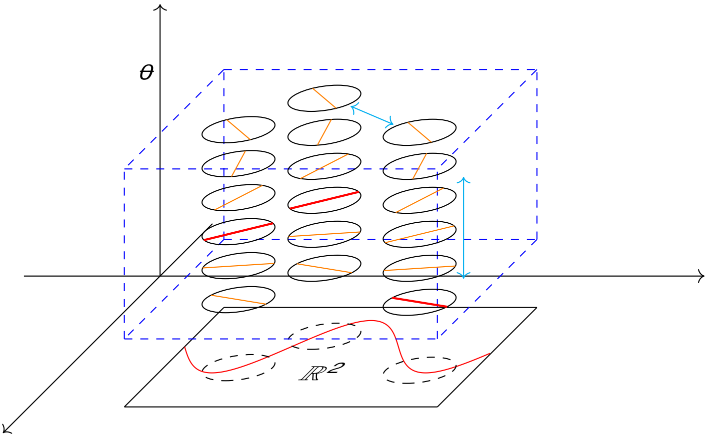

Figure: Hypercolumns of the Visual Cortex V1 under a stimulus (red curve)

The visual cortex V1

$$SE(2) = \left\{ \begin{bmatrix} \cos\theta & -\sin\theta & x \\ \sin\theta & \cos\theta & y \\ 0 & 0 & 1 \end{bmatrix} \;\middle|\; \begin{matrix} x, y \in \mathbb{R}, \\ \theta \in \mathbb{R}/2\pi\mathbb{Z} \end{matrix} \right\}$$

Using coordinates (x, y, θ) we see that $SE(2)$ as a space can be identified with the 3-dimensional cylinder $\mathbb{R}^2 \times S^1$.

$$X_1 = \cos(\theta)\partial_x + \sin(\theta)\partial_y, \qquad X_2 = \partial_\theta, \qquad X_3 = -\sin(\theta)\partial_x + \cos(\theta)\partial_y$$

$$[X_1, X_2] = -X_3, \qquad [X_2, X_3] = X_1, \qquad [X_1, X_3] = 0.$$

Let $\mathcal{H} = \text{span}\{X_1, X_2\}$ be the horizontal distribution, then $SE(2)$ endowed with $\mathcal{H}$ is a **bracket-generating sub-Riemannian manifold.**

Lift of curve from $\mathbb{R}^2$ to $SE(2)/\sim$

Consider $SE(2)/\sim$, where $(x, y, \theta) \sim (x, y, \theta + \pi)$.

A curve is lifted (by $\mathcal{L}$) from $\mathbb{R}^2$ to $SE(2)/\sim$ by adding a coordinate $\theta \in [0, \pi)$ corresponding to the angle between the orientation of the curve and the horizontal axis $y = 0$, measured counterclockwise.

A lifted curve is projected (by Π) to $\mathbb{R}^2$ simply by suppressing the third coordinate corresponding to orientation.

Curve completion in $SE(2)/\sim$

Completing a curve in $\mathbb{R}^2$ is equivalent to finding a minimizer in $SE(2)/\sim$, where $(x, y, \theta) \sim (x, y, \theta + \pi)$.

Proposition (Boscain, Charlot, Rossi - 2010)

For all boundary conditions $\gamma_0(b), \gamma_0(c) \in \mathbb{R}^2$ *with* $\gamma_0(b) \neq \gamma_0(c)$ *and* $\dot{\gamma}_0(b), \dot{\gamma}_0(c) \in \mathbb{R}^2 \setminus \{0\}$ *the cost* $J_\beta[\gamma]$ *admits a minimizer over the set*

$$\mathcal{D} = \left\{ \gamma \in C^2([b,c], \mathbb{R}^2) \middle| \begin{array}{c} \|\dot{\gamma}\|^2 + \|\dot{\gamma}\|^2 K_\gamma^2 \in L^1([b,c], \mathbb{R}) \\ \gamma(b) = \gamma_0(b), \gamma(c) = \gamma_0(c) \\ \dot{\gamma}(b) \approx \dot{\gamma}_0(b), \dot{\gamma}(c) \approx \dot{\gamma}_0(c) \end{array} \right\}$$

Lift of an image

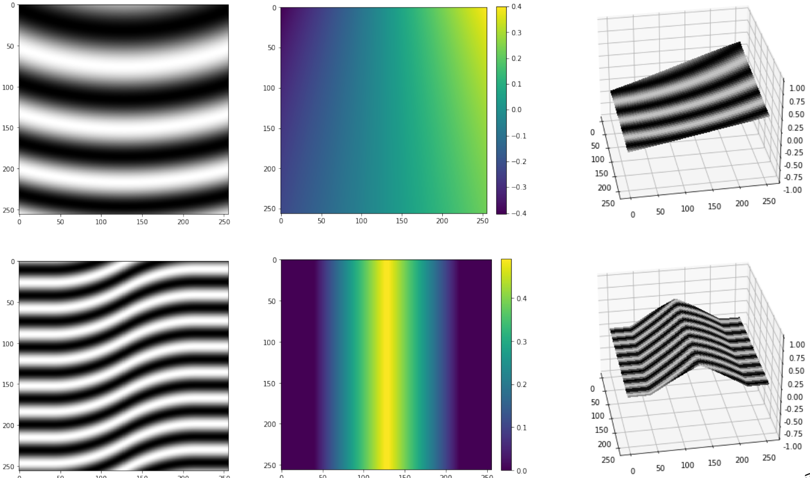

Beyond curve completion: image diffusion

- Not working with a single curve but with **many level curves.**
- Consider all possible admissible paths and model the controls by independent **Wiener processes** u_t and v_t obtaining the following SDE:

$$\begin{pmatrix} dx_t \\ dy_t \\ d\theta_t \end{pmatrix} = \sqrt{2} \begin{pmatrix} \cos\theta_t \\ \sin\theta_t \\ 0 \end{pmatrix} \circ du_t + \sqrt{2\beta} \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \circ dv_t$$

- The diffusion process associated to such SDE is

$$\frac{\partial \Psi}{\partial_t} = \Delta \Psi$$

where

$$\Delta = X_1^2 + \beta X_2^2 = \left(\cos\theta \frac{\partial}{\partial_x} + \sin\theta \frac{\partial}{\partial y} \right)^2 + \beta \frac{\partial^2}{\partial \theta^2}.$$

The Citti-Sarti-Boscain algorithm

1. **Smooth** the image by convolution with a Gaussian kernel (guarantees the image is generically a Morse function)

2. **Lift** $\mathbb{R}^2 \to SE(2)$

3. Solve the Cauchy problem

$$\Delta_\beta = X_1^2 + \beta X_2^2$$

$$\begin{cases} \partial_t u = \Delta_\beta u, \\ u(0, x, y, \theta) = \tilde{I}(x, y, \theta) \end{cases}$$

   over $SE(2)$

4. **Project** $SE(2) \to \mathbb{R}^2$

Some examples

(a) Original image (b) $T = 21$ with $\beta = 0.1$ (c) $T = 21$ with $\beta = 1$ (d) $T = 21, \beta = 10$

Figure: An image (a) is processed with the Citti-Sarti-Boscain algorithm for different values β (b,c,d)

Problem: no code available

Due to drastic and unfortunate events (allegedly), including but not limited to

- the laptop of the researcher that had the code was stolen
- the backup code was lost in a terrible fire

no code from the first paper survived to this day.

A codebase for the Boscain et al. was provided by Francesco Rossi: MATLAB implementation.

However some work was needed to replicate the results from the papers.

Public repository

v1diffusion Public

Pin · Unwatch 1

main · 1 branch · 0 tags · Go to file · Add file · Code

ballerin Refactored code and added initial jupyter notebook examples · e4a1403 3 weeks ago · 2 commits

images	Refactored code and added initial jupyter notebook examples	3 weeks ago
jupyter_notebooks	Refactored code and added initial jupyter notebook examples	3 weeks ago
.gitignore	first commit	2 months ago
README.md	first commit	2 months ago
__init__.py	first commit	2 months ago
evolution.py	Refactored code and added initial jupyter notebook examples	3 weeks ago
examples.py	first commit	2 months ago
metrics.py	first commit	2 months ago
operators.py	Refactored code and added initial jupyter notebook examples	3 weeks ago
processing.py	Refactored code and added initial jupyter notebook examples	3 weeks ago
requirements.txt	Refactored code and added initial jupyter notebook examples	3 weeks ago
transforms.py	Refactored code and added initial jupyter notebook examples	3 weeks ago
utils.py	first commit	2 months ago

README.md

v1diffusion

This is intended as a supporting package for the papers related to neurogeometry of vision of the primary visual cortex v1, and the image processing applications exploiting the geometrical structure of the space $SE(2)$.

References:

[1] : G Citti and A Sarti. "A cortical based model of perceptual completion in the Roto-translation space". In: J. Math. Imaging Vis. 24.3 (May 2006), pp. 307– 326

[2] : U Boscain et al. "Anthropomorphic Image Reconstruction via Hypoelliptic Diffusion". In: SIAM j. control optim. 50.3 (Jan. 2012), pp. 1309–1336.

[3] : U Boscain et al. "Hypoelliptic diffusion and human vision: A semidiscrete new twist". In: SIAM J. Imaging Sci. 7.2 (Jan. 2014), pp. 669–695.

In [1]:
```python
import numpy as np
import jax.numpy as jnp
import cv2
import matplotlib.pyplot as plt

import sys
sys.path.append('../../')
from v1diffusion import utils, transforms, operators, evolution, processing, metrics, examples
```

In [2]:
```python
x=np.linspace(0,255,256)
y=np.linspace(0,255,256)
X,Y = np.meshgrid(x,y)
```

In [3]:
```python
R = cv2.imread("../images/open_retinal_256.jpg",cv2.IMREAD_GRAYSCALE).astype(np.float32)
R/=255
R = 1-R
X,Y = np.meshgrid(x,y)
V = (X-128)**2+(Y-128)**2
R[V>15000]=0.7
R_h = cv2.GaussianBlur(R,(15,15),0.6)
R_hh = cv2.GaussianBlur(R,(15,15),1)
```

In [4]:
```python
R_p = R.copy()
x=np.linspace(0,255,256)
y=np.linspace(0,255,256)
X,Y = np.meshgrid(x,y)
V = (X-128)**2+(Y-128)**2
R_p[V>15000]=0.7
R_ph = cv2.GaussianBlur(R_p,(15,15),0.6)
```

In [5]:
```python
utils.imshow(R_h)
```

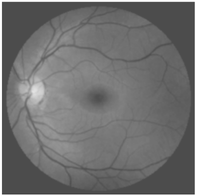

In [6]:
```python
#Usual WaxOn
```

Gaussian lift

In the work by Marcelja and Jones and Palmer the **similarity in behavior between simple cells and Gabor filters** is studied and presented. The output of a signal through the filter decays exponentially as the angle of the original signal differs from θ.

We can model each fiber as a normal distribution around the angle of the level curve, effectively "spreading" the input signal around the orientation of maximum response θ of the simple cells following a Gaussian distribution

$$\mathcal{L}_\sigma(I) = (I \circ \Pi) \cdot \exp\left(-\frac{(X_1(I \circ \Pi))^2}{2\sigma^2 |\nabla I|^2}\right) =$$

$$= (I \circ \Pi) \cdot \exp\left(-\frac{|\nabla I|^2 - (X_3(I \circ \Pi))^2}{2\sigma^2 |\nabla I|^2}\right)$$

Gaussian lift

Theorem

Define an operator $\Pi_\sigma : C^\infty(\mathrm{SE}(2), (0,1]) \to C^\infty(\mathbb{R}^2, (0,1])$ *by*

$$\Pi_\sigma(\tilde{I})(x,y) = \exp\left(\frac{1}{4\sigma} + \frac{1}{2\pi}\int_0^{2\pi} \ln \tilde{I}(x,y,\theta)d\theta\right).$$

Then $\Pi_\sigma(\mathcal{L}_\sigma(I)) = I$.

Moreover, the original lift can be considered as a limiting case when $\sigma \to 0$.

How to deal with blur: unsharp filtering

$$I \mapsto I + C(I - I * G_\sigma)$$

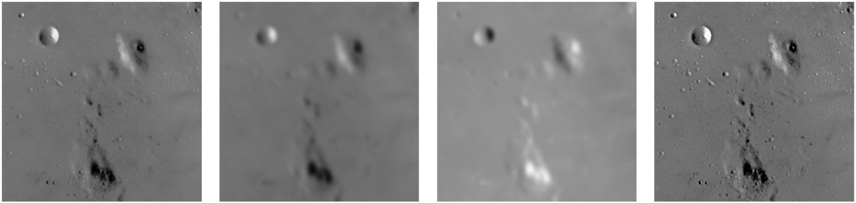

(a) Original image (b) Blurred image with $\sigma = 5$ (c) Negative of the blurred image (d) Sharpened image with $C = 1$

Figure: Example of usage of the unsharp filter applied to a low-contrast image of the surface of the moon

Vector fields X_1, X_2, and X_3

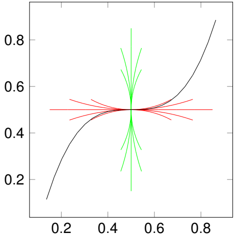

(a) Projections to $\mathbb{R}^2$ of the integral lines

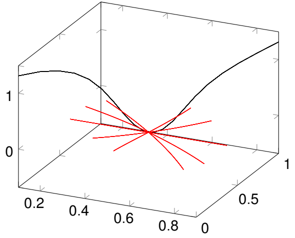

(b) Integral lines of $X_1^2 +_2^2$

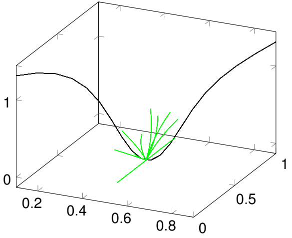

(c) Integral lines of $X_3^2 +_2^2$

Figure: Integral lines of the vector fields $X_1^2 + \beta X_2^2$ (red) and $X_3^2 + \beta X_2^2$ (green) for a polynomial curve, at point $\left(\frac{1}{2}, \frac{1}{2}\right)$, varying the coefficient β.

Sketch of the idea

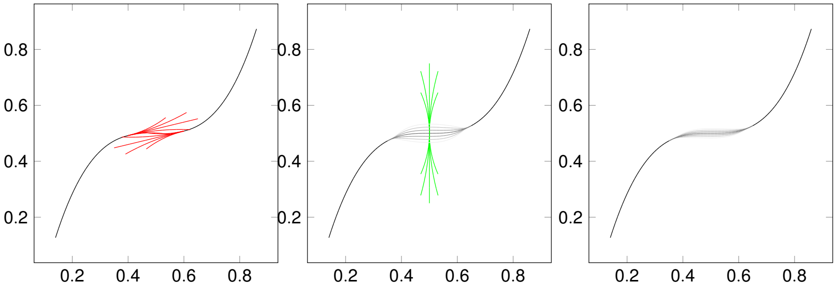

(a) `WaxOn` (b) `WaxOff` (c) Final result

Figure: Sketch of intuition behind `WaxOn-WaxOff`

Unsharp filtering on $SE(2)/\sim$

The **undesired blurring** is obtained from the Cauchy problem

$$\begin{cases} \partial_t u = \Delta_\beta u, \\ u(0,x,y,\theta) = \tilde{I}(x,y,\theta) \end{cases} \qquad \Delta_\beta = X_3^2 + \beta X_2^2$$

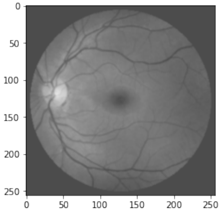

(a) Original image

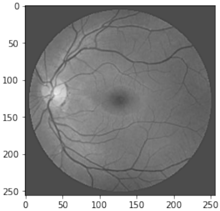

(b) $\mathbb{R}^2$ unsharp filter

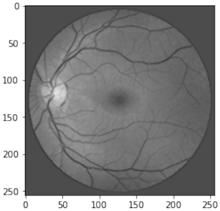

(c) $SE(2)$ unsharp filter

Figure: Retinal image (a) sharpened using the classical unsharp filter over $\mathbb{R}^2$ (b) and using the proposed sharpening method (c).

Retinal image enhancement through unsharp filtering

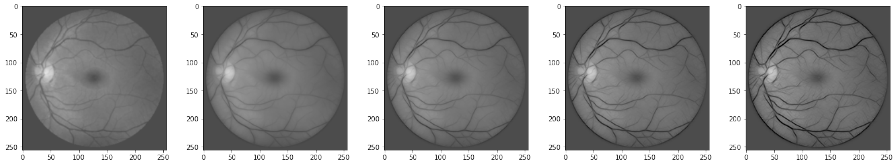

(a) Original image (b) $C = 0.5$ (c) $C = 1$ (d) $C = 1.5$ (e) $C = 2$

Figure: The original image (a) is processed with diffusion under $\Delta_\beta = X_1^2 + \beta X_2^2$ and sharpended with varying coefficients C (b,c,d,e)

WaxOn-WaxOff

- WaxOn: diffusion problem with $\Delta_\beta = X_1^2 + \beta X_2^2$
- WaxOff: inverse problem with $\Delta_\beta = X_3^2 + \beta X_2^2$

(a) Original image (b) $X_1^2 + \beta X_2^2$ diffusion (c) 1 cycle of WaxOn-WaxOff (d) 3 cycles of WaxOn-WaxOff

Figure: From (a) the CSB algorithm is applied to obtain (b). One iteration of WaxOn-WaxOff for small T_2 produces (c) while multiple iterations of WaxOn-WaxOff are sequentially applied to produce (d). Total diffusion time in (b) and (d) is the same.

References

Ballerin and Grong
'Geometry of the visual cortex with applications to image inpainting and enhancement', 2023
arXiv, 2308.07652

Citti and Sarti
'A cortical based model of perceptual completion in the Roto-translation space', 2006
J. Math. Imaging Vis. 24.3, pp. 307–326

Boscain, Charlot, and Rossi
'Existence of planar curves minimizing length and curvature', 2010
Proc. Steklov Inst. Math. 270.1, pp. 43-56

Boscain, Duplaix, Gauthier, and Rossi
'Anthropomorphic Image Reconstruction via Hypoelliptic Diffusion', 2012
SIAM j. control optim. 50.3, pp. 1309–1336

sR manifold

Definition

A sub-Riemannian manifold is a triplet $(M, \mathcal{H}, g)$ with M being a connected manifold, $\mathcal{H} \subset TM$ a linear subbundle and $g = \langle \cdot, \cdot \rangle$ a fiber-metric defined on on the subbundle $\mathcal{H}$.

We call $\mathcal{H} \subset TM$ in this definition the *horizontal distribution*. A sub-Riemannian manifold can be considered as a limiting case of a Riemannian manifold where the distances of vectors outside of $\mathcal{H}$ approach infinity. Curves $\gamma : [a, b] \to M$ with a finite length will then need to be *a horizontal curve*: an absolutely continuous curve satisfying $\dot{\gamma}(t) \in \mathcal{H}_{\gamma(t)}$ for almost every t. For such a curve, we can define its length by

$$\text{length}(\gamma) = \int_a^b \langle \dot{\gamma}(t), \dot{\gamma}(t) \rangle^{1/2} \, dt.$$

sR distance

We can then also introduce the corresponding sub-Riemannian distance by

$$d_g(x,y) = \inf\left\{\text{length}(\gamma) \ : \ \begin{array}{c} \gamma : [a,b] \to M \text{ horizontal} \\ \gamma(a) = x, \gamma(b) = y \end{array}\right\}.$$

In general, there might not be any curve connecting a point x and y, meaning that the distance above will be infinite. It is therefore typical to require the horizontal bundle sub-Riemannian manifold to be bracket-generating

Bracket generating distribution

$$\hat{\mathfrak{X}}_{\mathcal{H}} = \operatorname{span}\left\{[X_{i_1},[X_{i_2},[\cdots[X_{l=1},X_l]]\cdots]] \mid X_{i_j} \in \mathfrak{X}_{\mathcal{H}}, l = 1,2,3,\ldots,\right\},$$

where we interpret the case $l = 1$ simply as the vector field X_{i_1} itself. We then make the following definition.

Definition

We say that $\mathcal{H}$ is bracket-generating if for every $x \in M$,

$$T_x M = \{X(x) \,:\, X \in \hat{\mathfrak{X}}_{\mathcal{H}}\}.$$

Sub-Laplacian

Consider a second order operator L on a manifold M, which in local coordinates

$$L = \sum_{i,j=1}^{n} a_{ij}(x) \frac{\partial^2}{\partial x_i \partial x_j} + \sum_{j=1}^{n} b_j(x) \frac{\partial}{\partial x_j},$$

with $(a_{ij}(x))$ being positive semi-definite with a constant rank k. Such an operator can locally be written as $L = \sum_{=1}^{k} X_k^2 + X_0$. Define a sR structure on $(\mathcal{H}, g)$ on M by making $X_1, \ldots, X_k$ into a local orthonormal basis. If L is required to be symmetric, i.e. $\int_M f_1(Lf_2)d\mu = \int_M f_2(Lf_1)d\mu$ for any pair of smooth functions $f_1, f_2 \in C_0^\infty(M)$ of compact support, then the *symmetric* operator is unique with respect to a given volume density $d\mu$. We call this operators *the sub-Laplacian of* $(M, \mathcal{H}, g)$ *and* $d\mu$.

Hypoellipticity of Δ_β

Theorem

Let L be the sub-Laplacian of a sub-Riemannian structure $(M, \mathcal{H}, g)$ with volume element $d\mu$. Assume that $\mathcal{H}$ is bracket-generating. Then L and the heat operator $\partial_t - L$ are hypoelliptic. Furthermore, for the heat-semigroup e^{tL}, we have density

$$e^{tL} f(x) = \int_M p_t(x, y) f(y) \, d\mu,$$

where $p_t(x, y)$ is a smooth, strickly positive function that is symmetric in x and y. Furthermore, we have short time asymptotics

$$\lim_{t \downarrow 0} 2t \log p_t(x, y) = d_g(x, y).$$