Theorizing and Modeling in Neuroscience

Chapman University, January $16 - 17$ 2025

Geometrical models for the functional architecture of the visual area V1

Jean Petitot,
CAMS, EHESS, Paris,
jean.petitot@ehess.fr

Many thanks to Marco Panza and the organizers for their invitation.

I apologize for not being able to be with you, and I also apologize for my pitiful English.

I would also like to express my deepest sympathy at this time of dramatic fires in Los Angeles.

2. Introduction

I'm going to talk about neuromathematics, but in a rather special sense.

My purpose concerns *the internal geometry of the connectome* of some primary visual areas. It is an immanent internal neural geometry of connections.

I'm going to summarize the work I've done on the very specific internal geometry of these *functional architectures*. What I called in 1990s *Neurogeometry*.

Neurogeometry is more than a certain type of neural model. It seeks to answer a theoretical problem.

3. Introduction

The basic theoretical problem I would like to address concerning low-level vision (essentially V1) is the following.

Il we look at the intuitive phenomenal geometry (not at the internal geometry of underlying neural connections but at the phenomenal geometry) of natural visual percepts, it is an empirical evidence that the visual brain is able to perform a lot of *differential* routines: local orientations tangent to boundaries, curvatures, cusps, crossings, etc.

For example, a closed regular contour that is the boundary of a shape is the envelope of its tangents. It is the integral of its tangent directions. So, it is an object of intuitive differential geometry.

4. Introduction

How can such routines be *neurally* implemented? Indeed at their resolution scale (defined by the size of their receptive field), cortical visual neurons are "point-like" processors .

They code numerical values, for example by their firing rate (rate coding) or by their rank coding (Simon Thorpe).

But it seems impossible to compute differential routines with "point-like" processors.

More precisely :

Cortical visual neurons detect very *local* geometric cues at retinal positions. One of the main theoretical problem of low level vision is to understand how these local *infinitesimal* cues can be *detected* and *integrated* so as to generate the *global phenomenal* geometry of the images perceived, with all the well-known global spectacular phenomena studied since Gestalt theory (for example long range illusory contours of Kanizsa type). How neurons can do that !

So, just as we wonder how language, cognition, mind, consciousness can *emerge* from interactions between neurons linked by connections, we're going to ask how the intuitive differential geometry of visual phenomenal percepts can emerge from neural networks.

In a word, what are the neural correlates of differential and integral calculus? What can be the structure of neural functional architectures enabling such an emergence.
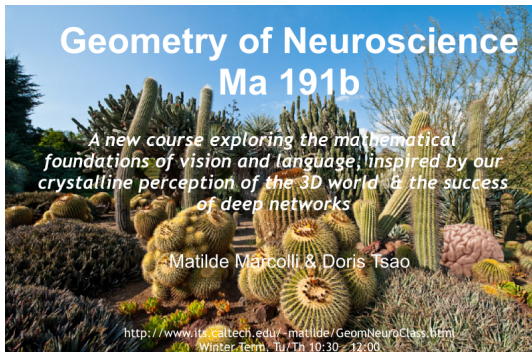
It's a "hard problem", an empirical, mathematical and philosophical foundational problem.

At the beginning, that I'm to talk about now has benefited greatly from my contacts with vision neuroscientists in Paris, such as Michel Imbert, Alain Berthoz and Yves Frégnac. And also with vision geometers as Jan Koenderink and David Mumford.

Specific research programs are now studying the problems raised by the neural correlates of differential geometry. For example, that of Alessandro Sarti and Giovanna Citti, or that of Matilde Marcolli and Doris Tsao.

"It is unimaginable that a viable representation of the brain/mind interface is possible without being incorporated into a broader mathematical framework."
-Mikhail Gromov, "Structures, Learning, and Ergosystems"

9. Experiments

At the experimental level, since the 1990s, revolutionary methods of "in vivo optical imaging" enabled to *visualize* the extremely special *connectivity* of the primary visual areas, that is their "functional architectures".

What we called at that time "*Neurogeometry*" is based on the discovery that these *hardwired* and *modular* functional architectures implement sophisticated mathematical structures such as the contact structure and the sub-Riemannian geometry of the jet spaces of plane curves. I will explain these terms.

Such structures provide a *geometrical reformulation* of differential calculus which explains how the visual brain, despite the fact that cortical neurons are point-like processors, can implement differential routines.

Experimentally, the story begins with the breakthrough recordings of V1 neurons in the early 60s by David Hubel and Torsten Wiesel (Nobel prizes in 1981).

These neurons detect a *preferred orientation p* crossing their receptive field centered on a retinal position a. When they are activated they fire and emit spikes and the spikes can be recorded using electrodes.

Orientation is what Hubel called an "engrafted" variable over a position $a = (x, y)$ of the retinal plane.
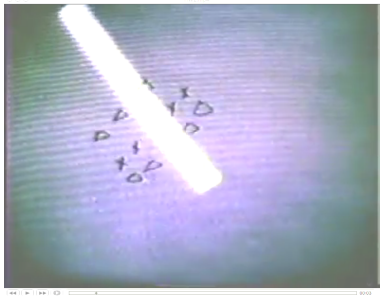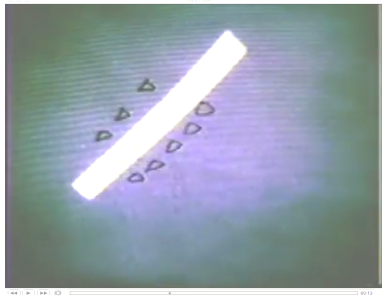
It is really a crucial experiment.

Here are two images of a 40s recording.
Left. A bar aligned along the preferred orientation : *noisy firing*.
Right. A bar orthogonal to the preferred orientation : *quiet, no firing*.

Moreover, Hubel and Wiesel discovered that

Neurons detecting all the orientations p at the same retinal position $a \in \mathbb{R}^2$ constitute an anatomically well delimited small neural module called an "orientation hypercolumn".

and that preferred orientations p vary smoothly with the retinal point a. So the (a, p) constitute an *orientation field* over the retinal plane.

Moreover, the retinotopy of the retino-geniculo-cortical way means that the *fiber bundle* $\pi : \mathbb{V} = \mathbb{R}^2 \times \mathbb{P}^1 \to \mathbb{R}^2$ is neurally implemented.

13. Braitenberg abduction

The first global reconstruction of an orientation field from the sparse local data provided by electrodes was *infered* abductively in 1979 by Valentino and Carla Braitenberg.

This was long before the introduction of modern *in vivo* optical imaging techniques.

They claimed:

> *"We believe that the most natural explanation of the facts observed would be in terms of orientations arranged with circular symmetry around centers, either radially or along concentric circles."*

14. Swindale's abduction

After Braitenberg, in an astonishing 1987 paper (still before the advent of optical imaging techniques), Nicholas Swindale reconstructed (for the cat's area 18), the "spatial layout" of the orientation map.
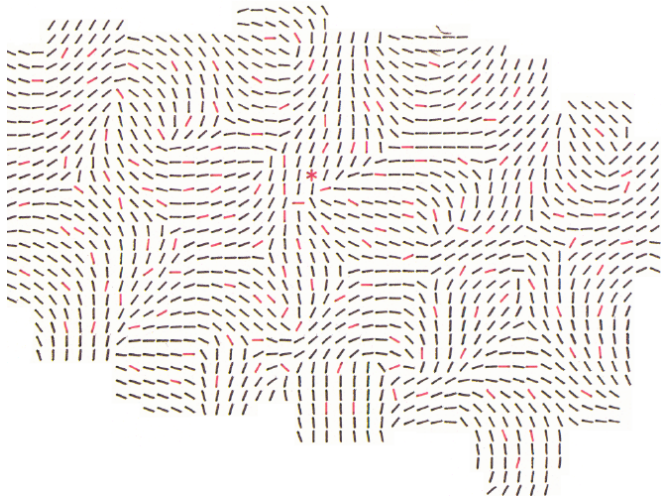
He thus confirmed Braitenberg's abduction.

His data came from electrodes separated by about $150 - 300 \mu m$ at a cortical depth of about $400 - 700 \mu m$.

He succeeded in *interpolating* between the prefered orientations measured at the different sites and reconstructed the "fine grained" map shown in the following figure.
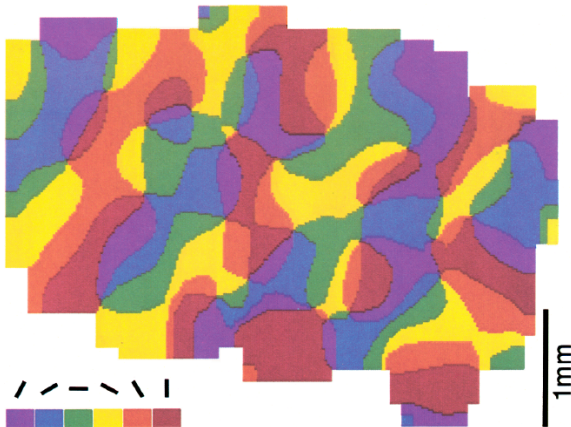
It was a great achievement.

16. Swindale's image 2

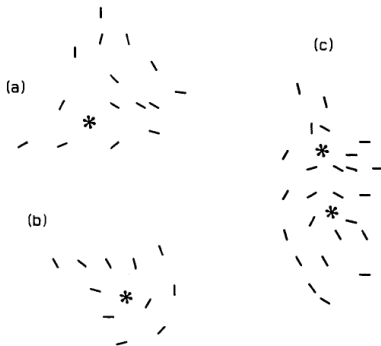Using a color code for directions, he got an orientation map.

This is a theoretical reconstruction from incomplete sparse data and not an empirical observation.

He even reconstructed the possible *singularities* of the orientation field: they can be end points or triple points.

What are now called *pinwheels*.

Braitenberg's and Swindale's abductions have been strikingly confirmed in the 1990s by brain imagery and techniques of "in vivo optical imaging based on activity-dependent intrinsic signals"(Amiram Grinvald and Tobias Bonhöffer).

They used the fact that the metabolic activity of cortical layers change their optical properties (differential absorption of oxyhemoglobin or deoxyhemoglobin whose fluorescence is an index of the local depolarisation of neurons).

This enables to acquire *in vivo* images of the activity of the superficial cortical layers.

As Kenichi Ohki and Clay Reid have pointed out,

> *"optical imaging revolutionized the study of functional architecture by showing the overall geometry of functional maps."*

The scale of observation is a "meso"-scale.

It is the "overall geometry of functional maps" we want to model using appropriate mathematical geometrical tools.
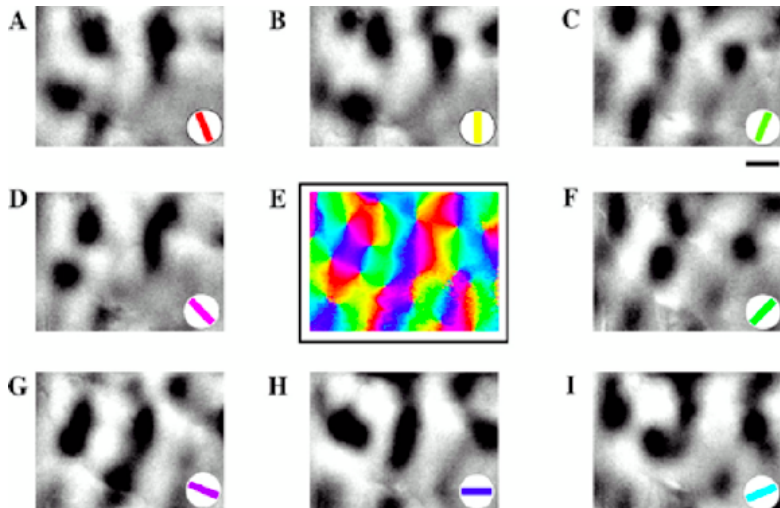
20. Orientation maps

Here is the functional architecture of the area $V1$ of a tree-shrew (tupaya) obtained by in vivo optical imaging (William Bosking with David Fitzpatrick's team at Duke University).

They used "grattings", that is large grids of parallel dark stripes translated in the visual field.

For every orientation (coded by the bottom-right color) they got a global map of activity (dark = active).
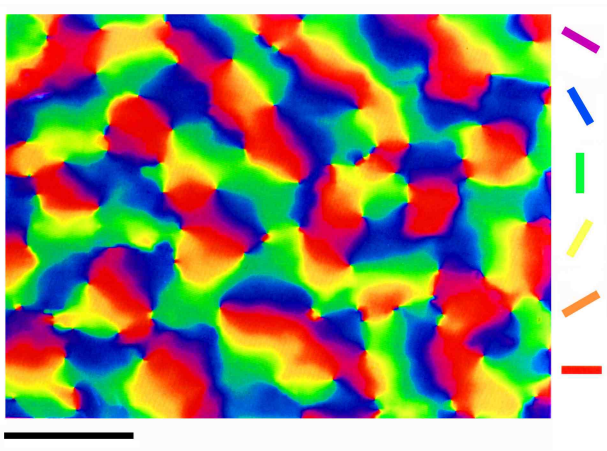
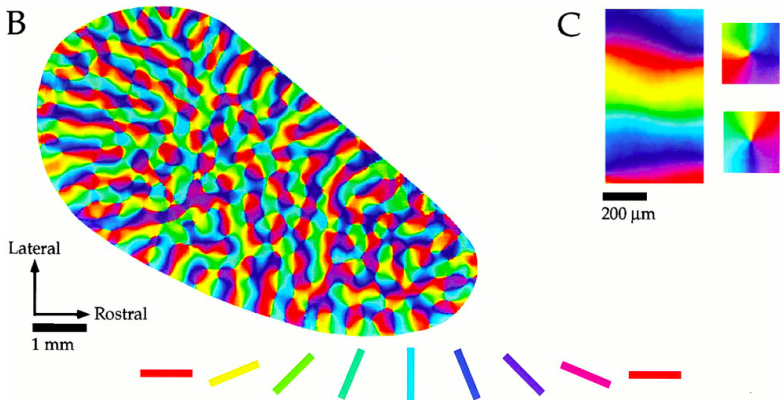This is now an empirical observation and not a theoretical reconstruction.

22. Pinwheels

Orientation maps with pinwheels are now well known. Here is the $V1$ area of the macaque by Blasdel and Salama,

Here is another image due to Bosking.

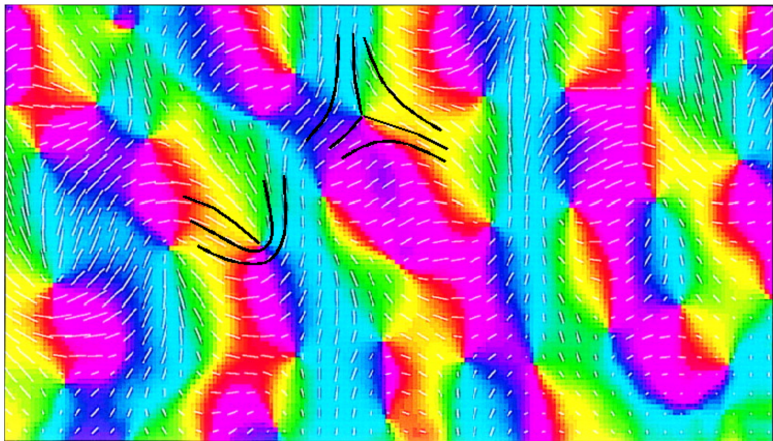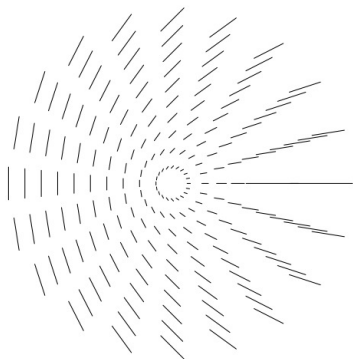In the following picture due to Shmuel (cat's area 17), orientations are coded by colors but are also represented by small white segments.

We observe very well the two types of generic singularities of 1D foliations in the plane (end points and triple points) anticipated by Braitenberg and Swindale.
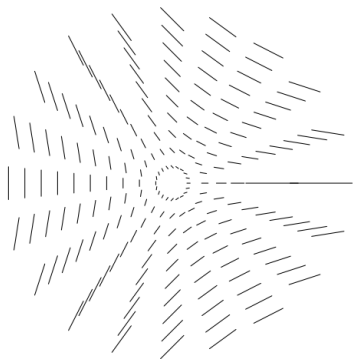
28. Pinwheels' structure

- The plane is $V1$,
- A colored point represents the mean of a small group of real neurons (meso-scale).
- Colors code for the preferred orientation at each point.
- The field of isochromatic lines (i.e. iso-orientation lines) is organized by a lattice of *singular points* (pinwheels) where all orientations meet (distant about 1200μm in cats and about 600μm in primates).
- There exists a "mesh" of the lattice of pinwheels (a sort of characteristic length).
- Pinwheels have a chirality.
- Adjacent pinwheels have opposed chirality.

29. Pinwheels

Left: local triviality. Right: two pinwheels of opposed chiralities.

200 μm
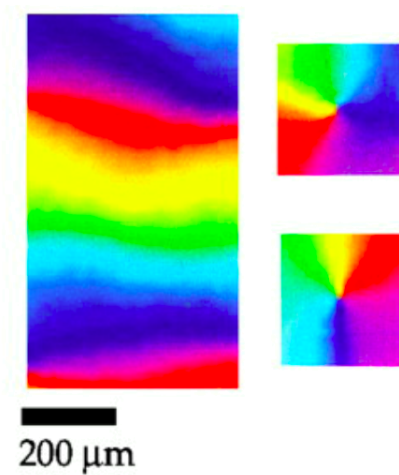
30. Micro-scale pinwheels

For a true "micro"-scale observation at the level of single neurons, you need more recent techniques such as "two-photon confocal microscopy" (Kenichi Ohki 2006).
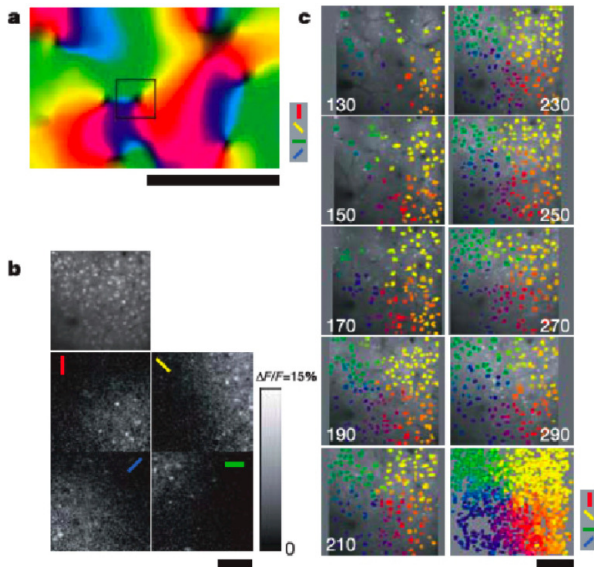
32. Interspecific functional architecture

A pinwheel organisation can be found in many species: cat, primate (marmoset), tupaya (tree shrew), prosimian Bush Baby, tawny owl, etc.

It is a widely interspecific functional architecture.

The following figure shows pinwheels in the $V1$ and $V2$ areas of the cat, (A) and (B), the marmoset (C) and (D), and the tawny owl (E) and (F).

33. Pinwheels in different species

1 mm

34. Pinwheels as blowing-up

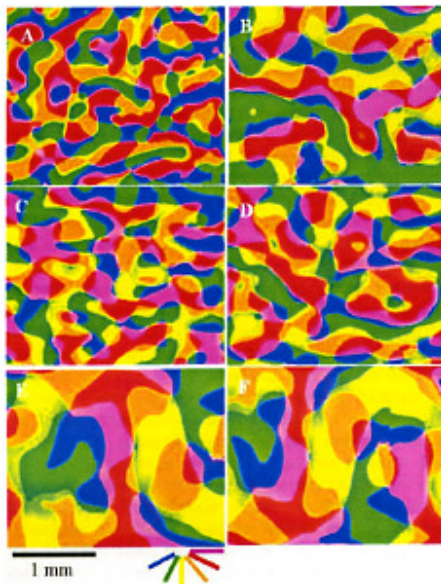Pinwheels can be interpreted geometrically as *blowing-up* of points a_i and the orientation field is the closure of a section σ of the fiber bundle $\pi : \mathbb{V} = \mathbb{R}^2 \times \mathbb{P}^1 \to \mathbb{R}^2$ defined over the open subset $\mathbb{R}^2 - \{a_i\}$.

Over the a_i the closure of σ is the "exceptionnal" fiber $\mathbb{P}^1_{a_i}$.

These exceptional fibers $\mathbb{P}^1_{a_i}$ are "contracted" and "folded" onto small wheels around the base points a_i.
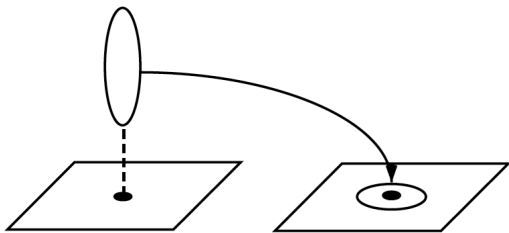
There is therefore a $3D \rightarrow 2D$ dimensional collapse : an orientation map is, in a way, a geometric object of "intermediate" dimension between 2 and 3, with a lattice of base points blown-up in parallel,

At the limit, when all the points of the base plane $\mathbb{R}^2$ are blown-up in parallel, we get the fiber bundle $\pi : \mathbb{V} = \mathbb{R}^2 \times \mathbb{P}^1 \rightarrow \mathbb{R}^2$.

So $\mathbb{V}$ can be considered as an *idealized continuous model* of the concrete neural $V1$ produced by biological evolution with its orientation field and lattice of pinwheels.

A lot of other experimental details can be found in my book
Elements of Neurogeometry (Springer, 2017)
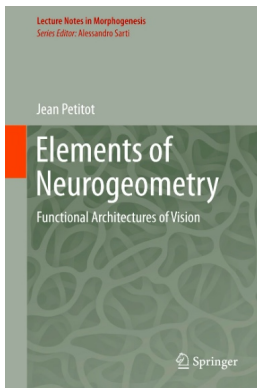
But now we run into the "hard problem" evoked in our Introduction.

The *geometry* of visual perception involves many *differential* computations. But neurons are (scaled) *point-like processors*. When they are activated, they emit spikes defining their "rate coding". And so, they can only code a single numerical value by means of their "firing rate".

Of course
(i) they are able to detect complex point-like cues and
(ii) they are connected and they can transmit their activity along their more or less inhibitory or excitatory connections.

But this is insufficient to *directly* implement differential routines.

There is therefore an *antinomy* at the root of a neurally implemented perceptual geometry.

How differential routines can be neurally implemented in networks of "point-like processors" since derivatives are not point-like entities?

The classical conceptions of "differentiation" and "integration" do not work.

Now, we have seen that biological evolution has introduced *new* post-retinal modules and layers that implement *new* "*engrafted*" variables *beyond* the two variables of retinal position.

We can therefore try to understand how a *special* connectivity *extended* to these new modules and layers can perform differential computations.

(i) It must certainly have a *very special* functional architecture.

(ii) But we must also know under what conditions a point-like functional architecture is able to implement a differential calculus if we add *new engrafted variables satisfying appropriate structural constraints*. It is a *mathematical* problem.

The hypothesis is therefore

Maybe point-like processors *can* implement an *alternative* formulation of differential calculus using "hidden derivatives" (Richard Montgomery) as new supplementary *independent* variables which can be implemented in point-like processors.

But these new "hidden" derivatives must satisfy very strong *constraints* in order to be interpretable as "true" derivatives (as in Hamiltonian mechanics where you introduce the momenta as new *independent* variables and force them to be dual to velocities using the symplectic 2-form).

For that, *the hardwired connectivity of the network must be extremely specific*.

43. Horizontal cortico-cortical connections

Let us come back to experiments and look at what are called *horizontal cortico-cortical connections* connecting neurons of different hypercolums.

The following image (also due to Bosking *et al.*) shows how a marker (biocytin) locally injected in a zone of about $100\mu m$ of the layer $2/3$ of $V1$ of a tupaya (tree shrew) diffuses along horizontal connections (black marks) in a selective, "patchy", anisotropic way.

Short-ranged diffusion is isotropic and corresponds to *intra*-hypercolumnar inhibitory connections.

On the contrary, long-ranged diffusion is highly anisotropic, and corresponds to excitatory *inter*-hypercolumnar connections

The injection site is upper-left in a green domain.
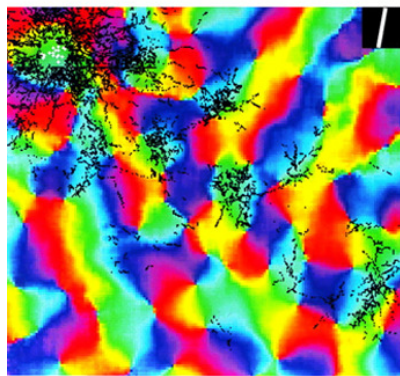
There are two main results:

1. The marked axons and synaptic buttons cluster in domains of the *same* color (same orientation) as the injection site, which means that horizontal connections connect neurons detecting approximatively parallel orientation and therefore implement neurally a *parallel transport*.

2. Furthermore, the striking global clustering along the top-left $\longrightarrow$ bottom-right diagonal means that horizontal connections connect neurons detecting not only almost *parallel* but also almost *aligned* "co-axial" orientations.
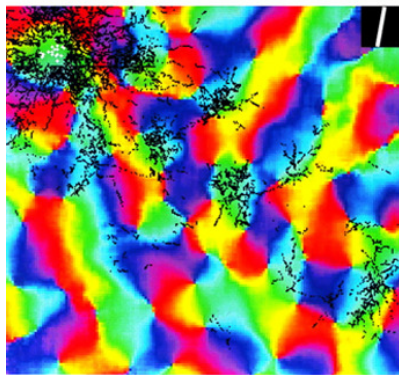
500 μm

0° 45° 90° 135° 180°

"The system of long-range horizontal connections can be summarized as preferentially linking neurons with co-oriented, co-axially aligned receptive fields." (W. Bosking)

We will see that this means that

the contact structure $\mathcal{C}$ of the fibration
$\mathbb{V} = \mathbb{R}^2 \times \mathbb{P}^1 \longrightarrow \mathbb{R}^2, (a, p) \mapsto a$ is neurally implemented in $V1$.

Let me underline that the hypothesis that the notion of *contact structure* must be involved in a natural way in visual perception was already explicitly formulated in 1989 by William Hoffman in his pioneering paper "The visual cortex is a contact bundle".

It was before "in vivo" optical imaging.

Neurophysiological results are corroborated by *psychophysical* experiments on curve integration.

In a breakthrough paper of 1993, David Field, Anthony Hayes and Robert Hess considered approximately aligned segments $v_i = (a_i, p_i)$ embedded in a background of randomly distributed segments lacking any global structure

However, subjects perceive very well a global alignment.

This striking phenomenon of *pop-out* (perceptual *saliency*) is due to a *low-level integration* processing.

After a lot of experimental measures, Field, Hayes and Hess concluded that the pop-out comes from a *specific connectivity*, which they called *association field*.

This connectivity is defined by what they called "joint conditions on positions and orientations":

1. two elements (a_1, p_1) and (a_2, p_2) are connected if one can *interpolate* between them a curve γ that is not too curved;

2. otherwise the two elements are not connected.

It has been shown later that the "joint constraints on positions and orientations" correspond to neural "horizontal" cortico-cortical connections.

51. Alternative differential calculus

Let us come now to the mathematical models.

Alternative versions of differential calculus *do exist* in "modern" differential geometry with Pfaff, Jacobi, Frobenius, Lie, Darboux, Cartan, Goursat, etc. They represent a great achievement of the "geometrization" of analysis during the XIXth century and the beginning of the XXth century.

They can solve the "hard problem" and it is why some structures they developed, such as jet-spaces, differential forms, contact structures, etc., are so basic in Neurogeometry,

The geometrization of the integrability conditions of differential equations has been a revolution similar to that accomplished by Galois for the resolution of algebraic equations.

A basic notion became that of a *contact element* (Lie's "Flächen-element"), that is a pair of a point x of a manifold M and a hyperplane K_x of the tangent space $T_x M$ of the manifold M at the point x.

If M is of dimension 2, K_x is an orientation p at x.

Sophus Lie developed this notion in great detail with a strong sense of novelty of this point of view.

53. Lie and Klein

In a beautiful text of 1894 "The Geometric Work of Sophus Lie", Felix Klein emphasized the importance of the change of perspective brought about with this "new, clear and penetrating view".

He explained that, inspired by Monge and Plücker, Lie introduced "new elements of space" *much more general than points* and, instead of applying the methods of differential analysis to geometry, was interested in "the reciprocal" and developed "the application of geometrical intuition to Analysis".

This is the key foundational issue: the foundations of differential calculus moved from Analysis to Geometry.

Now, a pair (a, p) of a retinal position a and an orientation p is a contact element. So, if it is possible to reformulate the differential calculus using contact elements, we can begin to understand how the visual brain can implement differential routines.

Hence the main hypothesis:

while the retinal cells detect positions, cortical neurons of $V1$ can detect contact elements and their cortico-cortical connectivity can implement a "geometry of integrability" for these contact elements.

To understand at the neurophysiological level how the visual system can implement an integro-differential calculus, we must therefore use the fact that, at the mathematical level, a differential calculus can be equivalent to a geometry of connectivity between contact elements.

It is why in the 1990s I began to develop the following relations between neurophysiology and geometry:

1. the "simple" cortical neurons of $V1$ detect contact elements (a, p);

2. the fibration $(a, p) \rightarrow a$ is neurophysiologically implemented by retinotopy and orientation hypercolumns;

3. the contact structure of the 1-jets of plane curves is neurophysiologically implemented (horizontal cortico-cortical connections);

4. the sub-Riemannian geometry of this contact structure is neurophysiologically implemented (illusory contours as geodesics).

Let us explain these concepts.

57. "Hidden derivatives"

The fibration $\mathbb{V} = \mathbb{R}^2 \times \mathbb{P}^1 \longrightarrow \mathbb{R}^2, (a, p) \mapsto a$ is not sufficient for interpreting p as an "hidden derivative" (Richard Montgomery), that is as a tangent to a curve.

For that, p must satisfy the fundamental Pfaff equation $\omega = dy - pdx = 0$ defining the contact structure of 1-jets of plane curves.

If we have a (regular) curve $\gamma = a(s)$ in the base plane $\mathbb{R}^2$ we can *lift* it in $\mathbb{V}$ as the skew curve

$$\Gamma = v(s) = (a(s), p(s)) = (x(s), y(s), p(s) = dy/dx)$$

The *point* $(x(s), y(s), p(s) = dy/dx)$ in $\mathbb{V}$ is called the *1-jet* of the curve at a and Γ is called the *Legendrian lift* of γ.

The notion of jet is very important because it is *a point coding a derivative* in a higher dimensional space if a certain condition is satisfied.

Adding a new dimension we can transform derivatives in "point-like" entities computable by point-like processors.

60. The contact structure of 1-jets

Conversely, a skew curve

$$\Gamma = v(s) = (a(s), p(s)) = (x(s), y(s), p(s))$$

in the $3D$ $\mathbb{V}$ is the *Legendrian lift* of its projection $\gamma = a(s)$ onto the base plane $\mathbb{R}^2$

- iff the "hidden" derivative $p(s)$ is the "real" derivative $p = dy/dx$ giving the tangent to the base curve γ at the point $a(s)$,

- iff it is an *integral curve* of the *contact structure* $\mathcal{C} = ker(\omega)$ of $\mathbb{V}$, where ω is the 1-differential form

$$\omega = dy - pdx$$

- Indeed, in the 3-space $\mathbb{V}$ of coordinates (x, y, p), $\omega = dy - pdx = 0$ means $p = dy/dx$ in the base plane (x, y).

- But $\omega = dy - pdx = 0$ defines a distribution of *planes* in $\mathbb{V}$. So the differential calculus of curves in the base plane $\mathbb{R}^2$ is translated into the geometry of a distibution of planes in the 3D-space $\mathbb{V}$.

62. The distribution of contact planes

The distribution $\mathcal{C}$ of contact tangent planes is maximally non integrable since the 3-form

$$\omega \wedge d\omega = (-pdx + dy) \wedge dx \wedge dp = -dx \wedge dy \wedge dp .$$

is a volume form, which is the opposite of the Frobenius integrability condition $\omega \wedge d\omega = 0$.

So, even if there exists a lot of integral *curves* of $\mathcal{C}$ (Legendrian lifts), there exists no integral *surface*.

63. The (polarized) Heisenberg group

It is interesting to note that the contact structure $\mathcal{C}$ is left-invariant for a group law making $\mathbb{V}$ isomorphic to the (*polarized*) *Heisenberg group* $\mathbb{H}_{pol}$. The product of 2 contact elements is given by :

$$(x, y, p).(x', y', p') = (x + x', y + y' + px', p + p') .$$

Its Lie algebra is generated by the basis of left-invariant fields $X_1 = \frac{\partial}{\partial x} + p\frac{\partial}{\partial y} = (1, p, 0)$ and $X_2 = \frac{\partial}{\partial p} = (0, 0, 1)$ with $[X_1, X_2] = (0, -1, 0) = -\frac{\partial}{\partial y} = -X_3$ (the other brackets $= 0$).

The basis $\{X_1, X_2\}$ of the distribution $\mathcal{C}$ is *bracket generating* (*i.e.* Lie-generates the whole tangent bundle $T\mathbb{V}$) (Hörmander condition).

$\mathbb{V} = \mathbb{H}_{pol}$ is a *nilpotent* group of step 2 (a Carnot group).

We can explain now how the contact structure of $\mathbb{V}$ is *implemented* in the *specific class* of cortico-cortical long-range horizontal connections.

Remember the key experimental discovery :

"the system of long-range horizontal connections can be summarized as preferentially linking neurons with co-oriented, co-axially aligned receptive fields."(W. Bosking)

This means that a chain of simple neurons (a_i, p_i) is a chain of *"horizontally" connected* simple neurons iff it is a discretization of the Legendrian lift of a not too curved base curve interpolating between the (a_i).

So, this means that,

up to some bound on curvature, the contact structure $\mathcal{C}$ is neurally implemented in $V1$.

66. Two vocabularies

So, we get a correspondance between two vocabularies, a neurophysiological one and a mathematical one.

| simple neurons	(scaled) contact elements (a, p)
R-G-C retinotopy	base space $\mathbb{R}^2$ of positions a
basic / "engrafted" variables	fiber bundle $\mathbb{R}^2 \times \mathbb{P} \to \mathbb{R}^2$
orientation hypercolumns and pinwheels	1-jet space $J^1 \subsetneq \mathbb{R}^2 \times \mathbb{P} \to \mathbb{R}^2$
◇ long-range horizontal connections, ◇ "co-oriented, co-axially aligned RFs", ◇ "joint constraints on positions and orientations" ◇ "good continuation"	Contact structure

With the tools of neurogeometry we can begin to understand how the very complex neurophysiology of visual neurons detecting *local* cues can generate the *global geometry* of the perceived images, with all the well-known phenomena studied since Gestalt theory, e.g. *long range illusory contours*.

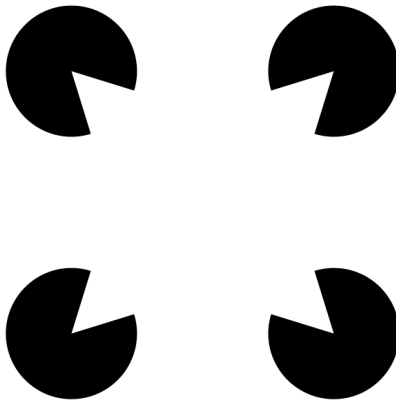Long range illusory contours are one of the most striking phenomenon of low level vision.

Consider for example the well-known Kanizsa triangle. *Local cues* as pacmen and end-points induce very long-range global illusory contours (what is called "modal completion").
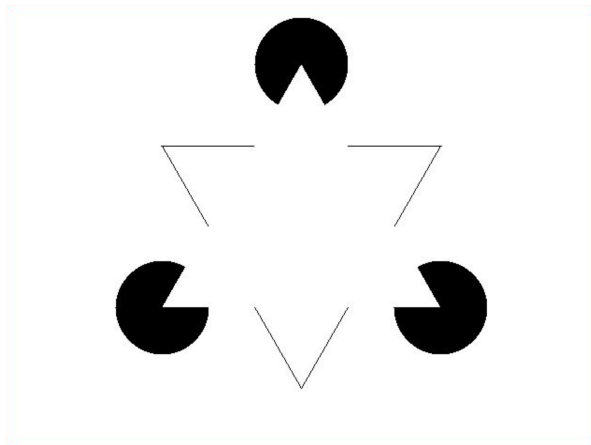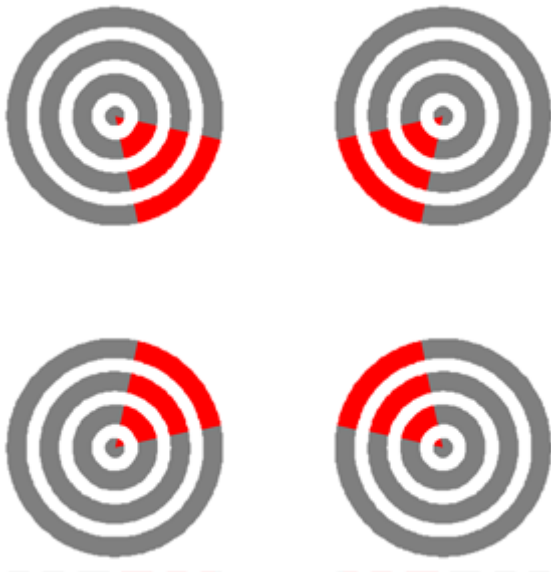
Illusory contours are particularly interesting when they are *curved*.

Furthermore, these contours act as *boundaries for a diffusion of color* inside the square (what is called the "neon" or "watercolor effect").

It is not easily seen on a screen but can be measured with adequate psychophysical methods.

We will see that illusory modal can be interpreted as *geodesics* of a contact structure for an appropriate *sub-Riemannian metric*.

Such sub-Riemannian models have many spectacular applications, in particular for *inpainting*, since to complete a corrupted image, we must construct the *illusory level sets* that can complete the missing parts.

The following picture shows how a highly corrupted image (left) can be very well restored using sub-Riemannian diffusion (Gauthier-Prandi inpainting based on our model).

The face of our colleague Jean-Paul Gauthier appears out of the blue.
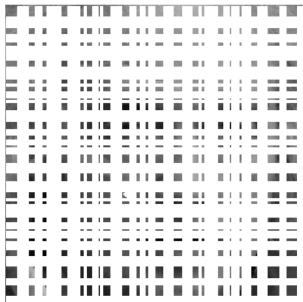
The idea is that of *variational models* for illusory contours. It has been introduced since the late 70s for $2D$ image processing.

1. Shimon Ullman (1976) explained that

   "A network with the local property of trying to keep the contours 'as straight as possible' can produce curves possessing the global property of minimizing total curvature."

2. Berthold Horn introduced in 1983 "the curves of least energy".

These models minimize an energy along curves *in the base plane*.

76. Mumford's elastica model. I

The best known is the *elastica* model proposed in 1992 by David Mumford.

The energy to minimize is:

$$E = \int_\gamma (\alpha\kappa^2 + \beta)ds$$

where γ is a smooth curve in $\mathbb{R}^2$.

But for *neural* models (and not only 2D image processing) it is natural to work in V1, that is with the *contact structure* and the *Legendrian lifts*.

It is here that sub-Riemannian geometry fully comes on stage.

The natural idea is to introduce sub-Riemannian metrics on the contact planes of $\mathbb{V}$ and look at geodesic models for curve completion and illusory modal contours.

We take the natural basis $\{X_1, X_2\}$ of the contact plane at the origin as an *orthonormal* basis and we *translate* it using left translations of the group structure (the polarized Heisenberg group $\mathbb{H}_{pol}$).

As the contact structure is left-invariant, we get that way a left-invariant metric on the contact planes.

As this metric is defined only on the contact planes and not on the complete tangent spaces of $\mathbb{V}$ it is called *sub-Riemannian*. But it enables to compute the length of the integral curves of the contact structure, that is of *Legendrian lifts* in the base plane.

In the 2000s these problems have been further explored with Alessandro Sarti and Giovanna Citti, and also Andrei Agrachev, Jean-Paul Gauthier, Ugo Boscain and Yuri Sachkov.

The sub-Riemannian geometry (geodesics, conjugate points, cut locus) of groups such as $\mathbb{H}_{pol}$ or $SE(2)$ is rather complex, even if the groups are elementary.

The sub-Riemannian geometry of the Heisenberg group $\mathbb{H}$ has been explained in the 1980s by Richard Beals, Bernard Gaveau and Peter Greiner.
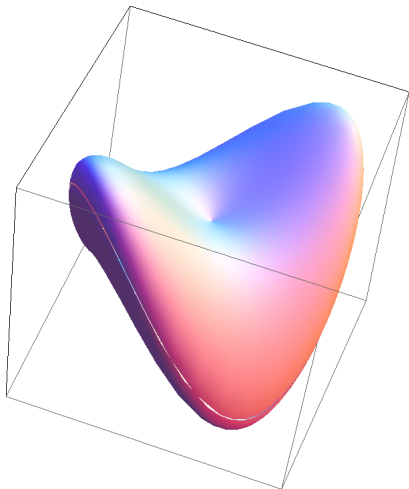
It can easily be adapted to the polarized $\mathbb{H}_{pol}$.

The sub-Riemannian sphere S and the wave front W are rather strange. One can compute them explicitly .

82. Image of the SR sphere of $\mathbb{H}$

They are displayed in the following figure. The external surface is the sub-Riemannian sphere S. It has a saddle form with singularities at the intersections with the y-axis.
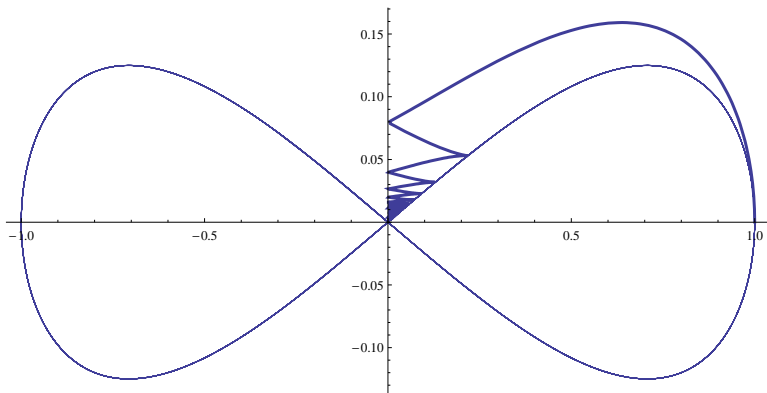
The internal part is $W - S$. It presents smaller and smaller circles of cusp singularities which converge to 0. Such a complex behavior is impossible in Riemannian geometry.
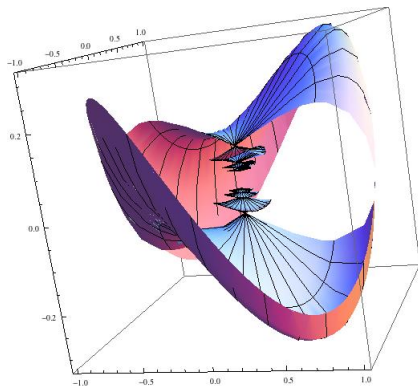
84. The cusps of W

The following figure displays the quarter of the wave front W for $\theta = 0$. Its equations are $x_1 = \frac{|\sin(\varphi)|}{\varphi}$, $p_1 = 0$, $y_1 = \frac{\varphi - \cos(\varphi)\sin(\varphi)}{4\varphi^2}$.
The cusps are on the curve of equation $x = \cos(\varphi)$,
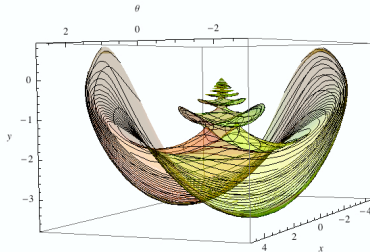$y = \frac{1}{4}\cos(\varphi)\sin(\varphi)$.

If you read French you will find a lot of mathematical details on all these structures in the second volume of my "Éléments de Neurogéométrie" :
hal.science/hal-04722501

Jean Petitot

ÉLÉMENTS DE NEUROGÉOMÉTRIE

Volume 2

Everything remains to be done in the field of neurogeometry. Firstly, to "go down" to the microphysical level of neurons and their spikes governed by Hodgkin-Huxley-type equations. Neurogeometry operates in fact at the meso neuronal level.

Then to go beyond $V1$ to higher retinotopic visual areas, from $V2$ to MT. We have indeed focused on $V1$.

But despite the very partial nature of the already worked out results, we hope to have shown how we can begin to understand the constitution of perceptual geometry from internal, "immanent" neurogeometric algorithms.