

RESEARCH

Open Access



A bio-inspired geometric model for sound reconstruction

Ugo Boscain^{1*} , Dario Prandi² , Ludovic Sacchelli³  and Giuseppina Turco⁴ 

*Correspondence:

ugo.boscain@upmc.fr

¹CNRS, LJLL, Sorbonne Université,
Université de Paris, Inria, Paris,
France

Full list of author information is
available at the end of the article

Abstract

The reconstruction mechanisms built by the human auditory system during sound reconstruction are still a matter of debate. The purpose of this study is to propose a mathematical model of sound reconstruction based on the functional architecture of the auditory cortex (A1). The model is inspired by the geometrical modelling of vision, which has undergone a great development in the last ten years. There are, however, fundamental dissimilarities, due to the different role played by time and the different group of symmetries. The algorithm transforms the degraded sound in an ‘image’ in the time–frequency domain via a short-time Fourier transform. Such an image is then lifted to the Heisenberg group and is reconstructed via a Wilson–Cowan integro-differential equation. Preliminary numerical experiments are provided, showing the good reconstruction properties of the algorithm on synthetic sounds concentrated around two frequencies.

1 Introduction

Listening to speech requires the capacity of the auditory system to map incoming sensory input to lexical representations. When the sound is intelligible, this mapping (‘recognition’) process is successful. With reduced intelligibility (e.g. due to background noise), the listener has to face the task of recovering the loss of acoustic information. This task is very complex as it requires a higher cognitive load and the ability of repairing missing input. (See [28] for a review on noise in speech.) Yet, (normal hearing) humans are quite able to recover sounds in several effortful listening situations (see, for instance, [27]), ranging from sounds degraded at the source (e.g. hypoarticulated and pathological speech), during transmission (e.g. reverberation) or corrupted by the presence of environmental noise.

So far, work on degraded speech has informed us a lot on the acoustic cues that help the listener to reconstruct missing information (e.g. [18, 31]); the several adverse conditions in which listeners may be able to reconstruct speech sounds (e.g. [2, 28]); and whether (and at which stage of the auditory process) higher-order knowledge (i.e. our information about words and sentences) helps the system to recover lower-level perceptual information (e.g. [22]). However, most of these studies adopt a phenomenological and descriptive approach. More specifically, techniques from previous studies consist in adding synthetic noise to speech sound stimuli, performing spectral and temporal analyses on the stimuli with noise and the same ones without it to identify acoustic differences, linking the results of these

© The Author(s) 2020. This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

analyses with the outcome from perceptual experiments. In some of these behavioural experiments, for instance, listeners are asked to identify speech units (such as consonants or words) when listening the noisy stimuli. Their accuracy scores provide a measure to the listeners' speech recognition ability.

As it stands, a mathematical model informing us on how the human auditory system is able to reconstruct a degraded speech sound is still missing. The aim of this study is to build a neuro-geometric model for sound reconstruction, stemming from the description of the functional architecture of the auditory cortex.

1.1 Modelling the auditory cortex

Knowledge about the functional architecture of the auditory cortex is scarce, and there are difficulties in the application of Gestalt principles for auditory perception. For these reasons, the model we propose is strongly inspired by recent advances in the mathematical modeling of the functional architecture of the primary visual cortex and the processing of visual inputs [9, 13, 24, 32], which recently yield very successful applications to image processing [10, 16, 20, 35]. This idea is not new: neuroscientists take models of V1 as a starting point for understanding the auditory system (see, e.g. [30] for a comparison, and [23] for a related discussion in speech processing). Indeed, biological similarities between the structure of the primary visual cortex (V1) and the primary auditory cortex (A1) are well-known to exist.

An often cited V1–A1 similarity is their 'topographic' organization, a general principle determining how visual and auditory inputs are mapped to those neurons responsible for their processing [38]. Substantial evidence for V1–A1 relation is also provided by studies on animals and on humans with deprived hearing or visual functions showing cross-talk interactions between sensory regions [41, 44]. More relevant for our study is the existence of receptive fields of neurons in V1 and A1 that allow for a subdivision of neurons in 'simple' and 'complex' cells, which supports the idea of a 'common canonical processing algorithm within cortical columns' [42, p. 1]. Together with the appearance in A1 of singularities typical of V1 (e.g. pinwheels) [34, 41], these findings speak in favour of the idea that V1 and A1 share similar mechanisms of sensory input reconstruction. In the next section we present the mathematical model for V1 that will be the basis for our sound reconstruction algorithm.

1.2 Neuro-geometric model of V1

The neuro-geometric model of V1 finds its roots in the experimental results of Hubel and Wiesel [25], which inspired Hoffman [24] to model V1 as a *contact space*.^a This model has then been extended to the so-called sub-Riemannian model in [8, 9, 13, 33]. On the basis of such a model, exceptionally efficient algorithms for image inpainting have been developed (e.g. [10, 15, 16]). These algorithms have now several medical imaging applications (e.g. [45]).

The main idea behind this model is that an image, seen as a function $f: \mathbb{R}^2 \rightarrow \mathbb{R}_+$ representing the grey level, is lifted to a distribution on $\mathbb{R}^2 \times P^1$, the bundle of directions of the plane.^b Here, P^1 is the projective line, i.e. $P^1 = \mathbb{R}/\pi\mathbb{Z}$. More precisely, the lift is given by $Lf(x, y, \theta) = \delta_{S_f}(x, y, \theta)f(x, y)$ where δ_{S_f} is the Dirac mass supported on the set $S_f \subset \mathbb{R}^2 \times P^1$ of points (x, y, θ) such that θ is the direction of the tangent line to f at (x, y) . Notice that, under suitable regularity assumptions on f , S_f is a surface.

When f is corrupted (i.e. when f is not defined in some region of the plane), the lift is corrupted as well, and the reconstruction is obtained by applying a deeply anisotropic diffusion adapted to the problem. Such diffusion mimics the flow of information along the horizontal and vertical connections of V1 and uses as an initial condition the surface S_f and the values of the function f . Mathematicians call such a diffusion the *sub-Riemannian diffusion* in $\mathbb{R}^2 \times P^1$, cf. [1, 29]. One of the main features of this diffusion is that it is invariant by rototranslation of the plane, a feature that will not be possible to translate to the case of sounds, due to the special role of the time variable.

In what follows, we explain how similar ideas could be translated to the problem of sound reconstruction.


1.3 From V1 to sound reconstruction

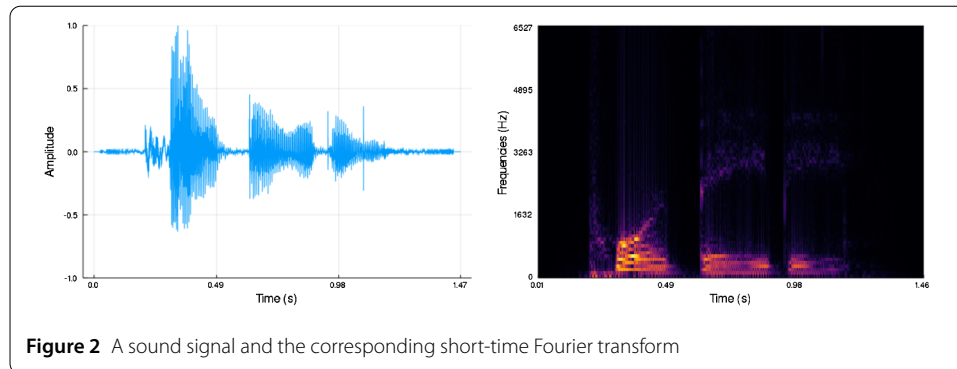
The sensory input reaching A1 comes directly from the cochlea [14]: a spiral-shaped, fluid-filled, cavity that composes the inner ear. Vibrations coming from the ossicles in the middle ear are transmitted to the cochlea, where they propagate and are picked up by sensors (so-called hair cells). These sensors are tonotopically organized along the spiral ganglion of the cochlea in a frequency-specific fashion, with cells close to the base of the ganglion being more sensitive to low-frequency sounds and cells near the apex more sensitive to high-frequency sounds, see Fig. 1. This early ‘spectrogram’ of the signal is then transmitted to higher-order layers of the auditory cortex.

Mathematically speaking, this means that when we hear a sound (that we can think as represented by a function $s : [0, T] \rightarrow \mathbb{R}$), our primary auditory cortex A1 is fed by its time–frequency representation^c $S : [0, T] \times \mathbb{R} \rightarrow \mathbb{C}$. If, say, $s \in L^2(\mathbb{R}^2)$, the time–frequency representation S is given by the short-time Fourier transform of s , defined as

$$S(\tau, \omega) := \text{STFT}(s)(\tau, \omega) = \int_{\mathbb{R}} s(t) W(\tau - t) e^{2\pi i t \omega} dt.$$

Here $W : \mathbb{R} \rightarrow [0, 1]$ is a compactly supported (smooth) window, so that $S \in L^2(\mathbb{R}^2)$. Since S is complex-valued, it can be thought as the collection of two black-and-white images, $|S|$ and $\arg S$. The function S depends on two variables, the first is time, that here we indicate with the letter τ , and the second is frequency, denoted by ω . Roughly speaking, $|S(\tau, \omega)|$





represents the strength of the presence of the frequency ω at time τ . In the following, we call S the sound image (see Fig. 2).

A first attempt to model the process of sound reconstruction into A1 is to apply the algorithm for image reconstruction described in Sect. 1.2. In a sound image, however, time plays a special role. Indeed,

1. While for images the reconstruction can be done by evolving the whole image simultaneously, the whole sound image does not reach the auditory cortex simultaneously, but sequentially. Hence, the reconstruction can be performed only in a sliding window.
2. A rotated sound image corresponds to a completely different input sound and thus the invariance by rototranslations is lost.

As a consequence, different symmetries have to be taken into account (see Appendix B) and a different model for both the lift and the processing in the lifted space is required.

In order to introduce this model, let us recall that, in V1, neural stimulation stems not only from the input but also from its variations. That is, mathematically speaking, the input image is considered as a real-valued function on a 2-dimensional space, and the orientation sensitivity arises from the sensitivity to a first order derivative information on this function, i.e. the tangent directions to level lines. This additional variational information allows lifting the 2-dimensional image space to the aforementioned contact space, and defining the sub-Riemannian diffusion [1, 11].

In our model of A1, we follow the same idea: we consider the variations of the input as additional variables. Input sound signals are time-dependent real-valued functions subjected to a short-time Fourier transform by the cochlea. As a result the A1 input is considered as a function of time and frequency. The first time derivative $v = d\omega/d\tau$ of this object, corresponding to the instantaneous chirpiness of the sound, allows adding a supplementary dimension to the domain of the input. As in the case of V1, this gives rise to a natural lift of the signal to an *augmented space*, which in this case turns out to be \mathbb{R}^3 with the Heisenberg group structure. (This structure very often appears in signal processing; see, for instance, [21] and Appendix B.)

As we already mentioned, the special role played by time in sound signals does not permit modeling the flow of information as a pure hypoelliptic diffusion, as was done for static images in V1. We thus turn to a different kind of model, namely Wilson–Cowan equations [43]. Such a model, based on an integro-differential equation, has been successfully applied to describe the evolution of neural activations. In particular, it allowed theoretically predicting complex perceptual phenomena in V1, such as the emergence of

hallucinatory patterns [12, 17], and has been used in various computational models of the auditory cortex [26, 37, 46]. Recently, these equations have been coupled with the neuro-geometric model of V1 to great benefit. For instance, in [4–6] they allowed replicating orientation-dependent brightness illusory phenomena, which had proved to be difficult to implement for non-cortical-inspired models. See also [39] for applications to the detection of perceptual units.

On top of these positive results, Wilson–Cowan equations present many advantages from the point of view of A1 modelling: (i) they can be applied independently of the underlying structure, which is only encoded in the kernel of the integral term; (ii) they allow for a natural implementation of delay terms in the interactions; and (iii) they can be easily tuned via few parameters with a clear effect on the results. On the basis of these positive results, we emulate this approach in the A1 context. Namely, we will consider the lifted sound image $I(\tau, \omega, \nu)$ to yield an A1 activation $a(\tau, \omega, \nu)$ via the following Wilson–Cowan equations:

$$\begin{aligned} \partial_t a(t, \omega, \nu) = & -\alpha a(t, \omega, \nu) + \beta I(t, \omega, \nu) \\ & + \gamma \int_{\mathbb{R}^2} k_\delta(\omega, \nu \| \omega', \nu') \sigma(a(t - \delta, \omega', \nu')) d\omega' d\nu'. \end{aligned} \quad (\text{WC})$$

Here (t, ω, ν) are coordinates on the augmented space corresponding to time, frequency, and chirpiness, respectively; $\alpha, \beta, \gamma > 0$ are parameters; $\sigma : \mathbb{C} \rightarrow \mathbb{C}$ is a non-linear sigmoid; $k_\delta(\omega, \nu \| \omega', \nu')$ is a weight modelling the interaction between (ω, ν) and (ω', ν') after a delay of $\delta > 0$. The presence of this delay term models the fact that the time-scale of the input signal and of the neuronal activation are comparable.

The proposed algorithm to process a sound signal $s : [0, T] \rightarrow \mathbb{R}$ is the following:

A. *Preprocessing*

- (a) Compute the time–frequency representation $S : [0, T] \times \mathbb{R} \rightarrow \mathbb{C}$ of s , via standard short-time Fourier transform (STFT);
- (b) Lift this representation to the Heisenberg group, which encodes redundant information about chirpiness, obtaining $I : [0, T] \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{C}$ (see Sect. 2.1 for details);

B. *Processing* Process the lifted representation I via Wilson–Cowan equations adapted to the Heisenberg structure, obtaining $a : [0, T] \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{C}$.

C. *Postprocessing* Project a to the processed time–frequency representation

$\hat{S} : [0, T] \times \mathbb{R} \rightarrow \mathbb{C}$ and then apply an inverse STFT to obtain the resulting sound signal $\hat{s} : [0, T] \rightarrow \mathbb{R}$.

Remark 1 All the above operations can be performed in real-time, as they only require the knowledge of the sound on a short window $[t - \delta, t + \delta]$.

Remark 2 Notice that in the presented algorithm we are assuming neural activations to be complex-valued functions, due to the use of the STFT. This is inconsistent with neural modelling, as it is known that the cochlea sends to A1 only the spectrogram of the STFT (that is, $|S|$), see [40]. When striving for a biologically plausible description, one can easily modify the above algorithm in this direction (i.e. by computing the lifted representation I starting from $|S|$ instead than S). However, during the post-processing phase, in order to

invert the STFT and obtain an audible signal, one then needs to reconstruct the missing phase information via heuristic algorithms. See, for instance, [19].

1.4 Structure of the paper

In Sect. 2, we present the reconstruction model. We first present the lift procedure of a sound signal to a function on the augmented space, and then introduce the Wilson–Cowan equations modelling the cortical stimulus. In Sect. 3, we describe the numerical implementation of the algorithm, together with some of its crucial properties. This implementation is then tested in Sect. 4, where we show the results of the algorithm on some simple synthetic signals. Such numerical examples can be listened at www.github.com/dprn/WCA1, and should be considered as a very preliminary step toward the construction of an efficient cortical-inspired algorithm for sound reconstruction. Finally, in Appendix B, we show how the proposed algorithm preserves the natural symmetries of sound signals.

2 The reconstruction model

As discussed in the introduction, the cochlea decomposes the input sound $s : [0, T] \rightarrow \mathbb{R}$ in its time–frequency representation $S : [0, T] \times \mathbb{R} \rightarrow \mathbb{C}$, obtained via a short-time Fourier transform (STFT). This corresponds to interpreting the ‘instantaneous sound’ at time $\tau \in [0, T]$, instead of as a sound level $s(\tau) \in \mathbb{R}$, as a function $\omega \mapsto S(\tau, \omega)$ which encodes the instantaneous presence of each given frequency, with phase information.

2.1 The lift to the augmented space

In this section, we present an extension of the time–frequency representation of a sound, which is at the core of the proposed algorithm. Roughly speaking, the instantaneous sound will be represented as a function $(\omega, \nu) \mapsto I(\tau, \omega, \nu)$, encoding the presence of both the frequency and the chirpiness $\nu = d\omega/d\tau$.

Assume for the moment that the sound has a single time-varying frequency, e.g.


$$s(\tau) = A \sin(\omega(\tau)\tau), \quad A \in \mathbb{R}. \quad (1)$$

If the frequency is varying slowly enough and the window of the STFT is large enough, its sound image (up to the choice of normalising constants in the Fourier transform) coincides roughly with

$$S(\tau, \omega) \sim \frac{A}{2i} (\delta_0(\omega - \omega(\tau)) - \delta_0(\omega + \omega(\tau))),$$

where δ_0 is the Dirac delta distribution centered at 0. That is, S is concentrated on the two curves $\tau \mapsto (\tau, \omega(\tau))$ and $\tau \mapsto (\tau, -\omega(\tau))$, see Fig. 3. Let us focus only on the first curve.

Figure 3 Short-time Fourier transform of the signal in (1), for a positive and increasing $\omega(\cdot)$



Because of the sensitivity to variations of the input, as discussed in Sect. 1, the curve $\omega(\tau)$ is lifted in a bigger space by adding a new variable $v = d\omega/d\tau$. In mathematical terms, the 3-dimensional space (τ, ω, v) is called the *augmented space*. It will be the basis for the geometric model of A1 that we are going to present.

Up to now the curve $\omega(\tau)$ was parameterized by one of the coordinates of the contact space (the variable τ), but it will be more convenient to consider it as a parametric curve in the space (τ, ω) . More precisely, the original curve $\omega(\tau)$ is represented in the space (τ, ω) as $t \mapsto (t, \omega(t))$ (thus imposing $\tau = t$). Similarly, the lifted curve is parameterized as $t \mapsto (t, \omega(t), v(t))$. To every regular enough curve $t \mapsto (t, \omega(t))$, one can associate a lift $t \mapsto (t, \omega(t), v(t))$ in the contact space simply by computing $v(t) = d\omega/dt$. Conversely, a regular enough curve in the contact space $t \mapsto (\tau(t), \omega(t), v(t))$ is a lift of planar curve $t \mapsto (t, \omega(t))$ if $\tau(t) = t$ and if $v(t) = d\omega/dt$. Now, defining $u(t) = dv/dt$, we can say that a curve in the contact space $t \mapsto (\tau(t), \omega(t), v(t))$ is a lift of a planar curve if there exists a function $u(t)$ such that

$$\frac{d}{dt} \begin{pmatrix} \tau \\ \omega \\ v \end{pmatrix} = \begin{pmatrix} 1 \\ v \\ 0 \end{pmatrix} + u(t) \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}. \quad (2)$$

Letting $q = (\tau, \omega, v)$, equation (2) can be equivalently written as the control system

$$\frac{d}{dt} q(t) = X_0(q(t)) + u(t)X_1(q(t)),$$

where the X_0 and X_1 are two vector fields in \mathbb{R}^3 given by

$$X_0 = \begin{pmatrix} 1 \\ v \\ 0 \end{pmatrix}, \quad X_1 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}.$$

Notice that the two vector fields appearing in this formula generate the Heisenberg group. However, we are not dealing here with a sub-Riemannian structure, since the space $\{X_0 + uX_1 \mid u \in \mathbb{R}\}$ is a line and not a plane. (One would get a plane by considering two controls, namely $\{u_0X_0 + u_1X_1 \mid (u_0, u_1) \in \mathbb{R}^2\}$.)

Following [9], when s is a general sound signal, we lift each level line of $|S|$. By the implicit function theorem, this yields the following subset of the contact space:

$$\Sigma = \{(\tau, \omega, v) \in \mathbb{R}^3 \mid v\partial_\omega |S|(\tau, \omega) + \partial_\tau |S|(\tau, \omega) = 0\}. \quad (3)$$

If $|S| \in C^2$ and $\text{Hess } |S|$ is non-degenerate, the set Σ is indeed a surface. Finally, the external input from the cochlea is given by

$$I(\tau, \omega, v) = S(\tau, \omega)\delta_\Sigma(\tau, \omega, v). \quad (4)$$

Here δ_Σ denotes the Dirac delta distribution concentrated at Σ . The presence of this distributional term is necessary for a well-defined solution to the evolution equation (WC). Such an equation is introduced in the next section.

2.2 Cortical activations in A1

On the basis of what described in the previous section and the well-known tonotopical organization of A1 (cf. Sect. 1), we propose to consider A1 to be the space of $(\omega, \nu) \in \mathbb{R}^2$. When hearing a sound $s(\cdot)$, the external input fed to A1 at time $t > 0$ is then given as the slice at $\tau = t$ of the lift I of s to the contact space. That is, hearing an ‘instantaneous sound level’ $s(t)$ reflects in the external input $I(t, \omega, \nu)$ to the ‘neuron’ (ω, ν) in A1 as follows: The ‘neuron’ receives an external charge $S(t, \omega)$ if $(t, \omega, \nu) \in \Sigma$, and no charge otherwise, where Σ is defined in (3).

We model the neuronal activation induced by the external stimulus I by adapting to this setting the well-known Wilson–Cowan equations. These equations are widely used and proved to be very effective in the study of V1 [12, 43]. According to this framework, the resulting activation $a : [0, T] \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{C}$ is the solution of the following equation with delay $\delta > 0$:

$$\begin{aligned} \partial_t a(t, \omega, \nu) = & -\alpha a(t, \omega, \nu) + \beta I(t, \omega, \nu) \\ & + \gamma \int_{\mathbb{R}^2} k_\delta(\omega, \nu \| \omega', \nu') \sigma(a(t - \delta, \omega', \nu')) d\omega' d\nu', \end{aligned} \quad (5)$$

with initial condition $a(t, \cdot, \cdot) \equiv 0$ for $t \leq 0$. Here $\alpha, \beta, \gamma > 0$ are parameters, k_δ is an interaction kernel, and $\sigma : \mathbb{C} \rightarrow \mathbb{C}$ is a (non-linear) saturation function, or sigmoid. In the following, we let $\sigma(\rho e^{i\theta}) = \tilde{\sigma}(\rho) e^{i\theta}$ where $\tilde{\sigma}(x) = \min\{1, \max\{0, \kappa x\}\}$, $x \in \mathbb{R}$, for some fixed $\kappa > 0$. The fact that the non-linearity σ does not act on the phase is one of the key ingredients in proving that this processing preserves the natural symmetries of sound signals, see Proposition 4 in Appendix B.

When $\gamma = 0$, equation (5) becomes the standard low-pass filter $\partial_t a = -\alpha a + I$, whose solution is the convolution of the input signal I with the function

$$\varphi(t) = \begin{cases} e^{-t\alpha} & \text{if } t > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Setting $\gamma \neq 0$ adds a non-linear delayed interaction term on top of this exponential smoothing, encoding the inhibitory and excitatory interconnections between neurons. Next section is devoted to the choice of the integral kernel k_δ .

Remark 3 In (5) we chose to consider a simple form for the interaction term. A more precise choice would indeed need to take into account the whole history of the process, for example, by considering

$$\int_{\tau}^{+\infty} e^{-\varrho(s-\tau)} \int_{\mathbb{R}^2} k_s(\omega, \nu \| \omega', \nu') \sigma(a(t-s, \omega', \nu')) d\omega' d\nu' ds, \quad \varrho > 0.$$

2.3 The neuronal interaction kernel

Considering A1 as a slice of the augmented space allows deducing a natural structure for neuron connections as follows. Going back to a sound composed by a single time-varying frequency $t \mapsto \omega(t)$, we have that its lift is concentrated on the curve $t \mapsto (\omega(t), \nu(t))$ such

that

$$\frac{d}{dt} \begin{pmatrix} \omega \\ v \end{pmatrix} = Y_0(\omega, v) + u(t)Y_1(\omega, v), \quad (6)$$

where $Y_0(\omega, v) = (v, 0)^\top$, $Y_1(\omega, v) = (0, 1)^\top$, and $u : [0, T] \rightarrow \mathbb{R}$.

As in the case of V1 [8], we model neuronal connections via these dynamics. In practice, this amounts to assuming that the excitation starting at a neuron $X_0 = (\omega', v')$ evolves as the stochastic process $\{A_t\}_{t \geq 0}$ naturally associated with (6). This is given by the following stochastic differential equation:

$$dA_t = Y_0(A_t)dt + Y_1(A_t)dW_t, \quad A_0 = (\omega', v'), \quad (7)$$

where $\{W_t\}_{t \geq 0}$ is a Wiener process. The generator of $\{A_t\}_{t \geq 0}$ is the second order differential operator

$$\mathcal{L} = Y_0 + (Y_1)^2 = v\partial_\omega + b\partial_v^2.$$

In this formula, the vector fields Y_0 and Y_1 are interpreted as first-order differential operators. Moreover, we added a scaling parameter $b > 0$, modelling the relative strength of the two terms.

It is natural to model the influence $k_\delta(\omega, v \parallel \omega', v')$ of neuron (ω', v') on neuron (ω, v) at time $\delta > 0$ as the transition density of the process $\{A_t\}_{t \geq 0}$. It is well-known that such transition density is obtained by computing the integral kernel at time δ of the Fokker–Planck equation corresponding to (7) that reads

$$\partial_t I = \mathcal{L}^* I, \quad \text{where } \mathcal{L}^* = -Y_0 + (Y_1)^2 = -v\partial_\omega + b\partial_v^2. \quad (8)$$

The existence of an integral kernel for (8) is a consequence of the hypoellipticity^d of $(\partial_t - \mathcal{L}^*)$. The explicit expression of k_δ is well-known, and we recall it in the following result, proved in Appendix A.

Proposition 1 *The integral kernel of equation (8) is*

$$k_\delta(\omega, v \parallel \omega', v') = \frac{\sqrt{3}}{2\pi b\delta^2} \exp\left(-\frac{g_\delta(\omega, v \parallel \omega', v')}{b\delta^3}\right), \quad (9)$$

where

$$g_\delta(\omega, v \parallel \omega', v') = 3(\omega - \omega')^2 - 3\delta(\omega - \omega')(v + v') + \delta^2(v^2 + vv' + v'^2).$$

3 Numerical implementation

For the numerical experiments, we chose to implement the proposed algorithm in Julia [7]. As already presented, this process consists in a pre-processing phase, in which we build an input function I on the 3D contact space, a main part, where I is used as the input of the Wilson–Cowan equation (WC), and a post-processing phase, where the reconstructed sound is recovered from the result of the first part.

In the following, we present these phases separately.

3.1 Pre-processing

The input sound s is lifted to a time–frequency representation S via a classical implementation of STFT, i.e. by performing FFTs of a windowed discretised input. In the proposed implementation, we chose to use a standard Hann window (see, e.g. [36])

$$W(x) = \begin{cases} \frac{1+\cos(2\pi x/L)}{2} & \text{if } |x| < L/2, \\ 0 & \text{otherwise.} \end{cases}$$

The resulting time–frequency signal is then lifted to the contact space through an approximate computation of the gradient $\nabla|S|$ and the following discretisation of (4):

$$I(\tau, \omega, \nu) = \begin{cases} S(\tau, \omega) & \text{if } \nu \partial_\omega |S|(\tau, \omega) = -\partial_\tau |S|(\tau, \omega), \\ 0 & \text{otherwise.} \end{cases}$$

Discretisation issues While the discretisation of the time and frequency domains is a well-understood problem, dealing with the additional chirpiness variable requires some care. Indeed, even if we assume that the significant frequencies of the input sound s belong to a bounded interval $\Lambda \subset \mathbb{R}$, in general the set $\{\nu \in \mathbb{R} \mid I(\tau, \omega, \nu) \neq 0\}$ is unbounded. Indeed, one can check that as (τ, ω) moves to a point where the countour lines of $|S|$ become vertical, the set of chirpinesses ν 's such that $\nu \partial_\omega |S|(\tau, \omega) = -\partial_\tau |S|(\tau, \omega)$ will converge to $\pm\infty$.

In the numerical implementation, we chose to restrict the admissible chirpinesses to a bounded interval $N \subset \mathbb{R}$. This set is chosen in a case by case fashion in order to contain the relevant slopes for the examples under consideration. Work is ongoing to automate this procedure.

3.2 Processing

Equation (WC) can be solved via a standard forward Euler method. Hence, the delicate part of the numerical implementation is the computation of the interaction term.

As is clear from the explicit expression given in Proposition 1, k_δ is not a convolution kernel. That is, $k_\delta(\omega, \nu \parallel \omega', \nu')$ cannot be expressed as a function of $(\omega - \omega', \nu - \nu')$. As a consequence, a priori we need to explicitly compute all values $k_\delta(\omega, \nu \parallel \omega', \nu')$ for (ω, ν) and (ω', ν') in the considered domain. As is customary, in order to reduce computation times, we fix a threshold $\varepsilon > 0$ and for any given (ω, ν) we compute only values for (ω', ν') in the compact set

$$K_\delta^\varepsilon(\omega, \nu) = \{(\omega', \nu') \mid k_\delta(\omega, \nu \parallel \omega', \nu') \geq \varepsilon\}.$$

The structure of $K_\delta^\varepsilon(\omega, \nu)$ is given in the following, whose proof we defer to Appendix A.

Proposition 2 *For any $\varepsilon > 0$ and $(\omega, \nu) \in \mathbb{R}^2$, we have that $K_\delta^\varepsilon(\omega, \nu)$ is the set of those $(\omega', \nu') \in \mathbb{R}^2$ that satisfy*

$$| \nu - \nu' |^2 \leq C_\varepsilon := -4b\delta \log\left(\frac{2\pi b\tau^2}{\sqrt{3}}\varepsilon\right),$$

$$\left| \omega' - \omega + \frac{\delta(\nu + \nu')}{2} \right| \leq \frac{\delta}{2\sqrt{3}} \sqrt{C_\varepsilon - |\nu - \nu'|^2}.$$

Remark 4 One has $C_\varepsilon \geq 0$ if and only if

$$\varepsilon \leq \frac{\sqrt{3}}{2\pi b\delta^2}.$$

Indeed, for any $(\omega, \nu) \in \mathbb{R}^2$, the right-hand side above corresponds to $\max k_\delta(\omega, \nu \|\cdot, \cdot)$, and thus $K^\varepsilon(\omega, \nu) = \emptyset$ for larger values of ε .

The above allows numerically implementing k_δ as a family of sparse arrays. That is, let $G \subset \Lambda \times N$ be the chosen discretisation of the significant set of frequencies and chirpinesses. Then to $\xi = (\omega, \nu) \in G$ we associate the array $M_\xi : G \rightarrow \mathbb{R}$ defined by

$$M_\xi(\xi') = \begin{cases} k_\delta(\xi \|\xi') & \text{if } \xi' \in K^\varepsilon(\xi), \\ 0 & \text{otherwise.} \end{cases}$$

Therefore, up to choosing the tolerance $\varepsilon \ll 1$ sufficiently small, the interaction term in (WC), evaluated at $\xi = (\omega, \nu) \in G$, can be efficiently estimated by

$$\int_{\mathbb{R}^2} w(\xi \|\xi') \sigma(a(t - \delta, \xi')) d\xi' \approx \sum_{\xi' \in K^\varepsilon(\xi)} M_\xi(\xi') a(t - \delta, \xi').$$

3.3 Post-processing

Both operations in the pre-processing phase are invertible: the STFT by inverse STFT, and the lift by integration along the ν variable (that is, summation of the discretized solution). The final output signal is thus obtained by applying the inverse of the pre-processing (integration then inverse STFT) to the solution a of (WC). That is, the resulting signal is given by

$$\hat{s}(t) = \text{STFT}^{-1} \left(\int_{-\infty}^{+\infty} a(t, \omega, \nu) d\nu \right).$$

The following guarantees that \hat{s} is real-valued and thus correctly represents a sound signal. From the numerical point of view, this implies that we can focus on solutions of (WC) in the half-space $\{\omega \geq 0\}$, which can then be extended to the whole space by mirror symmetry.

Proposition 3 *It holds that $\hat{s}(t) \in \mathbb{R}$ for all $t > 0$.*

Proof Let us denote

$$\hat{S}(t, \omega) = \int_{-\infty}^{+\infty} a(t, \omega, \nu) d\nu,$$

so that $\hat{s} = \text{STFT}^{-1}(\hat{S})$. Moreover, for any function $f(t, \omega, \nu)$, we let $f^*(t, \omega, \nu) := \bar{f}(t, -\omega, -\nu)$.

To prove the statement, it is enough to show that

$$a(t, \cdot, \cdot) \equiv a^*(t, \cdot, \cdot) \quad \forall t \geq 0. \quad (10)$$

This is trivially satisfied for $t \leq 0$, since in this case $a(t, \cdot, \cdot) \equiv 0$.

We now claim that if (10) holds on $[0, T]$ it holds on $[0, T + \delta]$, which will prove it for all $t \geq 0$. By definition of I and the fact that $S(t, -\omega) = \overline{S(t, \omega)}$, we immediately have $I \equiv I^*$. On the other hand, the explicit expression of k_δ in (9) yields that

$$k_\delta(-\omega, -v \parallel \omega', v') = k_\tau(\omega, v \parallel -\omega', -v').$$


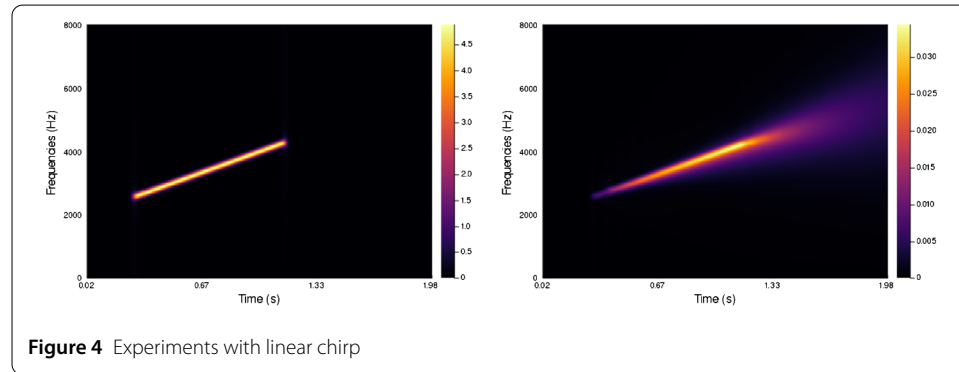
Then, for all $t \leq T + \delta$, we have


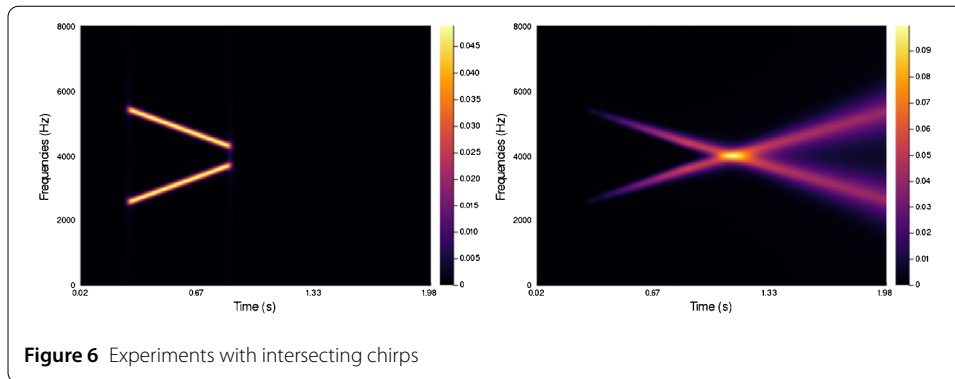
$$\begin{aligned} & \int_{\mathbb{R}^2} k_\delta(-\omega, -v \parallel \omega', v') \sigma(a(t - \tau, \omega', v')) d\omega' dv' \\ &= \int_{\mathbb{R}^2} k_\delta(\omega, v \parallel \omega'', v'') \sigma(a^*(t - \tau, \omega'', v'')) d\omega'' dv'' \\ &= \int_{\mathbb{R}^2} k_\delta(\omega, v \parallel \omega'', v'') \sigma(a(t - \tau, \omega'', v'')) d\omega'' dv''. \end{aligned}$$

A simple argument, e.g. using the variation of constants method, shows that these two facts imply the claim, and thus the statement. \square

4 Experiments

In Figs. 4–7 we present a series of experiments on simple synthetic sounds in order to exhibit some key features of our algorithm. These experiments can be reproduced via the code available at <https://www.github.com/dprn/WCA1>. For all experiments, the chosen delay is $\delta = 0.0625$ s and we present the STFT of the original and the processed sound. Each time, only the positive frequencies are shown: negative frequencies are recovered via the Hermitian symmetry of the Fourier transform on real signals.





The first example, Fig. 4, is a simple linear chirp such that the dominating frequency depends linearly on time (i.e. corresponding to $\omega(t) = \mu t$ for some $\mu \in \mathbb{R}$). One observes that the processed sound presents the same feature but for a longer duration. The parameters in the experiment ($\alpha = 55$, $\beta = 1$, $\gamma = 55$, $b = 0.05$) have been chosen to emphasize the effect of the modelling equation: the reconstruction should not present a tail that is as pronounced, however, this allows highlighting the diffusive effect along the lifted slope.

The second example, Fig. 5, corresponds to the same linear chirp as Fig. 4, that has been interrupted in its middle section, creating two disjoint linear chirps. The parameters are the same as in the previous experiment. Thanks to the transport effect of the algorithm, the gap between the two chirps is bridged in the processed signal. For this illustration, the interruption lasts about twice as long as the delay.

The third example, Fig. 6, consists of the sum of two linear chirps with different slopes. The slopes have been picked to suggest that linear continuations of the chirps should intersect. This is indeed what happens in the processed signal with parameters $\alpha = 53$, $\beta = 1$, $\gamma = 55$, $b = 0.01$. However, notice that the resulting crossing happens almost as a sum of the two chirps processed independently, with close to no interaction at the crossing. This is purely an effect of the lift procedure. The increasing chirp is (predominantly) lifted to a stratum corresponding to a positive slope, while the decreasing chirp is lifted to a negative slope stratum. De facto, their evolution under the Wilson–Cowan equation is decoupled in the 3D augmented space.

The fourth and last example, Fig. 7, corresponds to a non-linear chirp, roughly corresponding to choosing $\omega(\tau) = \sin(m\tau)$ in (1). The chosen parameters are $\alpha = 53$, $\beta = 1$, $\gamma = 55$, $b = 0.2$. The construction of the model favors linearity in the evolution of per-

ceived frequencies. We can observe how the more linear elements of the input result in more diffusion.

5 Conclusion

In this work we presented a sound reconstruction framework inspired by the analogies between visual and auditory cortices. Building upon the successful cortical inspired image reconstruction algorithms, the proposed framework lifts time–frequency representations of signals to the 3D contact space, by adding instantaneous chirpiness information. These redundant representations are then processed via adapted integro-differential Wilson–Cowan equations.

The promising results obtained on simple synthetic sounds, although preliminary, suggest possible applications of this framework to the problem of degraded speech. The next step will be to test the reconstruction ability of normal-hearing humans on originally degraded speech material compared to the same speech material after algorithm reconstruction. Such an endeavour will contribute to the understanding of the auditory mechanisms emerging in adverse listening conditions. It will furthermore help to deepen our knowledge on general organization principles underlying the functioning of the human auditory cortex.

Appendix A: Integral kernel of the Kolmogorov operator

The result in Proposition 1 is well-known. For example, by applying [3, Proposition 9] and letting $x = (\omega, \nu)$ and $x' = (\omega', \nu')$, one gets that the kernel is

$$k_\delta(\omega, \nu \| \omega', \nu') = \frac{1}{2\pi \sqrt{\det D_\delta}} \exp \left[-\frac{1}{2} (x' - e^{\delta A} x)^\top D_\delta^{-1} (x' - e^{\delta A} x) \right],$$

where

$$A = \begin{pmatrix} 0 & -1 \\ 0 & 0 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 0 \\ \sqrt{2b} \end{pmatrix},$$

and

$$D_\delta = e^{\delta A} \left[\int_0^\tau e^{-\sigma A} B B^* e^{-\sigma A^*} d\sigma \right] e^{\tau A^*}.$$

Direct computations yield

$$D_\delta = 2b \begin{pmatrix} \delta^3/3 & -\tau^2/2 \\ -\delta^2/2 & \tau \end{pmatrix} \quad \text{and} \quad \det D_\delta = \frac{b^2 \tau^4}{3}.$$

Therefore,

$$\frac{1}{2} D_\delta^{-1} = \frac{1}{b\tau^3} M, \quad \text{where } M = \begin{pmatrix} 3 & 3\delta/2 \\ 3\delta/2 & \delta^2 \end{pmatrix}.$$

Finally, the statement follows by letting

$$g_\delta(x \| x') = (x' - e^{\delta A} x)^\top M (x' - e^{\delta A} x).$$

We now turn to an argument for Proposition 2. Observe that $k_\delta(x\|x') \geq \varepsilon$ if and only if

$$g_\delta(x\|x') \leq \eta := -b\delta^3 \log\left(\frac{2\pi b\delta^2}{\sqrt{3}}\varepsilon\right).$$

Then, we start by solving $z^\top Mz \leq \eta$, for $z \in \mathbb{R}^2$. One can check that this is verified if and only if

$$|z_2| \leq \sqrt{\frac{4\eta}{\delta^2}} \quad \text{and} \quad \left|z_1 + \frac{\delta}{2}z_2\right| \leq \frac{\delta}{2\sqrt{3}}\sqrt{\frac{4\eta}{\delta^2} - z_2^2}.$$

Since $C_\varepsilon = 4\eta/\delta^2$, the statement follows by computing the above at $z = x' - e^{\tau A}x$.

Appendix B: Heisenberg group action on the contact space

Recall that the short-time Fourier transform of a signal $s \in L^2(\mathbb{R})$ is given by

$$S(\tau, \omega) := \text{STFT}(s)(\tau, \omega) = \int_{\mathbb{R}} s(t)W(\tau - t)e^{2\pi i t \omega} dt.$$

Here $W : \mathbb{R} \rightarrow [0, 1]$ is a compactly supported (smooth) window, so that $S \in L^2(\mathbb{R}^2)$. Fundamental operators in time–frequency analysis [21] are time and phase shifts, acting on signals $s \in L^2(\mathbb{R})$ by

$$T_\theta s(t) := s(t - \theta) \quad \text{and} \quad M_\lambda s(t) := e^{2\pi i \lambda t} s(t),$$

for $\theta, \lambda \in \mathbb{R}$. One easily checks that T_θ and M_λ are unitary operators on $L^2(\mathbb{R})$. By conjugation with the short-time Fourier transform, they naturally define the unitary operators on $L^2(\mathbb{R}^2)$ given by

$$\begin{aligned} T_\theta S(\tau, \omega) &= e^{-2\pi i \omega \theta} S(\tau - \theta, \omega), \\ M_\lambda S(\tau, \omega) &= S(\tau, \omega - \lambda). \end{aligned} \tag{11}$$

The relevance of the Heisenberg group in time–frequency analysis is a direct consequence of the commutation relation

$$T_\theta M_\lambda T_\theta^{-1} M_\lambda^{-1} = e^{-2\pi i \lambda \theta} \text{Id}.$$

Indeed, this shows that the operator algebra generated by $(T_\theta)_{\theta \in \mathbb{R}}$ and $(M_\lambda)_{\lambda \in \mathbb{R}}$ coincides with the Heisenberg group \mathbb{H}^1 via the representation $U : \mathbb{H}^1 \rightarrow \mathcal{U}(L^2(\mathbb{R}^2))$ defined by

$$U(\theta, \lambda, \zeta) = e^{-2\pi i \zeta} T_\theta M_\lambda. \tag{12}$$

The above discussion shows that the Heisenberg group can be regarded as the natural space of symmetries of sound signals. In particular, any meaningful treatment of these signals should respect such a symmetry. In the case of our model, this is the content of the following result.

Proposition 4 *The sound processing algorithm presented in this paper commutes with the Heisenberg group action (12) on sound signals. That is, if the input sound signal $s \in L^2(\mathbb{R}^2)$ yields \hat{s} as a result, then, for any $(\theta, \lambda, \zeta) \in \mathbb{H}^1$, the input $U(\theta, \lambda, \zeta)s$ yields $U(\theta, \lambda, \zeta)\hat{s}$ as a result.*

Proof We can schematically write the algorithm as:

$$\hat{s} = \text{STFT}^{-1} \circ \text{Proj} \circ \text{WC} \circ \text{Lift} \circ \text{STFT}(s).$$

Here Lift is the lift operator defined in Sect. 2.1, WC denotes the Wilson–Cowan evolution (WC), and Proj denotes the projection from the augmented space to the time–frequency representation, defined by

$$\text{Proj } a(t, \omega) = \int_{-\infty}^{+\infty} a(t, \omega, \nu) d\nu.$$

Observe that (11) shows that U induces a representation of \mathbb{H}^1 on $L^2(\mathbb{R}^2)$, the codomain of the STFT, which we will denote by \tilde{U} . Thus, to prove the statement it suffices to show that

$$[\text{Proj} \circ \text{WC} \circ \text{Lift}, \tilde{U}(\theta, \lambda, \zeta)] = 0, \quad \forall (\theta, \lambda, \zeta) \in \mathbb{H}^1.$$

Recall now that Lift associates with $S \in L^2(\mathbb{R}^2)$ a distribution of the form $\text{Lift}[S](\tau, \omega, \nu) = S(\omega, \nu)\delta_\Sigma(\tau, \omega, \nu)$ for some $\Sigma \subset \mathbb{R}^3$. Due to the fact that Σ is defined via the modulus of S , it is unaffected by the phase factors appearing in the representation \tilde{U} . That is, the lift of $\tilde{U}(\theta, \lambda, \zeta)S$ is given by

$$\begin{aligned} \text{Lift}[\tilde{U}(\theta, \lambda, \zeta)S](\tau, \omega, \nu) &= e^{-2\pi i(\zeta + \omega\theta)} \delta_\Sigma(\tau - \theta, \omega - \lambda, \nu) S(\tau - \theta, \omega - \lambda) \\ &= e^{-2\pi i(\zeta + \omega\theta)} \text{Lift}[S](\tau - \theta, \omega - \lambda, \nu). \end{aligned}$$

It is then immediate to check that $[\text{Proj} \circ \text{Lift}, \tilde{U}(\theta, \lambda, \zeta)] = 0$.

We are left to verify that the operator WC commutes with $\tilde{U}(\theta, \lambda, \zeta)$. The commutation is trivial for the linear terms. On the other hand, the non-linearity introduced in the integral term commutes with $\tilde{U}(\theta, \lambda, \zeta)$ thanks to the fact that $\sigma(\rho e^{i\phi}) = e^{i\phi} \sigma(\rho)$ for all $\rho > 0$, $\phi \in \mathbb{R}$. \square

Acknowledgements

The authors thank Jean-Paul Gauthier for stimulating discussions.

Funding

The authors acknowledge the support of the ANR project SRGI ANR-15-CE40-0018 and of the ANR project Quaco ANR-17-CE40-0007-01. This study was also supported by the IdEx Université de Paris, ANR-18-IDEX-0001, awarded to the last author, and by a public grant overseen by the French National Research Agency (ANR) as part of the program “Investissements d’Avenir” (reference: ANR-10-LABX-0083).

Abbreviations

Not applicable.

Availability of data and materials

Not applicable.

Ethics approval and consent to participate

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Consent for publication

Not applicable.

Authors' contributions

All authors have contributed equally. All authors read and approved the final manuscript.

Author details

¹CNRS, LJLL, Sorbonne Université, Université de Paris, Inria, Paris, France. ²Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des signaux et systèmes, 91190, Gif-sur-Yvette, France. ³Université Lyon, Université Claude Bernard Lyon 1, CNRS, LAGEPP UMR 5007, 43 bd du 11 novembre 1918, F-69100 Villeurbanne, France. ⁴CNRS, Laboratoire de Linguistique Formelle, UMR 7110, Université de Paris, Paris, France.

Endnotes

^a A 3-dimensional manifold M becomes a *contact space* once it is endowed with a smooth map $M \ni q \mapsto \mathcal{D}(q)$ where $\mathcal{D}(q)$ is a plane in the tangent space $T_q M$ passing from q . There is an additional requirement on this map. Locally one can always write $\mathcal{D}(q) = \text{span}\{X_0(q), X_1(q)\}$, where X_0 and X_1 are two smooth vector fields. Then at every point q one should require $\dim(\text{span}_q\{X_0, X_1, [X_0, X_1]\}) = 3$. Here $[\cdot, \cdot]$ is the Lie bracket of the vector fields. The main consequence of this condition is that no surface can be tangent to \mathcal{D} at all points.

By assigning to every $\mathcal{D}(q)$ an inner (Euclidean) product that is smooth as a function of q , we endow M with a *sub-Riemannian structure*. The simplest way of defining locally such a structure on a 3-dimensional manifold is to assign two vector fields X_0 and X_1 postulating, on the one hand, that $\mathcal{D}(q) = \text{span}\{X_0(q), X_1(q)\}$ (assigning in this way the contact structure) and, on the other hand, that they have norm one and are mutually orthogonal (assigning in this way the inner product).

The simplest example of sub-Riemannian structure on \mathbb{R}^3 is given by the so-called Heisenberg group for which the vector fields $X_0 = (1, \nu, 0)^T$ and $X_1 = (0, 0, 1)^T$ are orthonormal (here we write coordinates in \mathbb{R}^3 as (τ, ω, ν)). Such a structure is called Heisenberg group since defining $X_2 = (0, 1, 0)^T$ one has the Lie brackets $[X_0, X_1] = X_2$, $[X_0, X_2] = [X_1, X_2] = 0$, that are the commutation relations appearing in quantum mechanics.

^b Note that in mathematics, the term 'direction' corresponds to what neurophysiologists call 'orientation', and vice versa. In this study, we use the mathematical terminology.

^c Actually, its spectrogram $|S| : [0, T] \times \mathbb{R} \rightarrow [0, +\infty)$, see Remark 2.

^d That is, if f is a distribution defined on an open set Ω and such that $(\partial_t - \mathcal{L}^*)f \in C^\infty(\Omega)$, then $f \in C^\infty(\Omega)$.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 17 May 2020 Accepted: 8 December 2020 Published online: 04 January 2021

References

1. Agrachev A, Barilari D, Boscain U. A Comprehensive Introduction to Sub-Riemannian Geometry. Cambridge Studies in Advanced Mathematics. Cambridge: Cambridge University Press; 2020.
2. Assmann P, Summerfield Q. The Perception of Speech Under Adverse Conditions. New York: Springer; 2004. p. 231–308.
3. Barilari D, Boarotto F. Kolmogorov–Fokker–Planck operators in dimension two: heat kernel and curvature. SIAM J Control Optim. 2018.
4. Bertalmio M, Calatroni L, Franceschi V, Franceschiello B, Gomez Villa A, Prandi D. Visual illusions via neural dynamics: Wilson–Cowan-type models and the efficient representation principle. J Neurophysiol. 2020;PMID:32159409.
5. Bertalmio M, Calatroni L, Franceschi V, Franceschiello B, Prandi D. A cortical-inspired model for orientation-dependent contrast perception: A link with Wilson–Cowan equations. In: Scale Space and Variational Methods in Computer Vision. Cham: Springer; 2019.
6. Bertalmio M, Calatroni L, Franceschi V, Franceschiello B, Prandi D. Cortical-inspired Wilson–Cowan-type equations for orientation-dependent contrast perception modelling. J Math Imaging Vis. 2020.
7. Bezanson J, Edelman A, Karpinski S, Shah VB. Julia: A fresh approach to numerical computing. SIAM Rev. 2017;59(1):65–98.
8. Boscain U, Chertovskih R, Gauthier J-P, Remizov A. Hypoelliptic diffusion and human vision: a semi-discrete new twist on the Petitot theory. SIAM J Imaging Sci. 2014;7(2):669–95.
9. Boscain U, Duplaix J, Gauthier J-P, Rossi F. Anthropomorphic image reconstruction via hypoelliptic diffusion. 2010.
10. Boscain UV, Chertovskih R, Gauthier J-P, Prandi D, Remizov A. Highly corrupted image inpainting through hypoelliptic diffusion. J Math Imaging Vis. 2018;60(8):1231–45.
11. Bramanti M. An invitation to hypoelliptic operators and Hörmander's vector fields. SpringerBriefs in Mathematics. Cham: Springer; 2014.
12. Bressloff PC, Cowan JD, Golubitsky M, Thomas PJ, Wiener MC. Geometric visual hallucinations, Euclidean symmetry and the functional architecture of striate cortex. Philos Trans R Soc Lond B, Biol Sci. 2001;356(1407):299–330.
13. Citti G, Sarti A. A Cortical Based Model of Perceptual Completion in the Roto-Translation Space. J Math Imaging Vis. 2006;24(3):307–26.

14. Dallos P. Overview: Cochlear Neurobiology. New York: Springer; 1996. p. 1–43.
15. Duits R, Franken E. Left-invariant parabolic Evolutions on $SE(2)$ and Contour Enhancement via Invertible Orientation Scores. Part I: Linear Left-invariant Diffusion Equations on SE . *Q Appl Math*. 2010;68(2):255–92.
16. Duits R, Franken E. Left-invariant parabolic evolutions on $SE(2)$ and contour enhancement via invertible orientation scores. Part II: nonlinear left-invariant diffusions on invertible orientation scores. *Q Appl Math*. 2010;68(2):293–331.
17. Ermentrout GB, Cowan JD. A mathematical theory of visual hallucination patterns. *Biol Cybern*. 1979;34:137–50.
18. Fernandes T, Ventura P, Kolinsky R. Statistical information and coarticulation as cues to word boundaries: A matter of signal quality. *Percept Psychophys*. 2007;69(6):856–64.
19. Fienup J. Phase retrieval algorithms: a comparison. *Appl Opt*. 1982;21(15):2758–69.
20. Franken E, Duits R. Crossing-Preserving Coherence-Enhancing Diffusion on Invertible Orientation Scores. *Int J Comput Vis*. 2009;85(3):253–78.
21. Gröchenig K. Foundations of time-frequency analysis. Applied and Numerical Harmonic Analysis. Boston: Birkhäuser Boston; 2001.
22. Hannemann R, Obleser J, Eulitz C. Top-down knowledge supports the retrieval of lexical information from degraded speech. *Brain Res*. 2007;1153:134–43.
23. Hickok G, Poeppel D. The cortical organization of speech processing. *Nat Rev Neurosci*. 2007;8(5):393–402.
24. Hoffman WC. The visual cortex is a contact bundle. *Appl Math Comput*. 1989;32(2–3):137–67.
25. Hubel DH, Wiesel TN. Receptive fields of single neurons in the cat's striate cortex. *J Physiol*. 1959;148(3):574–91.
26. Loebel A, Nelken I, Tsodyks M. Processing of Sounds by Population Spikes in a Model of Primary Auditory Cortex. *Front Neurosci*. 2007;1(1):197–209.
27. Luce PA, McLennan CT. Spoken Word Recognition: The Challenge of Variation. New York: Wiley; 2008. p. 590–609.
28. Mattys S, Davis M, Bradlow A, Scott S. Speech recognition in adverse conditions: A review. *Lang Cogn Neurosci*. 2012;27(7–8):953–78.
29. Montgomery R. A tour of sub-Riemannian geometries, their geodesics and applications. Mathematical Surveys and Monographs. vol. 91. Providence: Am. Math. Soc.; 2002.
30. Nelken I, Calford MB. Processing Strategies in Auditory Cortex: Comparison with Other Sensory Modalities. Boston: Springer; 2011. p. 643–56.
31. Parikh G, Loizou PC. The influence of noise on vowel and consonant cues. *J Acoust Soc Am*. 2005;118(6):3874–88.
32. Petitot J, Tondut Y. Vers une neurogéométrie. Fibrations corticales, structures de contact et contours subjectifs modaux. *Math Sci Hum*. 1999;145:5–101.
33. Petitot J, Tondut Y. Vers une Neurogéométrie. Fibrations corticales, structures de contact et contours subjectifs modaux. 1999;1–96.
34. Polger TW, Shapiro LA, Press OU. The multiple realization book. Oxford: Oxford University Press; 2016.
35. Prandi D, Gauthier J-P. A semidiscrete version of the Citti–Petitot–Sarti model as a plausible model for anthropomorphic image reconstruction and pattern recognition. *SpringerBriefs in Mathematics*. Cham: Springer; 2017.
36. Press WH, Teukolsky SA, Vetterling WT, Flannery BP. Numerical Recipes. 3rd ed. Cambridge: Cambridge University Press; 2007.
37. Rankin J, Sussman E, Rinzel J. Neuromechanistic Model of Auditory Bistability. *PLoS Comput Biol*. 2015;11(11):e1004555.
38. Rauschecker JP. Auditory and visual cortex of primates: a comparison of two sensory systems. *Eur J Neurosci*. 2015;41(5):579–85.
39. Sarti A, Citti G. The constitution of visual perceptual units in the functional architecture of V1. *J Comput Neurosci*. 2015;38(2):285–300.
40. Sethares W. Tuning, Timbre, Spectrum, Scale. London: Springer; 2005.
41. Sharma J, Angelucci A, Sur M. Induction of visual orientation modules in auditory cortex. *Nature*. 2000;404(6780):841–7.
42. Tian B, Kuśmierek P, Rauschecker JP. Analogues of simple and complex cells in rhesus monkey auditory cortex. *Proc Natl Acad Sci*. 2013;110(19):7892–7.
43. Wilson HR, Cowan JD. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys J*. 1972;12(1):1–24.
44. Zatorre RJ. Do you see what I'm saying? Interactions between auditory and visual cortices in cochlear implant users. *Neuron*. 2001;31(1):13–4.
45. Zhang J, Dashtbozorg B, Bekkers E, Pluim JPW, Duits R, ter Haar Romeny BM. Robust retinal vessel segmentation via locally adaptive derivative frames in orientation scores. *IEEE Trans Med Imaging*. 2016;35(12):2631–44.
46. Zulfqar I, Moerel M, Formisano E. Spectro-Temporal Processing in a Two-Stream Computational Model of Auditory Cortex. *Front Comput Neurosci*. 2020;13:95.