



LUND UNIVERSITY

Haptic input features in cortical processing

Kesgin, Kaan

2026

Document Version:

Publisher's PDF, also known as Version of record

[Link to publication](#)

Citation for published version (APA):

Kesgin, K. (2026). *Haptic input features in cortical processing*. [Doctoral Thesis (compilation), Department of Experimental Medical Science]. Lund University, Faculty of Medicine.

Total number of authors:

1

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Haptic input features in cortical processing

KAAN KESGIN

EXPERIMENTAL MEDICAL SCIENCE | FACULTY OF MEDICINE | LUND UNIVERSITY



LUND
UNIVERSITY

Haptic input features in cortical processing

Haptic input features in cortical processing

Kaan Kesgin



LUND
UNIVERSITY

DOCTORAL DISSERTATION

by due permission of the Faculty of Medicine, Lund University, Sweden.

To be defended at Segerfalksalen,
on the 8th of October 2026 at 13.00.

Faculty opponent

Professor Boris Barbour, École Normale Supérieure, Paris, France

Organization: LUND UNIVERSITY

Document name: DOCTORAL DISSERTATION

Date of issue: 2026-10-08

Author(s): Kaan Kesgin
Horizon 2020 MSCA-ITN 'INTUITIVE', grant no. 861166

Sponsoring organization: EU

Title and subtitle: Haptic input features in cortical processing:

Abstract:

Touch is conventionally described as a funnel. Mechanoreceptors transduce skin deformation, labelled channels carry frequency and force centrally, and the primary somatosensory cortex extracts the features perception is built from. This thesis measures the two junctions at which that account places most of the discarding, the mechanics of the skin and the cortical representation of what the skin sends, and finds that neither behaves as the account requires.

Paper I recorded populations in primary somatosensory cortex with the skin bypassed, using reproducible electrotactile and visual patterns. The spontaneous state space proved extremely high-dimensional, tactile processing was modulated by visual input, and input-specific activity outlasted the stimulus. Paper II derives the hyperlet, a wavelet integrating a continuum of scales into a single transform, which localises short bursts better than existing super-resolution estimators at a fraction of their computational cost; it re-expresses a signal instead of filtering it. Paper III applies that decomposition pixel by pixel to high-speed imaging of the human fingertip resting on a vibrating haptic display. The skin generates harmonics of the drive, organised into domains that follow the papillary ridges, and the harmonic class expressed at a location reconfigures with the stimulus rather than being fixed by anatomy. Paper IV delivers comparable vibrotactile patterns to the rat forepaw with the skin intact. Fine haptic detail proves decodable far from the direct thalamocortical input, no neuron is tuned to frequency over the range tested, and pattern-specific decoding persists for hundreds of milliseconds past stimulus offset.

Neither end of the pathway performs the reduction attributed to it. The skin adds spectral structure contingent on the stimulus, which removes the stable frequency variable that cortical frequency tuning would require; the cortex then distributes the result and holds it beyond the input instead of compressing it into local, labelled representations. The apparent simplicity of touch is in part a property of measurement: of bulk mechanical recording, band-pass filtering, response averaging and dimensionality reduction, each of which discards the structure reported here.

Key words: Haptics, Somatosensory cortex, Tactile, Skin mechanics, Time–frequency analysis, Population coding, Neocortex

Classification system and/or index terms (if any)

Supplementary bibliographical information

Language: English

Number of pages: 74

ISSN: 1652-8220

ISBN: 978-91-8021-929-7

Recipient's notes

Price

Security classification

I, the undersigned, being the copyright owner of the abstract of the above-mentioned dissertation, hereby grant to all reference sources permission to publish and disseminate the abstract of the above-mentioned dissertation.

Signature

Date 2026-[MM-DD]

Haptic input features in cortical processing

Kaan Kesgin

Copyright pp 1–74 Kaan Kesgin

Paper 1 © the Authors, Creative Commons Attribution 4.0 International

Paper 2 © the Authors, Creative Commons Attribution 4.0 International

Paper 3 © by the Authors (Manuscript unpublished)

Paper 4 © by the Authors (Manuscript unpublished)

Faculty of Medicine

Department of Experimental Medical Science

ISBN 978-91-8021-929-7

ISSN 1652-8220

Printed in Sweden by Media-Tryck, Lund University

Lund 2026

This thesis has been partly produced with the help of the generative AI model Claude (Anthropic). I have processed the generated text and take full responsibility for the content. An account of how the model was used is given in the Methods section.

Table of contents

Preface	10
List of original papers	12
Author's contribution to the papers	13
Abbreviations	14
Abstract	15
Introduction	16
The skin transforms before it transduces	17
Compliance, shear and travelling waves.....	17
Nonlinearity at the contact	18
Papillary ridges as mechanical architecture	18
The mechanoreceptive afferents	18
How large is the space of possible touches?.....	19
What happens in between	20
The cuneate nucleus.....	21
Thalamic relays	21
The cortical end of the pathway	21
Somatosensory cortex and its supposed boundaries.....	21
Cortical state and the malleability of a response	22
The measurement problem.....	23
Bulk mechanics cannot see spatial structure	23
Filtering removes what decomposition preserves.....	24
Averaging and reduction remove what single trials preserve.....	25
Aims	26
Methods	27
Experimental	27
Ethical considerations.....	27
The haptic display	28
Stimulus design.....	28
High-speed imaging of the fingertip	29
Animals, anaesthesia and surgery	30

Neuropixels recordings	30
Analysis.....	31
Time–frequency decomposition.....	31
Quantifying the resolution of a transform.....	31
Per-pixel harmonic analysis.....	32
Spike sorting.....	32
Single-trial decoding	32
State-space characterisation	33
Alternatives considered	33
Use of generative AI	35
Results	36
Paper I: High-dimensional cortical signals reveal rich bimodal and working memory-like representations among S1 neuron populations	36
Paper II: Super-resolution time–frequency decomposition with hyperlets for neural spike analysis.....	37
Paper III: Spatially distributed, stimulus-dependent harmonic generation in the fingertip skin during interaction with a haptic display	42
Paper IV: Highly resolved representations of dynamic haptic inputs in non- primary sensory cortical neuron representations	44
Discussion	48
The skin expands rather than reduces.....	48
Harmonic generation as an expansion	48
Stimulus-dependence and the loss of frequency as a variable	49
Mechanics as part of the computation	49
The cortex distributes and sustains rather than localises	50
Fine haptic detail far from the primary area.....	50
Why no neuron is tuned to frequency	51
Persistence and a substrate for working memory	52
Cross-modal structure in a unimodal area	52
What the measurement determines.....	53
Re-expression versus reduction	53
The cost of dimensionality reduction	54
Situating the work: the plenhaptic function.....	54
Methodological considerations	55
Anaesthesia	55
Human fingertip and rat forepaw.....	56
Statistical reach of Paper III.....	56
What Paper IV cannot test	56
How far a better time–frequency representation gets you	57

The broader perspective	57
Implications for haptic engineering	57
Implications for how cortical representations are studied.....	58
Future directions.....	58
Concluding remarks	60
Populärvetenskaplig sammanfattning	62
Acknowledgements.....	64
Henrik Jörntell	64
Fredrik Bengtsson	64
Udaya Rongala	64
Johanna Norrlid.....	64
Ann-Sophie Alm.....	65
Sofie Skårup Kristensen	65
Anders Wahlbom and Jonas Enander	65
Fellow doctoral students	65
The INTUITIVE network.....	65
References	66

Preface

This thesis argues that haptic information is not funnelled as it moves inward. It is re-expressed. At the skin, and again in the cortex, the representation of a vibrotactile input becomes richer, more distributed and more dependent on context than the input itself was. The orderly picture drawn by much of the classical literature (frequency channels, labelled lines, a primary area that extracts features and hands them on) is, I will argue, substantially a consequence of measuring in a way that discards structure before looking for it.

In the introduction I build the problem from the outside in. I start at the fingertip, with the skin as a compliant, pre-tensioned, mechanically coupled medium that transforms the stimulus before it is transduced into neural signals by the mechanoreceptors, and with the theoretical question of how large the space of possible haptic interactions actually is. I note in passing what the cuneate nucleus and the thalamus are known to do to the signal on its way through, only so far as that bears on reading the two ends, and I review the accumulating evidence that tactile information is far more widely distributed across the cortex than the somatotopic account allows. The last part of the introduction is about measurement: why the tools conventionally used to look at skin mechanics and at neuronal activity (bulk impedance, band-pass filtering, response averaging, aggressive dimensionality reduction) each remove the structure I am trying to see, and why a different class of tool was needed first.

A methods section then gives an overview of the experimental and analytical approaches shared across the four papers: the time–frequency machinery, the haptic display and high-speed imaging, the acute *in vivo* recordings, and the single-trial decoding used in place of response averaging.

The thesis is a study of the two ends of that pathway: the point at which a mechanical event becomes a sensory signal, and the point at which that signal becomes a cortical representation. These are the two places the conventional account makes its strongest claims, and they can be measured directly, which is why the work is placed there.

The results are presented one paper at a time, in the order in which the work was done, which also traces the argument. Paper I establishes what a cortical neuron population does with a tactile input when the skin is bypassed altogether: the spontaneous state space is extremely high-dimensional, tactile processing is modulated by vision, and input-specific information outlives the input. Paper II derives the hyperlet and its transform, a super-resolution time–frequency decomposition that re-expresses a signal in a richer basis in-

stead of filtering information away. Paper III uses that decomposition to image the human fingertip in contact with a haptic display, and finds that the skin generates harmonics of what it is given, organises them into domains that follow the papillary ridges, and reconfigures those domains when the stimulus changes. Paper IV then delivers comparable vibrotactile inputs to the rat forepaw with the skin reinstated; fine haptic detail proves decodable at cortical sites far from the direct thalamocortical input, no neuron is tuned to frequency over the range tested, and pattern-specific information again outlives the stimulus.

The discussion is organised around the three claims these papers support jointly, not paper by paper: that the skin expands rather than reduces the information it is given, that the cortex distributes and sustains that information where the conventional account has it localising and discarding, and that seeing either of them requires measuring the signal before any step that discards part of it. It considers what the four papers can and cannot establish together, given that Paper III is human and passive while Papers I and IV are anaesthetised rat, and it places the findings against the theoretical framework of the plenhaptic function and of cross-modal invariants developed within this research environment.

List of original papers

This thesis is based on the following papers, referred to in the text by their Roman numerals.

- I. Sofie S. Kristensen, **Kaan Kesgin** and Henrik Jörntell. High-dimensional cortical signals reveal rich bimodal and working memory-like representations among S1 neuron populations. *Communications Biology*, 7(1):1043, 2024.
- II. **Kaan Kesgin** and Henrik Jörntell. Super-resolution time–frequency decomposition with hyperlets for neural spike analysis. *Pattern Recognition*, 172:112598, 2026.
- III. **Kaan Kesgin** and Henrik Jörntell. Spatially distributed, stimulus-dependent harmonic generation in the fingertip skin during interaction with a haptic display. Submitted.
- IV. **Kaan Kesgin** and Henrik Jörntell. Highly resolved representations of dynamic haptic inputs in non-primary sensory cortical neuron representations. Manuscript, under review at *Scientific Reports*.

Papers I and II are reproduced under the Creative Commons Attribution 4.0 International licence. Papers III and IV are unpublished manuscripts and are reproduced by permission of the authors.

Paper I was also included in the doctoral thesis of Sofie S. Kristensen.

Author's contribution to the papers

Paper I. SSK and HJ planned and designed the study and the main analyses, and SSK conducted the experiments and the analysis. I contributed to the main analysis and designed the convolutional neural network control analysis. Specifically, I built that network and its training procedure and established how its output was to be interpreted, and I supervised the principal component analysis and provided input to the interpretation and control of its results. I contributed to the composition of the manuscript by setting out what the different analysis methods implied and thereby the order in which the figures should be arranged, and I wrote the text describing the analysis methods. SSK and HJ wrote the article; all authors contributed to it and approved the submitted version.

Paper II. Conceptualization, methodology, software, formal analysis, investigation, data curation, validation, visualization, resources, writing of the original draft, and review and editing. I am solely responsible for the derivation of the hyperlet transform. I designed all tests and controls, generated all data, performed all analyses, produced all text and all figures, and authored the interactions with the reviewers.

Paper III. I designed the whole study, including the experiments and the analyses. I performed all experiments, performed all analyses, designed all figures and wrote all text.

Paper IV. I designed the experiment, performed all experiments and all analyses, and produced all figures. The text was co-authored with my supervisor, with myself in a main role, and I played a main role in the manuscript composition and in the interaction with the reviewers across two separate rounds of revision.

Abbreviations

S _I	Primary somatosensory cortex
V _I	Primary visual cortex
AP	Action potential
UP state	Depolarised phase of the cortical slow oscillation
PCA	Principal component analysis
kNN	<i>k</i> -nearest neighbours
CNN	Convolutional neural network
F _I	Harmonic mean of precision and recall, used here as a per-unit decoding score
STFT	Short-time Fourier transform
CWT	Continuous wavelet transform
SLT	Superlet transform
FASLT	Fractional adaptive superlet transform
HLT	Hyperlet transform
MMCE	Minimum mean cross-entropy

Abstract

Touch is conventionally described as a funnel. Mechanoreceptors transduce skin deformation, labelled channels carry frequency and force centrally, and the primary somatosensory cortex extracts the features perception is built from. This thesis measures the two junctions at which that account places most of the discarding, the mechanics of the skin and the cortical representation of what the skin sends, and finds that neither behaves as the account requires.

Paper I recorded populations in primary somatosensory cortex with the skin bypassed, using reproducible electrotactile and visual patterns. The spontaneous state space proved extremely high-dimensional, tactile processing was modulated by visual input, and input-specific activity outlasted the stimulus. Paper II derives the hyperlet, a wavelet integrating a continuum of scales into a single transform, which localises short bursts better than existing super-resolution estimators at a fraction of their computational cost; it re-expresses a signal instead of filtering it. Paper III applies that decomposition pixel by pixel to high-speed imaging of the human fingertip resting on a vibrating haptic display. The skin generates harmonics of the drive, organised into domains that follow the papillary ridges, and the harmonic class expressed at a location reconfigures with the stimulus rather than being fixed by anatomy. Paper IV delivers comparable vibrotactile patterns to the rat forepaw with the skin intact. Fine haptic detail proves decodable far from the direct thalamocortical input, no neuron is tuned to frequency over the range tested, and pattern-specific decoding persists for hundreds of milliseconds past stimulus offset.

Neither end of the pathway performs the reduction attributed to it. The skin adds spectral structure contingent on the stimulus, which removes the stable frequency variable that cortical frequency tuning would require; the cortex then distributes the result and holds it beyond the input instead of compressing it into local, labelled representations. The apparent simplicity of touch is in part a property of measurement: of bulk mechanical recording, band-pass filtering, response averaging and dimensionality reduction, each of which discards the structure reported here.

Introduction

The textbook account of touch is a chain of conversions. Contact deforms the skin; mechanoreceptors embedded in the tissue transduce that deformation into trains of action potentials; channels defined by receptor class carry those signals through the dorsal column nuclei and the thalamus; and the primary somatosensory cortex assembles, from what arrives, a representation of the object that was touched. It is a serviceable account, and it has organised a great deal of productive work. Its defining property is that every stage is a reduction. The mechanical event at the fingertip has more structure than the afferent volley, the afferent volley has more structure than what survives central processing, and the cortical representation is understood as the distilled residue: the features that mattered, separated from everything that did not.

This thesis takes that account seriously enough to ask where the reduction actually occurs. It is a question that can be put empirically, because the account localises the discarding at two junctions in particular. The first is the skin, which is treated as the point where a continuous mechanical field becomes a finite set of receptor activations. The second is the primary sensory cortex, which is treated as the point where those activations become labelled features. Both junctions are examined here, and neither is found to behave as the account requires.

There is a straightforward reason this has been easy to miss. Each of the phenomena described in this thesis lies below the resolution, or outside the scope, of the measurement conventionally used to look for it. Bulk mechanical measurements of the fingertip average over the contact area, so a transformation organised into spatial domains reads as a single averaged transformation. Band-pass filtering of an extracellular recording removes signal content in order to make spike detection tractable, so structure carried in the discarded bands is not available to be found. Averaging responses across repetitions of a stimulus removes the trial-to-trial variation in which state-dependence lives. Reducing a neuron population to a few principal components removes the higher dimensions in which the population may be doing most of its work. In each case the method is reasonable and the loss is deliberate; the difficulty is that the losses are correlated with precisely what this thesis set out to see. The introduction therefore ends where the work began, with the problem of measuring without first simplifying.

The skin transforms before it transduces

Whatever the central nervous system eventually extracts about a contacted object has first been filtered through the mechanics of the tissue in which the mechanoreceptors are embedded (Johansson and Flanagan, 2009). The fingertip is not a rigid probe but a layered, viscoelastic and pre-stressed structure whose deformation is far from a linear reproduction of its input (Wiertlewski and Hayward, 2012). This is not a caveat to be noted and set aside. If the tissue transforms the stimulus before any receptor responds to it, then the transformation is part of the sensory pathway, and any account of what the afferents encode is incomplete without it.

Compliance, shear and travelling waves

Because biological tissues are dynamically compliant, yielding under load and springing back, a mechanical event at one point does not stay at that point. Vibrations elicited at a contact propagate across the skin and decay at frequency-dependent rates (Rongala et al., 2024), and the extent of that propagation shapes which afferents are recruited and how strongly (Manfredi et al., 2012; Shao et al., 2016). A localised dynamic input hence becomes a distributed, dispersive event in the tissue, with different frequency components attenuated to different degrees as they travel.

The consequence is that mechanoreceptors are not point samplers of a local stimulus. They sit inside a medium that is carrying superimposed travelling waves from the contact, and they respond to the local differential deformations, or shear forces, that those waves produce (Ramasamy and Jörntell, 2026). The same reasoning extends beyond the skin: because the body is mechanically coupled throughout, muscle spindles and joint capsule afferents also report changes in contact state, and the classical division between exteroceptors and proprioceptors becomes correspondingly harder to sustain (Ramasamy and Jörntell, 2026).

This distributed mechanical behaviour is not merely a source of blur. In a network model of a cortical-like recurrent circuit receiving input through a simulated compliant skin, the dependencies that mechanical coupling creates among sensors turned out to be what made autonomous learning work: the network acquired representations that generalised to previously unseen inputs, and remained able to separate them after substantial loss of sensors or of internal connections (Rongala and Jörntell, 2026). Training the same network on inputs stripped of those dependencies degraded its performance instead. On that account the mechanics of the tissue are not noise imposed on the signal but part of the structure that makes the signal learnable.

Nonlinearity at the contact

The interaction between a finger and a surface is, in addition, genuinely nonlinear. When a finger slides over a periodic grating, the interaction force contains not only the fundamental imposed by the grating period but a series of harmonics of it, and the amplitudes of those harmonics decay with a non-integer power law, which is a signature of strong nonlinearity in the contact (Wiertlewski et al., 2011a,b). Related measurements of the lateral mechanical impedance of the fingertip across the frequencies and displacements relevant to touch show the same departure from a simple linear response (Wiertlewski and Hayward, 2012).

What these measurements do not settle, however, is where in the contact the nonlinearity resides. Bulk impedance measurements and single-point recordings integrate over the contact area, so they cannot distinguish a transformation that is uniform across the skin from one organised into distinct spatial domains, and they cannot say whether the behaviour at a given location is a fixed property of the local tissue or something that changes with the stimulus. A spatially resolved account has been missing, and it is the gap Paper III addresses.

Papillary ridges as mechanical architecture

The glabrous skin of the fingertip is not a smooth sheet. Its surface is organised into papillary ridges, and in contact with a flat surface those ridges determine where the skin actually touches and where it does not. The ridges are usually discussed in terms of friction and of the spatial sampling of surface features, but they are also mechanical architecture. They set the geometry of the contact band that is loaded when the surface moves, and with it the geometry over which shear is distributed. They also mean that the amount of tissue beneath a given point varies with its position relative to the ridge crest, so the local resonant properties should differ across the contact as well. If the contact is ridged, the deformation under a tangential drive will not in general be spatially uniform, and any nonlinearity in that deformation should be organised by the same geometry. This expectation is what Paper III sets out to test directly, by resolving the movement of the skin across the contact rather than integrating over it.

The mechanoreceptive afferents

The afferents that transduce this movement are conventionally divided by adaptation properties into slowly adapting and rapidly adapting classes, with the Pacinian channel responding to the highest frequencies. The division is real and useful, but it is coarser as a description of coding than it is often taken to be: the individual sensors have different preferred tunings across vibratory frequencies, yet each covers a wide frequency range and none of them is confined to a narrow band (Johansson et al., 1982; Ramasamy and Jörntell, 2026). Recordings from awake, freely moving animals make the same point

from the other direction: Pacinian afferents prove diverse in both tuning and sensitivity, and are driven as readily by the animal's own movement as by low-amplitude vibration reaching it from the environment (Turecek and Ginty, 2024). The vibrations produced during contact carry information about fine texture that can be read out through the rapidly adapting and Pacinian channels (Bensmaïa and Hollins, 2003; Weber et al., 2013), and both spatial and temporal codes contribute to the perception of natural textures (Weber et al., 2013).

Two features of this arrangement matter for what follows. First, a given mechanical event will generally recruit several classes at once, not one labelled line. Second, because each afferent samples the local movement of a medium that is itself transforming the stimulus, the signal available to a given afferent depends both on where it sits and on what the stimulus is. Whatever tuning is measured at the afferent is a property of the receptor and of the tissue around it, and those two contributions are not separable by recording from the afferent alone.

How large is the space of possible touches?

Behind the question of where information is reduced sits a prior one: how much there is to reduce. Vision has a clean answer to it. The plenoptic function returns the intensity of every light ray that could be registered at any position, from any direction, at any wavelength and any moment, and specifying one such ray takes seven numbers (Adelson and Bergen, 1991). No eye samples more than a thin projection of it, but the function itself is seven-dimensional, and that number describes the world and not the eye.

The same question was put to touch, and the answer that came back was awkward (Hayward, 2011). A contact is an interaction between two bodies of which at least one deforms, so describing it means relating the trajectories of the points of the one to the trajectories of the points of the other. Both are continua, and a map between continua needs infinitely many coordinates unless simplifying assumptions are imposed on the mechanics first. There is no seven for touch.

A finite count does become available if the question is asked at the skin instead of at the interaction (Ramasamy and Jörntell, 2026). An evolving contact drives differential movements of skin tissue compartments, those movements travel through the tissue as dispersive waves, and the wave field at a point can be decomposed at each instant into frequencies with an amplitude at each. Read this way the biologically relevant plenihaptic function takes six coordinates: frequency, amplitude, time, and three of ordinary physical space, locating a point inside the three-dimensional volume of the skin.

Two clarifications matter here, because the word dimension is doing more than one job in this thesis.

The first is that those six coordinates are a domain and not an inventory. They fix how many numbers are needed to say where and what one component of a mechanical signal is; they say nothing about how many are needed to specify a touch. A touch is an entire field over that domain, frequencies and amplitudes distributed through the tissue and changing from moment to moment, so the set of possible touches stays unbounded exactly as Hayward's argument requires. What the six coordinates buy is not a smaller answer but an addressable one, which is what allows a measurement to be aimed at part of it. Its constituents are then not static features but kinematic invariants, time-evolving activation patterns that recur consistently enough across contexts to be worth representing (Jörntell et al., 2026).

The second is that three of those six are physical space in the everyday sense, depth into the tissue and two coordinates across it. Wherever the plenhaptic function comes up again in this thesis those three are meant literally. Everywhere else, dimensionality refers to a state space with one axis per recorded neuron, and a count there is a claim about how many independent patterns a population expresses. The two quantities share a word and nothing else.

The distinction also settles what it would take for a stage to reduce anything. Position alone cannot answer that, since sampling a continuous mechanical field with a finite receptor population can only lose spatial resolution; a purely spatial account has its conclusion built in before any measurement is made. Frequency and time are what make the question empirical. A stage returning energy at frequencies its input did not contain, at places its input did not reach, has enlarged the region of the domain the signal occupies, and a stage discarding them has shrunk it, and which of the two happens can be measured instead of assumed.

Something must nonetheless perform the reduction, and it is important to know which stage does. Four of the six coordinates, frequency and amplitude as functions of time and of position across the contact, are what high-speed imaging of the moving skin can deliver. Hence Paper III is an attempt to sample that object at the first stage, not to assume its shape.

What happens in between

Between the skin and the cortex lie the cuneate nucleus and the thalamus. Neither is studied in this thesis and neither is surveyed here. What follows is the minimum needed to read the two ends: enough to establish what has already been done to the afferent signal before it reaches the cortex, and hence what can and cannot be attributed to the cortex itself.

The cuneate nucleus

Tactile afferent signals from the forelimb are first processed in the cuneate nucleus of the brainstem, and the processing there is not relay-like. Individual cuneate neurons integrate sensory quanta from multiple afferents (Bengtsson et al., 2013), and spatiotemporal skin strain distributions evoke spike responses in them with strikingly low trial-to-trial variability (Hayward et al., 2014). More pointedly for the framework used here, cuneate neurons were found to decompose tactile input into subsets of what are now called haptic kinematic invariants (Jörntell et al., 2014); that observation is where the idea originates (Jörntell et al., 2026; Ramasamy and Jörntell, 2026). It is direct evidence that reduction of the plenhaptic function begins before the thalamus and has a basis in the physics of contact. Intracellular dynamics in these neurons appear well suited to acquiring such representations by a self-stabilising learning process (Rongala et al., 2018).

The recombination at this stage is extensive. Tonic and phasic afferent signals are fused, not kept apart, which means that the receptor-class distinctions drawn at the periphery are already partly dissolved one synapse into the central nervous system (Jörntell et al., 2014; Hayward et al., 2014). The spinothalamic pathways provide a parallel route, but a more indirect one: it is relayed through neurons of the spinal cord, where the signal is mixed with input from several sensor types and shaped by intrinsic spinal network processing, and it tends to carry less precise tactile information (Craig, 2003; Dum et al., 2009).

Thalamic relays

The thalamus continues this process instead of reversing it. Individual thalamic neurons represent input from both rapidly and slowly adapting low-threshold mechanoreceptors, so that cortical responses to touch reflect subcortical integration, not the preservation of separate labelled channels (Emanuel et al., 2021). Tactile input patterns, moreover, are decodable widely among thalamic neurons, not only within a restricted somatosensory subset (Wahlbom et al., 2021). By the time afferent activity reaches the cortex it has passed through at least two stages at which nonlinear integration occurred, and relatively little of the structure of the original afferent signal should be expected to survive intact.

The cortical end of the pathway

Somatosensory cortex and its supposed boundaries

The conventional picture of primary somatosensory cortex is somatotopic and modality-specific: a map of the body surface, processing touch. A substantial body of work is difficult to reconcile with it. Using an electrotactile interface capable of delivering precisely reproducible, complex spatiotemporal patterns to the primary afferents of the skin (Oddo

et al., 2017), it has been shown that neocortical neurons decode tactile input patterns ubiquitously (Enander et al., 2019), and that somatosensory cortical neurons decode both pattern and location from digits on either side of the body (Enander and Jörntell, 2019). Responses to a given modality are not confined to its nominal primary area (Ghazanfar and Schroeder, 2006; Olcese et al., 2013; de Lafuente and Romo, 2006), and focal cortical lesions impair information processing in neurons at a distance from the lesion (Wahlbom et al., 2019).

The anatomical substrate for such spread is better documented than it once was. Individual corticocortical axons arborise far more widely than traditionally assumed, as data from the MouseLight database show (Szeier et al., 2025), and cortico-thalamo-cortical loops provide a further route by which activity evoked in one area can reach many others.

A competing interpretation must be taken seriously. Where responses to one modality appear in another modality's primary area, they may reflect behavioural responses that the stimulus triggered, and not the stimulus itself; this has been argued for sound-evoked activity in mouse visual cortex (Bimbard et al., 2023), and movement accounts for a great deal of the variance in single-trial cortical activity (Musall et al., 2019; Stringer et al., 2019). The argument does not dispose of the phenomenon, since responses that differ according to which stimulus was presented imply that the behaviour was itself stimulus-specific, but it does mean that experiments intending to characterise the passive propagation of sensory information through the cortical network have a reason to suppress behavioural reactions. That is one motivation for the lightly anaesthetised preparation used in Papers I and IV.

Cortical state and the malleability of a response

Even within a single area, the response of a cortical neuron to a fixed input is not fixed. Responses to identical tactile input patterns vary with the internal state of the network (Norrlid et al., 2021), remote cortical perturbation changes the network solutions that given tactile inputs elicit (Etemadi et al., 2022), and hippocampal output profoundly affects how tactile input patterns are interpreted in S1 (Etemadi et al., 2023). Adjacent neocortical neurons do not show repetitive correlation patterns with one another (Mogensen et al., 2019), which is to say that the population is not composed of redundant copies. Spontaneous activity, correspondingly, is neither low-dimensional nor idle (Stringer et al., 2019, 2021), and simultaneously recorded spikes and field potentials can convey partly independent information about the same input (Belitski et al., 2008). Comparable state-dependence has been described for colour processing in visual cortex (Kristensen and Jörntell, 2023). Where the fluctuations themselves originate is a further question. Channel current fluctuations account for them in cerebellar Purkinje cells (Nilsson and Jörntell, 2021), and modelling of closed-loop cortical and spinal networks indicates that the continual movement of the population state observed *in vivo* requires either dynamic sensory input or activity generated independently within the neurons themselves (Szeier and Jörntell, 2026).

Where sensory drive is absent, as in the spontaneous activity considered here, only the latter remains.

This variability is often treated as noise to be averaged away. The alternative reading, which this thesis adopts, is that it is the signature of a network whose internal state carries information of its own, and that a sensory input is interpreted in the context of whatever state it happens to arrive in. On that reading the correct experimental response is not to average over states but to characterise the state space and treat the input as a perturbation of it, which is the approach taken in Paper I. It also motivates analysing individual trials in place of mean responses, since averaging is the operation that removes the effect of interest. The point is not new: ensembles of cortical neurons were shown a quarter of a century ago to support prediction on single trials, with no averaging at all (Wessberg et al., 2000).

The functional argument for why a cortex should behave this way has been set out in terms of cross-modal invariants (Ramasamy and Jörntell, 2026). Unimodal invariants have a ground truth in physics and can be acquired pre-cortically; invariants that link modalities do not, and must be learned individually, retained over time, and applied only when relevant. Retention of that kind requires a network with recursive loops long enough to hold information in the absence of the input that produced it, which points to the thalamocortical system. If such invariants suffuse the cortical network and do not reside in identifiable convergence zones, then cross-modal structure and working-memory-like persistence should be detectable in ordinary sensory cortical populations. Papers I and IV test both expectations.

The measurement problem

Each of the claims above is invisible to at least one of the measurements conventionally used to investigate it. This was the practical obstacle that had to be cleared before the rest of the work could proceed, and it recurs in the same form at the periphery and at the cortex.

Bulk mechanics cannot see spatial structure

A measurement that integrates over the contact area returns one number where the question requires a map. If harmonic generation in the skin were organised into domains, an impedance measurement would report their average; if the organisation changed with the stimulus, the average might barely move. The same limitation applies to single-point recordings, which sample one location and cannot report whether it is representative. Resolving the question requires measuring the movement of the skin at many points at once, at a temporal resolution sufficient to recover the frequency content at each of them.

Filtering removes what decomposition preserves

For extracellular recordings the corresponding obstacle is band-pass filtering. Filtering is a destructive operation: it alters the information content of the signal so that spikes can be detected, and what is removed is not recoverable. A time–frequency decomposition does something different in kind. It re-expresses the same signal in another vocabulary, one in which events that overlap in time may be separable by their spectral signature, and it does so without discarding the content of the original.

The difficulty is that decomposition is subject to the uncertainty principle: resolution gained in time is paid for in frequency (Cohen, 1966). The short-time Fourier transform must choose a window length, and the choice is wrong for some part of any signal spanning a wide frequency range; the multiresolution Fourier transform was an early attempt to escape that single choice (Wilson et al., 1992). The continuous wavelet transform scales the window with frequency and so addresses the multiresolution problem directly (Morlet et al., 1982; Addison, 2018), but remains bound by the same constraint.

Two lines of work have pushed against it. One reshapes the distribution itself, as in the exponential kernels used to suppress cross-terms in multicomponent signals (Choi and Williams, 1989). The other combines several decompositions: taking the geometric mean of spectra computed at different window lengths yields a more localised positive time–frequency function, and that combination is optimal in the sense of minimising the mean cross-entropy to the underlying energy distribution (Loughlin et al., 1994). Applying the same idea to wavelets gives the superlet transform (Moca et al., 2021) and its fractional adaptive variant. The remaining problems are practical: the ensemble is computationally expensive, and if the scales are spaced too widely a feature can be suppressed whenever any single scale returns a near-zero response. Paper II closes that gap.

The reach of these methods is worth noting, because it bears on how widely the result should be expected to matter. Time–frequency decomposition underpins work as varied as fingerprint enhancement (Chikkerur et al., 2007), seismic segmentation (Hashemi, 2012), speech and music classification (Bhattacharjee et al., 2020), long-horizon forecasting (Luo et al., 2025; Zhao et al., 2025), and clinical signal analysis (Wu et al., 2021; He et al., 2022); super-resolution variants have since been applied to seizure detection, cardiac arrhythmia classification, phonon dynamics and radio-frequency fingerprinting (Tripathi et al., 2023, 2022; Gentry et al., 2022; Baldini and Dimc, 2025).

Why this matters for the biology is that low-impedance multichannel probes now record many neurons at once (Buzsáki, 2004; Jun et al., 2017; Steinmetz et al., 2021; Chung et al., 2019), and the recorded signal is a mixture of the characteristic waveforms of many sources. The problem is formally the cocktail party problem, and it is soluble to the extent that the sources have separable time–frequency signatures. Any loss of fidelity in the representation propagates: it affects which spikes are detected, which are assigned to which unit, and hence which population is analysed at all.

Averaging and reduction remove what single trials preserve

The third obstacle is analytical rather than instrumental. Averaging responses across repetitions is the standard way of establishing that a neuron responds to a stimulus, and it discards exactly the trial-to-trial variation that cortical state-dependence produces (Norrlid et al., 2021; Etemadi et al., 2022). A measure that operates on individual trials, and asks whether responses to different inputs form separable clusters, addresses a different and arguably more relevant question: not whether the mean response differs, but whether the identity of the input can be recovered from a single presentation.

Dimensionality reduction poses a related problem one level up. Embedding a population recording in a small number of dimensions makes it tractable and usually makes it look orderly, but every embedding carries an implicit assumption about the computational architecture of the circuit, and that many different embeddings all return workable answers shows that none of them has been established as the right one (Ramasamy and Jörntell, 2026; Spanne and Jörntell, 2015). Give a recording enough independent degrees of freedom and some projection of it will track some external variable, so a demonstrated correlation is only worth as much as the constraints placed on the embedding beforehand (Kristensen et al., 2024; Roland, 2023). The argument here is not against reduction, which this thesis uses throughout, but against treating the number of retained dimensions as a presentational convenience when it is a scientific claim.




Figure 1. From skin mechanics to cortical representation. The conventional account treats each arrow as a reduction. This thesis examines the two stages shown in red and finds that neither reduces: the skin adds spectral structure that depends on the stimulus, and the cortical representation is high-dimensional, distributed and persists beyond the input. The relay stages between them are not studied here; they are drawn in because what they are known to do to the signal constrains what can be attributed to the cortex.

Taken together, these three obstacles describe a single situation. The conventional tools for looking at touch each simplify the signal before examining it, and they simplify it in ways that would conceal spatially organised mechanical structure at the periphery and high-dimensional, state-dependent, persistent structure at the cortex. The four papers of this thesis were an attempt to look without simplifying first, and to find out what is there when one does.

Aims

The overall aim of this thesis is to establish where, along the path from a mechanical event at the fingertip to a cortical population response, haptic information is actually reduced, and to show that at the two junctions where reduction is conventionally assumed, it does not occur.

Reaching that aim required a measurement problem to be solved first, and the specific aims follow from that dependency:

- **Paper I.** To characterise the dimensionality of the spontaneous state space of primary somatosensory cortical neuron populations, and to test whether tactile representations in that state space are modulated by another modality and whether they outlast the input, all with mechanotransduction bypassed.
- **Paper II.** To derive a time–frequency decomposition capable of separating short, overlapping, similar bursts, and to establish that it re-expresses a signal in a form more accessible to a classifier than the band-pass filtered voltage from which information has already been removed.
- **Paper III.** To determine, with spatial resolution across the contact area, how the fingertip skin transforms a controlled vibrotactile input before transduction: specifically, whether that transformation is uniform across the skin, whether it is linear, and whether it is a fixed property of the tissue or contingent on the stimulus.
- **Paper IV.** To test whether fine nuances of dynamic vibrotactile input, delivered through the intact mechanotransduction step, are represented at cortical sites remote from the direct sensory thalamocortical input, and whether cortical neurons encode such input in terms of frequency.

Methods

This section gives an overview of the experimental and analytical methods used in the four papers; the full detail is in the papers themselves. Three choices recur throughout and are worth stating once. Inputs were made highly reproducible, so that variability in the response could be attributed to the system under study, not to variation in what it was given. Inputs were made dynamic and long-lasting instead of brief and monotonous, because a static stimulus tests how much of the afferent signal is preserved, which is a different question from how the network transforms it. And measurements were taken at the highest resolution the method allowed, because the structure reported in this thesis exists above the resolution at which the conventional analysis operates.

Experimental

Ethical considerations

All animal experiments were approved in advance by the Malmö/Lund animal ethics committee, and all procedures were performed in accordance with the relevant guidelines and regulations. Paper I was performed under ethical permit M13193-2017. The work in Paper IV is reported in accordance with the ARRIVE guidelines. Animals were maintained in the Lund University animal facilities on a 12 h light/dark cycle, housed two or three to a cage and habituated before the experiments, with free access to food and water.

All recordings were acute and terminal, and the animals were killed with an overdose of pentobarbital at the end of the experiment without regaining consciousness. The acute anaesthetised preparation was chosen because the question asked here concerns the passive propagation of sensory information through the cortical network, which in an awake animal is confounded by behavioural reactions to the stimulus (Bimbard et al., 2023; Musall et al., 2019). Animal numbers were kept to those needed to establish the effects reported: five animals in Paper I and eight in Paper IV, one insertion per animal in the latter.

The human imaging in Paper III was performed on one healthy adult volunteer who gave written informed consent. The measurement was non-invasive throughout: the finger rested passively on the surface of the stimulator and was imaged optically.

The haptic display

Papers III and IV used the same class of vibrotactile stimulator, a custom haptic display built around a wide-band electrodynamic actuator (Actronika SAS) coupled to a rigid plexiglass plate. The actuator combines a Halbach magnet array, generating regions of quasi-uniform magnetic field, with a pair of counter-oriented flat coils, so that actuation arises from the mutual Laplace force between the two modules. A unidirectional four-leaf-spring suspension, whose bidirectional stiffness was tuned through the thickness and width of the leaves, was used to keep the frequency response flat and repeatable across the working band and to prevent distortion under load. The design target was a flat response between 50 and 500 Hz at peak-to-peak displacements up to 100 μm .

In Paper III the plate was driven tangentially, in the plane of the contact, so that the fingertip was loaded by lateral shear, not by normal compression. Drive waveforms were synthesised digitally at 30 kHz. An accelerometer mounted on the plate and sampled at 100 kHz was used to pre-compensate the drive: the measured acceleration yielded an amplitude-correction factor that flattened the surface acceleration spectrum, and the corrected waveform was the one delivered.

That accelerometer signal carries the weight of the central claim in Paper III. Measuring the rigid plate alone, it showed no harmonics of the driving frequency. It hence serves as a harmonic-free reference for the input, and any harmonic content subsequently found in the movement of the skin cannot have originated in the stimulus or the stimulator.

Stimulus design

Tone bursts and chirps (Paper III)

Two stimulus families were used. Tone bursts were Morlet-windowed, with centre frequencies stepped from 75 to 475 Hz in 25 Hz increments, each burst approximately 300 ms long with about 500 ms between bursts; the number of cycles was chosen to fill the burst duration. The harmonic analysis concentrates on three representative centre frequencies, 75, 100 and 126 Hz. Chirps were linear frequency sweeps from 80 to 480 Hz with total durations of 1, 3, 6 and 10 s, presented singly and as overlaid combinations of two or three sweeps of different rates, and were used for the filtering analysis. Each sequence was repeated three times, and stimulus onsets were marked with digital triggers so that the optical recording could be aligned to the stimulus timeline.

Vibrotactile tunes (Paper IV)

Eight waveform sequences, referred to as tunes, were designed for the haptic stimulation and sampled at 30 kHz to match the recording. Each tune comprised seven tones of varying pitch and length; each tone was a sinusoid multiplied by a Gaussian window whose width was set by a per-tone scale factor. Tune 1, for example, used the frequencies 112, 100, 134, 142, 100, 89 and 84 Hz with scale factors 1.5, 0.5, 1.0, 1.75, 1.25, 0.75 and 0.25; Tune 2

used frequencies at double those values with the same scale factors, and Tune 3 a fixed randomised reordering of the frequencies of Tune 1.

The design matters for the interpretation. Consecutive tones within a tune always differ in frequency, and the frequency content overlaps substantially between tunes, so the tunes cannot be told apart by their spectra alone. Their peak-to-peak amplitude was held constant while pulse durations and intervals varied, so what distinguishes them is temporal structure. Patterns were presented in random order with a random relaxation period of 0.9–1.1 s between presentations.

Electrotactile and visual patterns (Paper I)

Paper I bypassed mechanotransduction in order to deliver patterns of maximal reproducibility at the level of afferent activation (Oddo et al., 2017). Tactile stimulation was delivered through four pairs of intracutaneous needle electrodes inserted into the volar side of the second digit of the forepaw, 2–3 mm between pairs and 1 mm or less within a pair, as pulses of 0.5 mA and 0.14 ms duration. The four patterns mimic hypothetical afferent activation patterns for four types of mechanical skin-object interaction, lasted between 200 and 340 ms, and were delivered in pre-defined random order at intervals of about 1.8 s, each repeated 200 times.

Visual stimulation used four adjacent 7000 mcd light-emitting diodes in red, blue, green and yellow, positioned at 30° from the midline of the animal and 4–5 cm from the left eye. Four patterns were built as multicolour pulse trains containing two pulses of each colour, each pulse 5 ms long at 1 mW with 25 ms to the onset of the next, giving a total duration of 215 ms; each was repeated 200 times. In the visuo-tactile experiments the visual component was a non-patterned train of five pulses in which all four colours were activated together, beginning simultaneously with the tactile pattern.

High-speed imaging of the fingertip

The skin in contact with the plate was imaged from above, through the transparent plate, with a high-speed camera at 1000 frames/s. Frames were stored as 512 × 512 grayscale images, and a 300 × 300 pixel patch covering the contact region was used for the burst analysis; the field of view spanned the fingertip–plate contact, of the order of several millimetres across. For each pixel the frame-to-frame brightness defined a time-continuous movement signal sampled at 1000 Hz, and this signal was the input to the time–frequency analysis.

What that brightness physically reports matters for the interpretation of the results, because it is contact intensity and not displacement. Each pixel images the ridged contact band being sheared beneath it, which is what allows the parity of the harmonics at a pixel to be read as a signature of its position on the contact geometry.

Animals, anaesthesia and surgery

Papers I and IV used adult male Sprague-Dawley rats, weighing 320–430 g and 300–350 g respectively. Animals were sedated with isoflurane in air at 3%, for one to two minutes in Paper I and two minutes in Paper IV, and general anaesthesia was then induced with an intraperitoneal mixture of ketamine and xylazine. The induction doses differed between the two studies; the maintenance regime did not. Anaesthesia was maintained through a catheter in the femoral vein delivering Ringer acetate and glucose mixed with ketamine and xylazine in a 20:1 ratio, at approximately 5 mg/kg/h ketamine.

A craniotomy was made over the right somatosensory cortex, 2×2 mm over the forepaw area in Paper I and 4×4 mm spanning AP -2.0 to $+2.0$ and ML 1.5 – 5.0 relative to bregma in Paper IV, and a grounding electrode was placed in the muscle of the neck or back. After probe insertion the exposed cortex was covered with agarose, at 1% in Paper I and 0.5% in Paper IV, to prevent dehydration and reduce tissue movement.

Depth of anaesthesia was monitored throughout in two ways: by regularly confirming the absence of withdrawal reflexes to noxious pinch of the hind paw, and by continuously observing the irregular alternation of sleep spindles with more desynchronised epochs that characterises sleep (Niedermeyer and Lopes da Silva, 2005). The intention was a light but sufficient anaesthetic depth.

Neuropixels recordings

Extracellular recordings were made with Neuropixels silicon probes (Jun et al., 2017): version 2.0 in Paper I and 1.0 acquisition hardware in Paper IV, in the latter case with a PXIe-1071 chassis and PXI-6133 I/O (National Instruments) and export through SpikeGLX. Signals were sampled at 30 kHz per channel. In Paper IV the amplifier gain was 500 and the deepest 384 recording sites were used. Recordings were time-synchronised to the stimulator through the trigger input of the I/O module.

In Paper I the probe was inserted in the forepaw region of S1, at coordinates -1.0 to 1.0 mm relative to bregma. Of its 384 channels only the upper 180 were analysed, these being the ones within cortex, which in the rat is approximately 1.8 mm deep. Paper IV made eight vertically oriented insertions, one per animal and approximately one hour of recording each, targeted predominantly at regions outside the paw representation of S1, with one insertion at paw-S1 for comparison.

Anatomical assignment in Paper IV requires a caveat that is made explicitly in the paper. Insertion sites were referenced to bregma and assigned by overlaying the coordinates onto a cortical atlas (Enander and Jörntell, 2019); post-hoc histological reconstruction of the electrode tracks was not performed. The anatomical labels should be read as targeted regions, corroborated functionally by the latency and F1 distributions, not as histologically verified ones.

Analysis

Time–frequency decomposition

The transform derived in Paper II is used as the analytical engine of Paper III. Its construction begins from the observation that combining continuous wavelet transforms at several scales through a geometric mean yields a more localised time–frequency representation, and that this geometric mean is the estimator minimising the mean cross-entropy to the underlying energy distribution. Paper II takes the scale increment to zero, so that the product over discrete scales becomes a geometric product integral, and solves that integral analytically. The result is a single wavelet, the hyperlet, which delivers the super-resolution of an ensemble in one transform.

The decomposition re-expresses the signal rather than removing content from it, which is what makes it an alternative to band-pass filtering and not a refinement of it. It is also cheap: because it replaces an ensemble of convolutions with one, it is fast enough to be applied densely, which is what made a decomposition at every pixel of an imaging record tractable in Paper III. Both properties are relied on throughout.

Against them stands a trade-off. The hyperlet confines energy to a compact central patch at the price of a faint peripheral halo, and that halo inflates the classical Heisenberg-area metric even as the central lobe becomes cleaner and more information-dense. The transform is hence well suited to isolating short bursts from dense backgrounds and correspondingly less suited to confining long cyclical events within a narrow window.

In Paper III the per-pixel decomposition was evaluated on a 1–425 Hz grid in 1 Hz steps. For the chirp analyses, where the requirement is to follow a slow sweep instead of resolving a short transient, a conventional adaptive Morlet continuous wavelet transform was sufficient and was used instead.

Quantifying the resolution of a transform

Paper II evaluates the hyperlet against the short-time Fourier transform, the continuous wavelet transform and the fractional adaptive superlet transform on seven complementary metrics computed on the kernels and transforms: Heisenberg area, integrated side-lobe level, joint Rényi entropy of order 2, Gini sparsity index, mutual information, Wasserstein-2 distance to a centred isotropic Gaussian, and box-counting fractal dimension.

Because these metrics are abstract, an operational test was added: two identical superimposed signal bursts were gradually separated in time and in frequency, and the separation at which each method could resolve them as two events was recorded. This is the property the application actually needs.

Synthetic signals with known components were used to probe resolution across the spectrum, and signals with unknown, rapidly changing components were used to test whether the higher resolving power improves the description of events that were not designed to be found. Runtime was benchmarked on CPU with multithreaded convolutions, taking the median of eight trials per point.

Per-pixel harmonic analysis

For each pixel in Paper III, the fundamental and its low-order harmonics, from f through $5f$, were recovered from the response to a burst of about 300 ms. Pixels were then assigned to one of four classes according to which harmonics carried substantial energy: all, only odd, only even, or none. Two further quantities were derived across the contact: the rate at which harmonic energy decays across harmonic orders, and whether a pixel changes class between stimuli of different driving frequency.

This classification is a coarse, threshold-based summary of a continuous phenomenon, and is treated as such in the interpretation. Its purpose is to make a spatial pattern visible, not to assert that the four classes are natural kinds.

Spike sorting

Recordings were sorted with Kilosort (Pachitariu et al., 2024), version 2.5 in Paper I and version 4 in Paper IV. Sorted units were then curated manually in Phy on the basis of spike waveform shape, firing pattern and inter-spike interval violations. In Paper I units firing below 0.8 Hz, units confounded with stimulation artefacts, and units with more than 1% of their inter-spike intervals shorter than 2 ms were removed. In Paper IV the manual curation was applied to the units that had already passed the decoding criterion described below.

The hyperlet was not used as the spike-sorting front end, although that was the intention when Paper II was written. It demonstrates that the decomposition improves the separability of spike waveforms for a generic classifier, but a sorter is more than a representation with a classifier attached: automated unsupervised clustering, drift correction and an interface for error correction would all have to be built before it could replace an established one. The hyperlet is one part of such a system, not a substitute for it, which is why the recordings here were sorted conventionally.

Single-trial decoding

Papers I and IV both classify individual trials instead of comparing averaged responses. The reason is set out in the introduction: averaging removes the trial-to-trial variation in which cortical state-dependence lives, and the question of interest is not whether a mean

response differs but whether the identity of the input can be recovered from a single presentation.

The procedure is the same in outline. The discrete spike times of a unit are converted to a time-continuous signal by convolution with a Gaussian kernel; the resulting representations are reduced by principal component analysis; and a k -nearest-neighbour classifier, operating on Euclidean distances between data points, is asked which stimulus produced each trial. Performance is reported as a macro-F1 score against the chance level set by the number of classes, which in Paper IV is 12.5% for the eight tunes.

Paper IV specifies two hyperparameters explicitly. The Gaussian convolution width was selected per unit from a fixed grid of six biologically motivated values, 3, 5, 10, 15, 20 and 30 ms, by maximum mean cross-validated F1 across 100 train/test shuffles, so that each neuron's preferred temporal scale is reflected in the analysis. The number of neighbours k was tuned within each individual shuffle by a small grid search on the training fold only, never on the test fold. A unit was retained if its mean macro-F1 across the 100 stratified shuffles exceeded chance at $p < 0.05$, with the p -value computed as the fraction of shuffles at or below chance.

State-space characterisation

Paper I treats the population, and not the unit, as the object of analysis. Principal component analysis was applied to the activity distributions across the recorded population during the 300 ms window preceding stimulus onset, to define the state space of spontaneous activity. Sensory inputs were then treated as perturbations of that space, and the question asked was whether an input displaced the population state detectably, and whether different inputs produced different displacements. Decoding within that space used the same k -nearest-neighbour approach.

The number of retained components is a substantive choice, not a presentational one. Instead of reducing to a small fixed number, the analysis retained the components required to account for 95% of the variance of the spontaneous data, which in the five experiments meant between 27 and 55 components. A convolutional neural network was trained on the same data as an independent control, on the reasoning that a fundamentally different analysis converging on the same conclusion is stronger evidence than either alone.

Alternatives considered

The faculty asks that a thesis discuss the methods critically and weigh the alternatives that were not taken. Five choices in this work had defensible alternatives, and each was a trade.

Anaesthetised rather than awake recording

Awake recording is the obvious alternative and would remove the concern that anaesthesia distorts cortical dynamics. It was not chosen because the question here is whether sensory information propagates through the cortical network, and in an awake animal a sensory input can produce cortical activity indirectly, by triggering a behavioural response that itself drives the cortex (Bimbard et al., 2023; Musall et al., 2019; Stringer et al., 2019). Under light anaesthesia with low-intensity inputs, that route is largely closed, and what remains is closer to the passive propagation the thesis is about. The cost is real: ketamine and xylazine suppress cortical excitability and intracortical communication, so the persistence reported in Papers I and IV is likely to underestimate what an awake animal would show. The direction of that bias is at least known, which is why the finding is reported as conservative, not as anaesthesia-dependent.

Electrophysiology rather than calcium imaging

Calcium imaging can record from an order of magnitude more neurons than a Neuropixels probe, and would have strengthened any claim about population dimensionality by sheer sample size. It was not used because the temporal resolution of the indicator is far below that of spiking, and the estimate of complexity in Paper I depends on richness across time scales: the same paper shows the measured dimensionality collapsing when temporal resolution is degraded. Imaging would hence have traded away the variable being measured. Electrical recording keeps millisecond resolution at the cost of sampling a few tens of neurons per experiment.

Atlas targeting rather than histological reconstruction

Insertion sites in Paper IV were assigned by referencing coordinates to bregma and overlaying them on a cortical atlas, without reconstructing the electrode tracks histologically. Histology would have given anatomically verified labels and is the stronger method. What the present design offers instead is functional corroboration: response latencies and F1 distributions separate the insertion targeted at paw-S1 from the others in the direction the targeting predicts. The claims in Paper IV are correspondingly framed in terms of physiological distance from the direct thalamocortical input and not in terms of named areas, which is what the evidence supports.

Kilosort rather than a hyperlet-based front end

Papers I and IV were sorted with established versions of Kilosort (Pachitariu et al., 2024), although a hyperlet-based front end was the intended route when Paper II was written and is the reason the spike application appears in that paper at all. What the work established is that the gap between a better representation and a working sorter is wider than it looks: drift correction, unsupervised clustering and an interface for error correction would all have to be built first, each a research problem of its own, and doing so was a larger undertaking than these studies could carry. Using a widely adopted sorter also makes the unit yield comparable with other work. Re-sorting the Paper IV recordings with a hyperlet

front end is the natural test of whether the improvement demonstrated on hybrid data transfers to real recordings.

Linear reduction rather than nonlinear embedding

Principal component analysis is a linear method, and a range of nonlinear embeddings would produce lower-dimensional and often more visually orderly descriptions of the same population data. The argument for staying linear is set out in Paper I and is not that nonlinear methods fail. It is that they succeed too easily: many different embeddings work, they disagree, and each encodes an unverified assumption about the computational architecture of the circuit (Kristensen et al., 2024; Ramasamy and Jörntell, 2026; Spanne and Jörntell, 2015; Roland, 2023). A linear method retaining enough components to account for 95% of the variance makes fewer such assumptions and is less likely to discard structure because it does not fit a low-dimensional picture. The cost is that the resulting description is harder to visualise and less compact.

Use of generative AI

Generative AI (Claude, Anthropic) was used to draft and revise parts of the framework narrative and to prepare its typesetting. It was not used to generate, analyse or interpret experimental data, and no figures were produced with it. The script that generates the cover artwork was also written with its assistance; the artwork is a synthetic illustration and contains no data. All text has been reviewed and revised by the author.

Results

The papers are numbered as they are on the notification for the thesis defence, which is roughly the order in which the work was done. That order also happens to trace the argument. Paper I establishes the cortical phenomenon with the skin taken out of the loop. Paper II derives the analytical instrument. Paper III turns that instrument on the skin itself, where it is the analysis engine of every principal result. Paper IV then returns to the cortex with the skin reinstated, and its central result is only interpretable in the light of Paper III.

The instrument was meant to reach further than it did. The intention behind Paper II was that a hyperlet front end would replace band-pass filtering in the spike-sorting pipeline, and so serve the cortical recordings of Paper IV as well as the imaging of Paper III. Working the problem through showed that a spike sorter is a substantially larger undertaking than a better time–frequency representation, and that the hyperlet is one component of such a system and not a substitute for it. Paper IV was therefore sorted with an established sorter. That is a result of the work rather than an omission from it, and it is taken up in the discussion.

Paper I: High-dimensional cortical signals reveal rich bimodal and working memory-like representations among S1 neuron populations

The first paper asks how complex the signals of a cortical neuron population actually are. Tactile input was delivered as highly reproducible electrotactile spatiotemporal patterns through intracutaneous needle electrodes (Oddo et al., 2017), visual input as multicolour LED pulse trains, and in a subset of experiments the two were combined. Populations of neurons in primary somatosensory cortex were recorded with Neuropixels probes at millisecond resolution. The analysis first mapped the state space of spontaneous activity using principal components of the prestimulus window, then treated the sensory inputs as perturbations of that space and asked whether, and how specifically, they displaced it.

The spontaneous state space proved extremely high-dimensional: accounting for 95% of its variance required a large fraction of the components available, and that dimensionality collapsed when the temporal resolution of the recording was reduced. Tactile inputs

displaced the state reliably, but so did visual inputs, and the population separated fine nuances between individual visual patterns, in a cortical area conventionally treated as unimodal. When visual and tactile inputs were combined, the same tactile patterns produced different population covariance patterns than they did alone. Finally, input-specific displacements of the state persisted well after the stimulation had terminated, leaving information about which input had just occurred resident in the ongoing activity.

The analysis drew on five experiments with 36 to 70 simultaneously isolated S1 neurons each, 274 neurons in total. Accounting for 95% of the variance of the spontaneous activity required 50, 27, 55, 42 and 28 components in the five experiments, which is 68, 81, 79, 62 and 78% of the number needed to account for all of it. The state space is hence high-dimensional relative to the size of the population recorded, and that dimensionality collapsed when the temporal resolution of the data was reduced.

In the illustrated experiment, decoding separated evoked from spontaneous activity with an F1 score of 93% for tactile input and 81% for visual input, against a chance level of 50%; shuffling the labels collapsed performance to a narrow distribution around chance. Across the three visuo-tactile experiments and four tactile patterns, F1 scores for separating the same tactile pattern with and without accompanying visual input averaged 0.73 (SD 0.027) against 0.50 for shuffled data. A convolutional network trained on the same data as an independent control confirmed that evoked and spontaneous distributions differed, though at a lower level than the full set of components achieved with the nearest-neighbour approach.

Information about which tactile pattern had been delivered remained decodable throughout the 1000 ms window tested. F1 scores in successive post-stimulus windows stayed significantly above the pre-stimulus level, at 0.363 (SD 0.054) in the 201–400 ms window and still 0.329 (SD 0.028) at 601–800 ms.

What this contributes to the argument

Because the electrotactile interface bypasses mechanotransduction, none of the complexity seen here can have been introduced by the skin. Paper I fixes a baseline: the cortical representation of a tactile input is high-dimensional, is not confined to its nominal modality, and outlives its stimulus, all without any peripheral mechanics in the loop. Everything that follows is an attempt to find out what the skin adds to that, and what survives when it is put back.

Paper II: Super-resolution time–frequency decomposition with hyperlets for neural spike analysis

Answering that question required an instrument. Time–frequency decomposition is the standard way to unmix signals from multiple sources with distinct characteristics, but ex-




Figure 2. Summary of results from Paper I. Principal component analysis of the S1 neuron population, showing the dimensionality of the spontaneous state space and the separation of evoked from spontaneous activity. Reproduced from Paper I.

isting tools localise short, overlapping bursts poorly. The problem is acute in extracellular recordings from low-impedance multi-electrode arrays, where each neuron contributes a characteristic waveform mixed with the signals of hundreds of other sources. Existing super-resolution approaches, notably the superlet transform and its fractional adaptive variant, combine several continuous wavelet transforms at different scales through a geometric mean. This works, but it is computationally expensive, and it can suppress a feature whenever any one scale returns a near-zero response, a real risk when the scales are widely spaced.

Paper II removes the discretisation. By taking the scale increment to zero, the geometric mean over a set of wavelets becomes a geometric product integral, which can be solved analytically into a single wavelet: the hyperlet. The resulting transform retains the minimum mean cross-entropy optimality of the discrete construction while collapsing an ensemble of convolutions into one.

The transform was evaluated against the short-time Fourier transform, the continuous wavelet transform and the fractional adaptive superlet transform on seven metrics, and on an operational test in which two identical superimposed bursts were gradually separated in time and in frequency, down to separations of 2 ms and 2 Hz. The hyperlet resolved the pair at smaller separations than the alternatives while requiring roughly a twentieth of the operations of the fractional adaptive superlet transform.

The trade-off is visible in the metrics themselves and is worth stating plainly. The hyperlet confines energy to a compact central patch at the cost of a faint peripheral halo, and that halo inflates the classical Heisenberg area even as the central lobe becomes cleaner and more information-dense. The transform is hence suited to isolating discrete events from dense backgrounds, and correspondingly less suited to confining long cyclical events within a narrow window.

Table 1. Kernel and transform quality metrics, reproduced from Table 3 of Paper II. Heisenberg area (HA), integrated side-lobe level (ISL), joint Rényi entropy H_2 , Gini sparsity index, mutual information $I(T; F)$, Wasserstein-2 distance W_2 to a centred unit Gaussian, and box-count fractal dimension (FD). Lower is better except for ISL, I and Gini. The three best values in each column are in bold. HA and ISL are undefined for FASLT.

Method	Parameters	HA	ISL	H_2	Gini	$I(T; F)$	W_2	FD
STFT	$W = 0.032$	1.026	26.64	8.73	0.998964	2.21e-3	1.231	1.664
STFT	$W = 0.128$	1.026	32.66	9.78	0.997119	2.91e-5	1.254	1.655
STFT	$W = 0.512$	1.026	38.69	11.14	0.988750	-1.71e-7	1.228	1.659
CWT	$c = 5$	1.000	27.79	8.80	0.998887	9.08e-3	1.239	1.710
CWT	$c = 16$	1.000	32.84	9.64	0.997432	1.44e-2	1.258	1.676
CWT	$c = 64$	1.000	38.86	10.99	0.990066	1.53e-2	1.235	1.676
HLT	$\zeta = 1.0, n = 0.99$	299	11.77	8.85	0.998794	7.37e-3	1.026	1.840
HLT	$\zeta = 5.0, n = 0.99$	144	19.82	8.78	0.998779	1.83e-2	1.063	1.829
HLT	$\zeta = 0.35, n = 0.5$	1.59e4	39.18	8.73	0.998069	3.44e-2	1.149	1.717
FASLT	$c_1 = 3, o = [1, 22]$	-	-	9.25	0.998300	2.63e-3	1.254	1.659
FASLT	$c_1 = 5, o = [1, 30]$	-	-	9.73	0.997276	7.25e-3	1.254	1.658
FASLT	$c_1 = 1, o = [5, 40]$	-	-	8.94	0.998760	1.04e-3	1.250	1.656

The hyperlet takes the best value on four of the seven metrics (Table 1): the highest integrated side-lobe level, the highest mutual information, the lowest Wasserstein-2 distance to a centred Gaussian, and the equal-lowest joint Rényi entropy, where it ties with the narrowest short-time Fourier window. Read together these say that almost all the energy sits in one tight ridge whose shape is near-Gaussian and therefore close to optimally compact, with little ambiguity between when an event occurred and at what frequency.

The one metric it loses on is, however, the revealing one. Its Heisenberg area is two to four orders of magnitude worse than the classical transforms, and the reason is the spectral-sharpening kernel: it produces a narrow main ridge but pushes a very low-level skirt far into the periphery. That skirt carries almost no energy, so it barely affects entropy or sparsity, but it inflates both σ_t and σ_f , and since the Heisenberg area is their product it explodes. The classical metric penalises a halo heavily even when the core is pin-sharp. Which of these matters depends on the task. Where discrete events must be isolated from dense backgrounds, algorithms look at peak energy cells and their immediate neighbours, the halo is irrelevant, and clean narrow peaks with low side lobes improve detection directly.

Two applications follow. For the neural case, ten synthetic spike clusters with known ground truth were inserted into real extracellular cortical recordings across eight channels, and twenty occurrences of each were classified twice: once from the band-pass filtered voltage, once from the hyperlet decomposition. Six principal components accounted for about 80% of the variance of the decomposed data against 60% for the filtered voltage, and the filtered data needed 17 components to reach the level the decomposition reached with six. Nearest-neighbour classification on six components was 45% more accurate, in relative terms, on the decomposition. Note that six components is the number the current version of Kilosort uses, so the comparison is made at the operating point of an existing sorter.

The likely reason is that band-pass filtering distorts a waveform according to what else is present: superimposing an oscillation below the filter cut-off changes the filtered result, so the same spike can fail to be recognised as the same spike depending on whether a field potential happened to coincide with it. A decomposition does not have that property because it removes nothing.

The second application is acoustic, decomposing an orchestral recording, where the hyperlet again gave sharper temporal localisation and greater frequency separation than the alternatives. Given that human frequency and time discrimination operate in the range of 2–10 Hz and 2–10 ms, the improvement is of a size that matters for auditory work.

This is a demonstration of what the representation affords, not a spike sorting pipeline; the distinction is returned to in the discussion.




Figure 3. Summary of results from Paper II. Decomposition of a series of signal events by STFT, CWT, FASLT and the hyperlet transform at matched parameters. Reproduced from Paper II.

What this contributes to the argument

Band-pass filtering removes information; it alters the content of the signal in order to make it tractable. A time–frequency decomposition does something categorically different: it expresses the same information in another vocabulary, one a downstream classifier can read more easily. Paper II establishes that the second option is computationally affordable at the scale of real recordings and that it measurably improves what can be recovered. Paper III rests on it directly: without a decomposition cheap enough to run at every pixel of an imaging record, the spatial map of harmonic content could not have been made at all. The intended second application, a spike-sorting front end for the cortical recordings, is where the same paper establishes the limit of what a transform alone can do. Beyond both, and as the discussion will argue, re-expressing instead of discarding is also a description of what the biological system itself does at each stage.

Paper III: Spatially distributed, stimulus-dependent harmonic generation in the fingertip skin during interaction with a haptic display

The mechanical compliance of the fingertip shapes vibrotactile information before any receptor transduces it, but how that transformation is organised across the contact area was unknown. Bulk impedance measurements and single-point recordings average over the contact and cannot say whether the transformation is uniform, nor whether what happens at a given location is a fixed property of the local tissue or depends on the stimulus.

A fingertip rested passively on the plexiglass surface of a custom haptic display delivering tangential vibratory tone bursts. The drive was pre-compensated using a co-located accelerometer. Critically, that accelerometer, measuring the rigid plate alone, showed no harmonics of the driving frequency. Any harmonic content in the skin response hence cannot originate in the stimulus or the stimulator. The moving skin was imaged from above at 1000 frames per second, resolving the papillary ridges, and the brightness signal of every pixel was decomposed with the transform of Paper II.

Four results follow. First, the skin generates harmonics: local movement carries substantial energy at integer multiples of a frequency that was not present in the input. Second, that harmonic content is spatially organised into coherent domains aligned with the papillary ridges, in which the movement is dominated by all, only odd, only even, or no harmonics. Third, the harmonic class at a given location is not an anatomical property: many locations switch between odd-dominant and even-dominant behaviour when the driving frequency changes, and the rate at which harmonic energy decays across orders varies markedly across the surface. Fourth, the skin low-pass filters the input, and the attenuation applied to any one frequency depends on which other frequencies are simultaneously present.

The harmonic analysis concentrates on three representative burst frequencies, 75, 100 and 126 Hz, drawn from a series stepped between 75 and 475 Hz. Pixels a fraction of a millimetre apart behave qualitatively differently: one may be dominated by the fundamental with an orderly decaying harmonic series while its neighbour expresses harmonics that exceed the fundamental.

The parity switching between stimuli is not scattered. Locations that were odd-dominant under the 75 Hz stimulus became even-dominant under the 100 Hz stimulus, and separately locations that were even-dominant became odd-dominant, and both switching populations form spatially coherent bands along the ridges, not isolated pixels. Coherence of that kind is the argument against classification noise. Decay of harmonic energy across orders was summarised by fitting an exponential and taking the decay coefficient; the resulting maps were again spatially structured, with the most rapidly decaying behaviour concentrated along particular ridge domains, and they differed between stimuli.

The filtering result is quantitative. Presenting two chirps of different sweep rates simultaneously, the slowly swept component was attenuated approximately 20% less than the rapidly swept component over the 100–225 Hz range, so the attenuation applied to a frequency depends on the context in which it arrives.

Classification used the relative power in ± 20 Hz bands around the fundamental and its second and third harmonics. A pixel was labelled as expressing all harmonics when the series decayed in order, with the fundamental above the second above the third; odd harmonics when the fundamental exceeded the third and the third exceeded the second; even harmonics when the second dominated; and no harmonics otherwise. The resulting maps were highly structured, not random. The classes formed coherent elongated domains running parallel to the papillary ridges, with bands of odd-dominant, even-dominant and full-harmonic behaviour interleaved, and the same banded organisation appeared at every stimulus frequency tested. The pattern was not, however, identical across stimuli: both the relative areas occupied by each class and the precise placement of the bands shifted with the driving frequency (Figure 4).

These observations admit a direct mechanical reading. Because brightness reports the intensity of skin–plate contact, each pixel images the ridged contact band being sheared beneath it: a pixel over a contact maximum sits at a turning point of that motion and returns even harmonics, one on a contact slope returns odd harmonics, and the balance between them shifts as the shear amplitude changes with frequency. The parity map is therefore a signature of the sheared ridge contact geometry and not a map of distinct mechanical regimes. The harmonic content itself, however, is real, and it is the skin that introduces it.

What this contributes to the argument

This is the first of the two junctions at which touch is supposed to be simplified, and it is not a simplification. The skin adds spectral components that were not in the stimulus, distributes them non-uniformly across the contact according to its own microarchitecture,

Spatial distribution of skin harmonics across the fingertip




Figure 4. Summary of results from Paper III. Spatial distribution of harmonic classes across the fingertip contact for two driving frequencies, showing the ridge-aligned domains and their reorganisation with the stimulus. Reproduced from Paper III.

and rearranges that distribution when the stimulus changes. Whatever a mechanoreceptor receives is hence a locally and contingently transformed version of the event, not a scaled copy of it. Frequency does not survive the interface as an independent variable, which is what makes the central result of Paper IV intelligible. And part of what has been read as neural computation is, to some degree, already performed by tissue mechanics.

Paper IV: Highly resolved representations of dynamic haptic inputs in non-primary sensory cortical neuron representations

Earlier work from this environment showed that neurons far from primary somatosensory cortex, even in VI, can separate fine nuances of tactile input patterns (Enander and Jörn-tell, 2019; Enander et al., 2019). Those studies, like Paper I, used an electrotactile interface, which delivers highly reproducible patterns directly to the afferents but bypasses mechanotransduction entirely. Paper IV asks whether the same holds when the skin is left in the loop.

A vibratory plate was placed in contact with four digits of the rat forepaw and delivered eight stimulation patterns, or “tunes”, each a series of Gaussian pulses at specific frequencies, with consecutive pulses differing in frequency and with large partial overlaps in the frequencies used across tunes. Peak-to-peak amplitude was held constant while pulse durations and intervals varied between patterns, so the tunes differed in their temporal structure, not in their spectral content alone. Neuropixels insertions were made in separate animals under light anaesthesia, targeted predominantly at regions outside the paw representation of S₁, with one insertion at paw-S₁ for comparison. Instead of averaging responses, each trial was represented as a Gaussian-convolved spike train and classified with a PCA-kNN decoder, with per-unit macro-F₁ tested against the 12.5% chance level set by the eight balanced classes.

Three findings stand out. Large numbers of neurons separated the tunes regardless of recording location, including at sites whose response latencies place them at substantial physiological distance from the direct sensory thalamocortical input. No neuron showed a frequency-specific response over the range tested, roughly 80 to 300 Hz; what the neurons tracked instead was the acceleration of the input magnitude, and the same tune played back at different speeds was encoded as more or less disjunct instead of as a scaled version of itself. Finally, the combined activity of small populations of neighbouring neurons supported time-resolved decoding of which pattern was currently playing, and that pattern-specific decoding persisted for hundreds of milliseconds past stimulus offset, most markedly in populations recorded closer to the direct thalamocortical input.

Across the eight experiments, 410 units passed the per-unit significance criterion. Because different neurons may carry decodable information on different temporal scales, the Gaussian convolution width was selected per unit from a six-point grid committed in advance; the median preferred width was 12 ms, consistent with the 12–40 ms range previously reported for S₁ cortical neurons responding to electrotactile patterns (Oddo et al., 2017). The preferred width was uncorrelated with F₁, which argues against the selection inflating scores by fitting noise.

The experiment targeted at paw-S₁ behaved as its position predicts. It showed the shortest response latencies, 35–50 ms, the largest number of units passing the inclusion criterion, and the largest number of units with high F₁ scores. Among units scoring 0.3 and above there was an overall inverse relationship between response latency and F₁; below that level no clear relationship was apparent. This is the functional evidence standing in for histological verification of how directly a recording site receives thalamocortical input.

Time-resolved decoding was computed in 50 ms steps within populations of ten neurons. Pattern-specific information persisted past stimulus offset for different durations at different sites: throughout the roughly 750 ms post-offset window displayed in the paw-S₁ experiment, above the +2 SD baseline band until approximately 250 ms past offset in another, and above the +1 SD band until approximately 550 ms past offset in a third. The gradient follows proximity to the direct thalamocortical input.




Figure 5. Summary of results from Paper IV. Time-resolved decoding within small populations, showing pattern-specific information persisting beyond stimulus offset. Reproduced from Paper IV.

What this contributes to the argument

The absence of frequency tuning is the result that most needs explaining, and Paper III explains it. If the skin redistributes energy across harmonics in a way that is spatially heterogeneous and contingent on the stimulus, then no stable frequency variable survives the interface for a cortical neuron to be tuned to. The absence is not a failure of the cortex to extract frequency; it is a property of what the cortex is given. Paper IV also closes the loop on Paper I: the high-dimensional, distributed, persistent representation seen with the skin bypassed is, however, still there when the skin is reinstated, so the periphery is not the source of that complexity; it is an additional, structured contribution to it.

Discussion

The conventional account of touch is one of progressive reduction. Skin deformation is transduced, labelled channels carry frequency and force centrally, and the primary somatosensory cortex extracts the features from which perception is assembled. Each stage is assumed to discard what is not needed. The four papers in this thesis examine the two junctions at which most of that discarding is supposed to happen, the mechanics of the skin and the cortical representation of what the skin sends, and neither behaves as the account requires. What follows is organised around the three claims this supports: that the skin expands rather than reduces, in the spatiotemporal sense set out below, that the cortex distributes and sustains where the account has it localising and discarding, and that seeing either of these requires measuring the signal before any step that discards part of it.

The skin expands rather than reduces

Harmonic generation as an expansion

The fingertip in Paper III returns spectral content that the stimulus did not contain. Because the accelerometer measuring the driven plate showed no harmonics of the drive, the harmonics observed in the movement of the skin can only have been introduced by the tissue itself. That the finger–surface interaction is nonlinear is not new: the interaction force during sliding contact carries harmonics of the fundamental imposed by a grating, with amplitudes decaying at a non-integer rate (Wiertlewski et al., 2011a,b), and the lateral mechanical impedance of the fingertip departs from a linear response across the frequencies relevant to touch (Wiertlewski and Hayward, 2012).

What Paper III adds is a spatial account. Bulk impedance measurements and single-point recordings return one number where the question requires a map, and an averaged transformation is compatible with a great many underlying arrangements. Resolving the movement at every pixel of the contact shows that the arrangement is not uniform: harmonic content is organised into coherent domains that follow the papillary ridges, within which movement is dominated by all, only odd, only even, or no harmonics of the drive.

The sense in which this counts as an expansion needs stating exactly, because in one sense it cannot be one. Spatially the tissue can only lose, since a continuous mechanical field

sampled by a finite receptor population has fewer places to be read than to exist. The expansion is in frequency and in time. The plate was driven within a narrow band; what the tissue presents to the receptors covers a wider one, covers it differently from one place on the contact to the next, and differs again from moment to moment. Put in the terms of Chapter 1, the coordinates of the plenoptic domain are fixed at six, but the region of that domain the signal occupies is larger downstream of the tissue than upstream of it. Hence the features that survive to the receptors are kinematic invariants, activation patterns evolving in time, not static ones (Jörntell et al., 2026). Note that no information has been created here; a nonlinear process has redistributed energy that was already present. But describing what arrives at the receptor sheet takes more than describing what was delivered to the plate, and that is the reverse of what the funnel account predicts of this stage.

Stimulus-dependence and the loss of frequency as a variable

The consequence that carries furthest is not the harmonics themselves but their contingency. A location returning even harmonics at one driving frequency returns odd harmonics at another, and the attenuation the skin applies to any one frequency depends on which other frequencies accompany it. Neither observation is available to a measurement that presents one frequency at a time and averages over the contact.

Frequency is hence not conserved across the interface as an independent variable. What reaches a receptor is jointly determined by the input spectrum, the local contact geometry and the shear amplitude, and those contributions cannot be separated by recording from the afferent alone. This has a direct implication for how afferent tuning should be read. A tuning curve measured at a mechanoreceptor is a property of the receptor together with the tissue around it, and since the tissue's contribution is itself stimulus-dependent, the tuning need not be stable across the stimuli an animal actually encounters.

What I measured is optically resolved movement at the surface of the skin, not strain at the depth where the receptors sit, and pixel brightness reports contact intensity, not displacement. The parity map is read here as a signature of the sheared ridge contact geometry, not as a map of distinct mechanical regimes. What is established is that the tissue transforms the input nonlinearly and non-uniformly; what is inferred, and remains to be shown directly, is that the afferent population inherits that transformation.

Mechanics as part of the computation

If the tissue performs a transformation that is nonlinear, spatially structured and contingent on the stimulus, the boundary usually drawn between mechanics and computation becomes hard to defend. This aligns with the argument that skin tissue dynamics are integral to tactile sensing and not a nuisance preceding it (Rongala et al., 2024), and with evidence that vibrations propagating across the hand shape the afferent population that is recruited (Manfredi et al., 2012; Shao et al., 2016).

It also supplies an empirical counterpart to a modelling result. A cortical-like recurrent network learning autonomously through Hebbian rules acquired representations that generalised to unseen inputs, and retained that ability after substantial loss of sensors or internal connections, only when its input carried the dependencies that mechanical coupling produces; training on inputs stripped of those dependencies degraded it (Rongala and Jörntell, 2026). Note that Paper III establishes dependencies of the required kind as real, measurable and richer than the model assumed, since they are reorganised by the stimulus, not fixed by the tissue.

There is a design implication that follows without any neuroscience. If the fidelity of a haptic percept depends on a medium that redistributes energy across frequency and space in a manner contingent on the whole stimulus, then actuator density and precision are not the only limits on a display, and compensating for the actuator while ignoring the skin corrects the smaller of the two transformations.

The cortex distributes and sustains rather than localises

Fine haptic detail far from the primary area

Paper IV finds large populations of neurons separating fine-nuanced vibrotactile patterns at sites whose response latencies place them well away from the direct thalamocortical input. This extends to genuinely haptic stimulation what had previously been shown with electrotactile patterns delivered directly to the afferents (Enander and Jörntell, 2019; Enander et al., 2019; Oddo et al., 2017; Wahlbom et al., 2021), and it does so with mechanotransduction intact, which was the one step those earlier studies bypassed.

The finding belongs to a broader revision of cortical parcellation (Wallace et al., 2004; Ghazanfar and Schroeder, 2006; Olcese et al., 2013). Responses attributable to a given modality are recorded well outside its nominal primary area (de Lafuente and Romo, 2006), focal lesions impair processing in neurons at a distance from the lesion (Wahlbom et al., 2019), and the anatomical substrate is better documented than it was: individual corticocortical axons arborise far more widely than the textbook picture allows (Szeier et al., 2025), and cortico-thalamo-cortical loops provide a second route.

The strongest objection, however, is that distant responses reflect behaviour instead of sensation. Sound-evoked activity in mouse visual cortex has been argued to originate in the behavioural response the sound triggers (Bimbard et al., 2023), and movement accounts for much of the variance in single-trial cortical activity (Musall et al., 2019; Stringer et al., 2019). Two things limit the objection here. First, responses that differ according to which stimulus was presented imply that whatever mediates them is itself stimulus-specific, so the objection relocates the specificity instead of removing it. Second, the preparation was chosen to close that route: light anaesthesia and low-intensity inputs leave little scope for a be-

havioural reaction to intervene. What is observed is closer to passive propagation through the network than to a sensorimotor loop.

Why no neuron is tuned to frequency

The negative result of Paper IV is where the peripheral and central halves of this thesis meet. No neuron showed a frequency-specific response over the range tested, roughly 80 to 300 Hz; what the neurons tracked was the acceleration of input magnitude, and the same tune played at different speeds was represented as disjunct and not as a scaled version of itself.

Read alone, this is a puzzle, and one that invites the wrong explanation, that the cortex has discarded frequency. Read after Paper III it is close to a prediction. If the skin redistributes energy across harmonics in a manner that varies with location and with the rest of the stimulus, then no stable frequency variable survives the interface for a neuron to be tuned to. The absence is a property of what the cortex is given, not a failure of the cortex to extract what it was given. That the neurons instead tracked the acceleration of input magnitude is consistent with local shear forces dominating the afferent drive, which is the quantity Paper III measures.

This bears on the strongest contrary result. Neurons in mouse forelimb somatosensory cortex have been reported to encode vibration frequency through spike rates tuned to a preferred value, in a manner explicitly likened to pitch representation in auditory cortex (Prsa et al., 2019). That finding was obtained with sustained sinusoidal vibration whose frequency was stepped between conditions, and the disagreement is less about the cortex than about the stimulus. A long-duration pure tone is a form of skin stimulation an animal will never encounter, and it holds fixed the very thing that varies in natural contact. An experiment of that design can only return an answer to the question it asks: present a stationary frequency and a tuning curve is what there is to find. What the relevant afferents actually encounter, recorded in awake and freely moving animals, is the animal's own movement and low-amplitude vibration arriving from the environment, across a population diverse in tuning and sensitivity (Turecek and Ginty, 2024). The tunes used in Paper IV were built the opposite way, with the frequency changing from tone to tone and the tunes overlapping in spectral content, so that frequency alone could not identify the stimulus.

Paper III is a human fingertip resting passively and Paper IV is the rat forepaw under a vibrating plate. The two are linked by using the same class of haptic display, so the mechanical input is of the same kind, but Paper III does not measure Paper IV's periphery, and the inference is one of mechanism, not of measurement. The absence of frequency tuning would in any case be overdetermined. Successive nonlinear recombination along the pathway would degrade a frequency label even if the skin were linear: tonic and phasic afferent signals are fused in the cuneate nucleus, where subsets of haptic kinematic invari-

ants are segregated (Jörntell et al., 2014; Bengtsson et al., 2013; Hayward et al., 2014), and individual thalamic neurons already represent input from both rapidly and slowly adapting mechanoreceptors (Emanuel et al., 2021). Paper III shows that the degradation begins earlier than any of these stages, not that the later stages are irrelevant.

Persistence and a substrate for working memory

Papers I and IV converge on the same phenomenon from opposite directions. Input-specific information outlives its stimulus: by hundreds of milliseconds in non-primary cortex under genuine haptic drive, and on a comparable or longer timescale in S1 under electrotactile drive. The finding survives a change of preparation, of stimulation route and of recording site, which suggests it is a property of the cortical network and not of either protocol.

The behaviour matches standard definitions of working memory closely enough to be worth naming as such. The classical criterion, that neurons remain active when the information they code for is no longer present (Funahashi et al., 1989; Goldman-Rakic, 1996), is satisfied, and the phenomenon is described in comparable terms across a wide range of preparations (Haque et al., 2021). It is also what would be expected of a recurrent network that is not noise-free: a perturbation of a recursive system persists until noise dissipates it (Faisal et al., 2008), so a finite decay is what the account predicts, not an anomaly. Related observations point the same way, with sensory input followed by internal replay that gradually alters population responses (Marks and Goard, 2021), and with stable representations emerging from episodic activity that is itself variable (Xia et al., 2021).

The theoretical significance is what makes this more than a curiosity. Cross-modal invariants have no ground truth in physics and must be learned, retained and applied only when relevant, which requires a network able to hold information in the absence of the input that produced it; only the thalamocortical system, with recursive loops of in principle unbounded length, appears to have the necessary infrastructure (Ramasamy and Jörntell, 2026). If such invariants suffuse the cortical network and do not reside in identifiable convergence zones, working-memory-like persistence should be detectable in ordinary sensory cortical populations, not only in association areas. Papers I and IV report exactly that, in S1 and in cortex remote from S1 respectively.

Cross-modal structure in a unimodal area

Paper I finds S1 separating fine nuances between individual visual patterns, and finds tactile processing altered when visual input accompanies it. Similar crossmodal associations have been reported in cortical somatosensory cells (Zhou et al., 2020), and the case for a neocortex that is multisensory in principle, not by exception has been made from several directions (Ghazanfar and Schroeder, 2006; Wallace et al., 2004; Olcese et al., 2013).

Together with the demonstration that a given tactile input is interpreted differently depending on the internal state at the moment it arrives (Norrlid et al., 2021; Etemadi et al., 2022, 2023), and that cortical processing varies with motivational state (Matteucci et al., 2022), this is difficult to reconcile with a modality-specific primary area. The position this thesis takes is that functional localisation is not so much wrong as underdetermined. A region can have a reliable and privileged input, giving the fastest and strongest responses to its own modality, while carrying detailed information about much else. What the localisation account gets wrong is the inference from privileged input to exclusive content. On this reading the right question is not which area represents touch, but how a network that represents a great deal at once keeps its representations usable.

What the measurement determines

Re-expression versus reduction

The same methodological choice recurs in all four papers. Paper II replaces band-pass filtering, which removes information, with a decomposition that re-expresses the same information in a basis a classifier can read. Paper III replaces bulk and single-point mechanical measurement with a per-pixel map. Papers I and IV replace response averaging with single-trial decoding, and Paper I declines to reduce the population data to a few components, retaining instead the dimensions needed to account for 95% of the variance. In each case the structure reported here would have been invisible under the conventional alternative.

Stating this as a principle invites an obvious objection: a method that preserves more structure also preserves more noise, and the burden falls on the analysis to show that what survived is systematic. Each paper discharges that burden differently, and it is worth being explicit about how. Paper III has the harmonic-free accelerometer, which establishes that the harmonics cannot have come from the stimulus; the ridge alignment of the domains is a further constraint, since noise has no reason to organise itself along the skin's microarchitecture. Paper IV tests every unit against a shuffled null across 100 stratified train/test splits and tunes its hyperparameters on training folds only, so that decoding above chance cannot be an artefact of fitting. Paper I adds a convolutional network as an independent control on the principal-component analysis, on the reasoning that two dissimilar analyses agreeing is stronger evidence than either alone. Paper II is benchmarked against established transforms on seven metrics and, more usefully, on an operational test of when two superimposed bursts can be told apart.

The cost of dimensionality reduction

The dominant approach to population recordings is to find an efficient low-dimensional embedding. The objection developed in Paper I is not that such embeddings fail but that many of them succeed, return different answers, and each carries an unverified assumption about the computational architecture of the circuit (Jazayeri and Ostojic, 2021; Ramasamy and Jörntell, 2026; Spanne and Jörntell, 2015). A recording with enough independent degrees of freedom will have some projection that tracks almost any variable one cares to test it against, which is why the constraints imposed on the embedding matter more than the fit it achieves (Kristensen et al., 2024; Roland, 2023).

The argument should be kept narrow, because it is easy to overstate. Dimensionality reduction is not a mistake; this thesis uses principal component analysis throughout, and in some circuits a low-dimensional description can be characterised with real precision (Chaudhuri et al., 2019; Khona and Fiete, 2022). That precision should not be mistaken for evidence that the underlying system is low-dimensional. It is a calculated convenience: a reduction accepted in order to gain some quantitative purchase on a complexity whose true extent is unknown. The claim is that the number of retained dimensions is a substantive scientific choice, not a display convenience. Paper I shows how easily it is made by accident: the measured dimensionality of the same data collapses when the temporal resolution of the recording is degraded, so a methodological decision about binning implicitly determines a priori a conclusion about the complexity of the circuit.

A related point applies to what counts as noise. Trial-to-trial variability in cortical responses is routinely averaged away, but variability of comparable character appears throughout behaviour, in movement (Stergiou and Decker, 2011; Cavanaugh et al., 2017; Wu et al., 2014), in force production (Nagamori et al., 2021) and in speech (Tang et al., 2022), and it is regulated, not merely tolerated. Treating the neural counterpart as a nuisance term assumes an answer to a question that has not been asked.

Situating the work: the plenhaptic function

Described as a map between two continua of moving points, a haptic interaction needs unboundedly many coordinates, and any organism must hence be sensitive to a drastically reduced subset of what could be measured (Hayward, 2011). Asked at the skin instead, the same function becomes six-coordinate, frequency, amplitude, time and a position in the three-dimensional tissue, with kinematic invariants as its constituents (Ramasamy and Jörntell, 2026).

Paper III samples that object empirically: frequency and amplitude as functions of time and of location across the contact. The sampling is partial, resolving the two-dimensional surface, not the tissue volume in which the receptors actually sit. What it shows is that the first reduction step does not reduce. The skin does not hand the receptors a subset of the

plenahaptic function; it hands them a transformed version of the input, spectrally larger than what it received and organised by its own microarchitecture.

Where, then, does the reduction occur? The two junctions this thesis measures were chosen because they are where the conventional account places most of it, and neither performs it. That leaves the relay stages, and there the existing evidence already points somewhere specific: subsets of theoretically defined haptic kinematic invariants have been found segregated in the cuneate nucleus (Jörntell et al., 2014), a reduction with a basis in the physics of contact and of a kind neither stage measured here performs. Read together with the results here, the picture that emerges is not a funnel narrowing steadily from skin to cortex but something more localised: an expansion at the periphery, a selective extraction of invariants in the brainstem, and a high-dimensional, distributed representation at the far end.

Methodological considerations

Anaesthesia

Papers I and IV both used ketamine and xylazine. The positive argument has been given: the question concerns passive propagation of sensory information through the cortical network, and an awake preparation admits an indirect route in which the stimulus drives behaviour and the behaviour drives the cortex (Bimbard et al., 2023; Musall et al., 2019; Stringer et al., 2019). Light anaesthesia with low-intensity inputs largely closes that route.

The choice of agent also matters. The agent was chosen for its property of preserving patterns of sequential neuronal recruitment across diverse forms of UP state (Luczak and Barthó, 2012), which is the justification given in Paper I. Depth was monitored continuously by the absence of withdrawal reflexes and by the characteristic alternation of spindles with desynchronised epochs (Niedermeyer and Lopes da Silva, 2005).

The caveat is real and should not be softened. Ketamine and xylazine suppress cortical excitability and intracortical communication, and cortical dynamics differ between anaesthetised and waking states in ways that are directly relevant to sensory processing (Wright et al., 2021). The persistence reported in Papers I and IV is most likely a conservative estimate of what an awake animal would show, and not an artefact of the anaesthetic. That the bias runs in a known direction is what makes the finding usable; it is not a substitute for the awake experiment, which remains the obvious next test.

Human fingertip and rat forepaw

Paper III is a human fingertip and Papers I and IV are the rat forepaw, and the thesis draws a mechanistic link across that gap. The link is that Papers III and IV used the same class of haptic display, so the mechanical input is of the same kind, delivered through glabrous skin in both cases. What is not claimed is that Paper III measures the mechanics of the rat forepaw. The ridge architecture, contact area and tissue dimensions all differ, and the parity domains reported in Paper III are specific to the geometry that produced them. The transferable claim is the general one: that a compliant ridged contact under tangential drive generates stimulus-dependent harmonic structure, and so that frequency is not conserved across the interface. Demonstrating this in the rat forepaw directly would remove the inference, and is a well-defined experiment.

Statistical reach of Paper III

Note that Paper III reports one subject. The findings were internally reproducible, consistent across the three stimulus frequencies analysed and spatially coherent across the contact, which is what licenses treating them as a mechanism demonstration. The across-subject distribution of harmonic and filtering behaviour is unestablished, and there are obvious reasons to expect it to vary: skin hydration, ridge geometry, contact force and finger size all differ between individuals. A cohort study would need to establish that harmonic generation occurs in every subject, that the domain organisation follows the ridges in every subject, and that parity switching with driving frequency is general, not idiosyncratic. The first of those seems likely on mechanical grounds; the third is the one that most needs testing, because it carries the argument about frequency.

What Paper IV cannot test

The single-insertion-per-animal design demonstrates that fine haptic information is locally decodable at each recorded site, including at sites remote from the direct thalamocortical input. It cannot show that those local representations share a common format. Training a decoder on activity from one cortical site and applying it without retraining to another would test that, and would require simultaneous multi-area recordings in one animal, which the design does not afford. The distinction matters for interpretation: what is established is that the information is present in many places, not that the cortex uses one code for it. The stronger claim is frequently assumed in this literature and is not supported by these data.

Anatomical labels in Paper IV are likewise targeted, not histologically verified, and their corroboration is functional, resting on response latencies and F1 distributions.

How far a better time–frequency representation gets you

Paper II was written with a specific destination in mind. If band-pass filtering distorts a spike waveform according to what else happens to be present, and a decomposition does not, then replacing the filter at the front of a spike-sorting pipeline should recover units that current sorters miss, and the cortical recordings of this thesis were the obvious place to try it. Paper II establishes the first half of that argument: on hybrid data with known ground truth, six components of the decomposition carry as much variance as seventeen components of the filtered voltage, and a generic classifier placed on them is substantially more accurate.

The second half did not follow, and the reason belongs here as a result and not as an omission. A spike sorter is not a representation with a classifier attached. It requires automated unsupervised clustering, drift correction over hours of recording, template matching that survives waveform change, and an interface for human error correction, and each of these is a research problem in its own right (Rey et al., 2015; Chung et al., 2017; Rossant et al., 2016). Building one was a larger undertaking than this thesis could absorb alongside the experimental work, and what Paper II contributes to that problem is a component: a front end that is demonstrably better than filtering, waiting on the rest of a pipeline. Papers I and IV were sorted with Kilosort (Pachitariu et al., 2024) accordingly, which has the incidental benefit of keeping unit yields comparable with other work.

The obvious test is to re-sort the Paper IV recordings with a hyperlet front end and ask whether the improvement demonstrated on hybrid data transfers to real recordings, and in particular whether it changes the number of units recovered. That question has consequences beyond convenience, since discrepancies between the number of neurons theoretically available to a probe and the number actually recovered bear on any claim about population dimensionality. Cell-type assignment from spike waveform (Trainito et al., 2019) would also benefit from a representation that has not been distorted by filtering.

That the transform has been taken up outside this work is some evidence that it generalises. The singular superlet formulation on which it builds has since been applied to seizure detection, cardiac arrhythmia classification, phonon dynamics and radio-frequency fingerprinting (Tripathi et al., 2023, 2022; Gentry et al., 2022; Baldini and Dimc, 2025).

The broader perspective

Implications for haptic engineering

If the skin actively and nonlinearly transforms what it is given, generating harmonics, redistributing energy in space, and filtering each frequency in a manner that depends on the rest of the stimulus, then the fidelity of a haptic display depends on the integrative mechanics

of the skin as much as on the density and precision of its actuators. Pre-compensating the actuator, as was done here, corrects the smaller transformation and leaves the larger one untouched. A display that modelled the skin's response, or exploited it, might achieve a more faithful percept than one treating the tissue as a linear medium. The spatial organisation reported in Paper III suggests this is not merely a matter of equalising a transfer function, since the transformation differs across the contact and changes with the stimulus.

Implications for how cortical representations are studied

If fine haptic information is broadly distributed and outlives its stimulus, then asking where touch is represented is less productive than asking how a recursive network holds and transforms it. That is the question the cross-modal invariant framework poses (Ramasamy and Jörntell, 2026), and the demonstration that such representations can emerge from autonomous learning driven by the dependencies of a compliant body begins to answer it (Rongala and Jörntell, 2026; Rongala et al., 2018). It also connects to accounts in which the cortex is treated as a prediction-generating system, not a feature detector (Clark, 2013; Bar, 2007; Loeb and Fishel, 2014), and to observations that behavioural decisions draw on activity distributed brain-wide and not localised to a task-specific area (Mimica et al., 2023; Khilkevich et al., 2024). The methodological corollary is that experiments designed to find localised, feature-tuned responses will tend to find them, because the analyses that make data tractable are the same ones that remove the evidence against localisation.

Future directions

Four experiments follow directly from what is reported here.

The first is to measure the skin mechanics of the rat forepaw under the same display used in Paper IV, which would replace the cross-species inference at the centre of this thesis with a direct measurement. The second is to record afferent activity while imaging the skin, so that the harmonic content available at a location can be compared with what the afferents at that location actually encode; this is the step that would convert Paper III's mechanical claim into a claim about the sensory signal. The third is simultaneous multi-area cortical recording, which is what a test of representational generalisation across cortical sites requires, and which would settle whether the distributed decodability reported in Paper IV reflects a shared code. The fourth is to repeat the persistence measurements in awake animals, where the anaesthetic caveat does not apply and where the prediction, given that ketamine suppresses cortical communication, is that the effect should be larger, not smaller.

Two methodological developments would support all four. Building the hyperlet into a complete spike-sorting pipeline would test whether better time–frequency representation recovers units that current sorters miss. And extending the per-pixel imaging approach to active touch, where the finger moves instead of resting passively, would bring the measure-

ment closer to the condition under which the haptic system actually evolved (Ramasamy and Jörntell, 2026).

Concluding remarks

This thesis set out to find where haptic information is reduced on its way from a mechanical event at the fingertip to a cortical population response. It examined the two junctions at which the reduction is conventionally assumed to occur, and found that neither performs it.

Paper I established the baseline with the skin bypassed: recorded through direct electro-tactile activation, the spontaneous state space of a primary somatosensory population is extremely high-dimensional, its tactile representations are modulated by vision, and input-specific information outlives the input. Paper II derived the hyperlet, a wavelet that integrates a continuum of scales analytically and so achieves super-resolution in a single transform, at a small fraction of the computational cost of the ensemble methods it replaces; its role here was to make it affordable to re-express a signal instead of filtering it. Paper III used that decomposition to map, pixel by pixel, what the fingertip skin does to a controlled vibratory input, and found that the skin generates harmonics that were not present in the stimulus, arranges them into domains that follow the papillary ridges, and rearranges those domains when the stimulus changes, so that the attenuation applied to any one frequency depends on which others accompany it. Paper IV delivered comparable dynamic inputs to the rat forepaw with the skin reinstated, and found fine haptic detail decodable at cortical sites far from the direct thalamocortical input, no neuron tuned to frequency over the range tested, and pattern-specific information persisting for hundreds of milliseconds after the stimulus ended.

Read together these papers show that neither end of the pathway re-expresses less than it receives. The skin returns frequencies its input never contained and organises them by its own microarchitecture, which is why frequency does not survive the interface as an independent variable and why the cortex has no stable frequency to be tuned to. The cortex in turn spreads the resulting information widely and holds it past the stimulus instead of compressing it into local, labelled representations. Neither of the two stages measured here does the work of reduction that the conventional account assigns to it. Whatever selection the system performs appears instead to happen between them, in the brainstem relay where subsets of haptic invariants have been found segregated, instead of distributed evenly along the pathway as the funnel picture implies.

A methodological claim runs alongside the biological one, and the two are not separable. Each result reported here required declining a conventional simplification (filtering, spatial averaging, response averaging, dimensionality reduction), and each would have been

invisible had that simplification been accepted. The orderliness of the classical account of touch is, to a degree that this thesis argues is substantial, a property of how touch has been measured.

What I hope the work makes possible is a change of emphasis in two places. For the engineering of haptic displays, it suggests that the medium deserves as much design attention as the actuator. A display that compensates its own transducer and then treats the skin as though it simply passed the signal on is correcting the smaller of two transformations, and the larger one is spatially organised and changes with what is being played. For the study of cortical representation, it suggests that the productive question is not where touch is represented but how a recursive network holds and reshapes it, and that answering it will require being willing to look at data before simplifying them. Neither point is settled by four papers. Both, I think, are now harder to set aside.

Populärvetenskaplig sammanfattning

Den här avhandlingen består av fyra vetenskapliga arbeten. Målet är att förstå vad som händer med känselinformation på vägen från fingertoppen till hjärnbarken. Den vanliga bilden är en tratt: huden känner av ett tryck, nervsignaler sorteras i separata kanaler på vägen in, och hjärnbarken plockar till slut ut de egenskaper vi upplever. I varje steg antas något sorteras bort. Frågan här är var det bortsorterandet faktiskt sker.

En del av svaret har varit svårt att se därför att mätmetoderna själva förenklar. Mäter man hudens rörelse som ett medelvärde över hela kontaktytan kan man inte se om olika ställen betar sig olika. Filtrerar man en hjärnsignal för att hitta nervcellernas urladdningar kastar man samtidigt bort en del av informationen. Tar man medelvärdet av många upprepningar försvinner just den variation som kan vara det intressanta.

I arbete I undersökte vi hur sammansatt aktiviteten är hos en grupp nervceller i den del av hjärnbarken som anses hantera känsel. Vi gav elektriska beröringsmönster direkt till nerverna i tassens, och ibland samtidigt ljusblixtar, och spelade in från flera tiotal nervceller på en gång. Tre saker framkom. Aktiviteten var betydligt mer mångfacetterad än väntat. Nervcellerna reagerade inte bara på beröring utan kunde också skilja olika ljusmönster från varandra, trots att området anses vara enbart ett känselområde. Och informationen om vad som just hänt fanns kvar i aktiviteten långt efter att stimuleringen upphört.

Arbete II handlar om verktyget. För att kunna se detaljer i en signal behöver man kunna beskriva den både i tid och i frekvens samtidigt, och det finns en fysikalisk gräns för hur bra det går. Vi härledde en ny matematisk metod, kallad hyperlet, som beskriver signalen på ett rikare sätt i stället för att filtrera bort delar av den, och som är ungefär tjugo gånger snabbare än den bästa jämförbara metoden. Det gjorde det praktiskt möjligt att analysera mycket stora datamängder punkt för punkt.

I arbete III använde vi den metoden på huden. Vi filmade en fingertopp mot en vibrerande glasplatta med tusen bilder i sekunden och analyserade rörelsen i varje enskild bildpunkt. Resultatet var oväntat. Huden skickar inte bara vidare den ton den får, utan skapar övertoner, alltså nya frekvenser som inte fanns i signalen. Den gör det dessutom olika på olika ställen, i band som följer fingeravtryckets åsar, och vilka övertoner ett visst ställe ger ifrån sig ändras när man byter frekvens på det som spelas upp. Huden förenklar alltså inte signalen. Den gör den rikare, och på ett sätt som beror på vad som spelas.

I arbete IV följde vi signalen vidare in i hjärnan. Vi gav vibrationsmönster till tassens på sövd råtta och spelade in från hundratals nervceller, de flesta långt utanför det område som bru-

kar kallas känselbarken. Även där kunde nervcellerna skilja mycket lika vibrationsmönster från varandra, vilket betyder att informationen är spridd och inte samlad på ett ställe. Ingen enda nervcell var däremot inställd på en bestämd frekvens. Det är svårt att förstå i sig, men begripligt tillsammans med resultatet från arbete III: om huden ändrar frekvensinnehållet olika på olika ställen finns det ingen stabil frekvens kvar att vara inställd på. Precis som i arbete I fanns informationen om vad som just spelats kvar i flera tiondels sekunder efter att stimuleringen upphört.

Sammantaget pekar arbetena åt samma håll. Känselinformation trattas inte ihop på vägen in, den skrivs om. Huden lägger till struktur, och hjärnbarken sprider ut den och håller kvar den i stället för att komprimera den. Att detta inte setts tidigare beror till en del på att de vanliga mätmetoderna förenklar innan man tittar. Kunskapen har två praktiska riktningar. För den som bygger utrustning som ska återge känsel, till exempel i handkontroller eller skärmar, betyder det att hudens egen mekanik är minst lika viktig som hur exakt själva vibratorn arbetar. För hjärnforskningen betyder det att frågan om var känseln sitter är mindre användbar än frågan om hur ett nätverk håller kvar och omvandlar information.

Acknowledgements

Five years in one lab leaves a long list of people who made the work possible, and a shorter list of people without whom it would not have started or finished. Both are below.

Henrik Jörntell

Thank you for giving me a problem rather than a protocol. You let me spend a long time deriving the hyperlet before it was at all clear where it would lead, and then you saw further with it than I did: what began as a way of looking at a signal became something much larger, largely because you insisted it could be. That trust, both in me and in the value of taking on big problems rather than safe ones, is why most of this thesis exists. You have been as generous with support outside the work as within it, and I have felt that throughout these years. I will miss the conversations that began as a quick question and took the rest of the afternoon. Thank you for your time, for your curiosity, and for taking me on in the first place.

Fredrik Bengtsson

Thank you for taking on the role of assistant supervisor, and for the part you have played in keeping this department a place where this kind of work gets done.

Udaya Rongala

For the discussions, which were always worth having, and for the good times along the way. Thank you.

Johanna Norrlid

You taught me how to do the surgeries. Everything in this thesis that required an animal on the table rests on what you showed me, patiently and more than once. Thank you.

Ann-Sophie Alm

Thank you for the laboratory work, and for being the person who knew where everything was kept and who to ask about anything the rest of us could not work out.

Sofie Skårup Kristensen

We shared an office, a project and most of a PhD. Paper I is yours and I am glad to have been part of it. Thank you for the years of arguing about analysis, and for the company through the INTUITIVE project.

Anders Wahlbom and Jonas Enander

You built and taught the experimental approach that the recordings in this thesis rest on, and I benefited from it long after you had both moved on. Thank you, and good luck in your current positions.

Fellow doctoral students

Szilvia Szeier and Astrid Mellbin. Thank you for the shared office, the journal clubs, and the commiseration over everything a PhD entails.

The INTUITIVE network

This work was carried out within the EU Horizon 2020 MSCA-ITN ‘INTUITIVE’ consortium, which supplied most of its collaborations and much of its equipment. Thank you in particular to Vincent Hayward, whose framing of the problem the third paper of this thesis is essentially an attempt to measure, and to Actronika for the haptic display that Papers III and IV depend on.

References

- Addison, P. (2018). Introduction to redundancy rules: the continuous wavelet transform comes of age. *Philos. Trans. Royal Soc. A: Math. Phys. Eng. Sci*, 376:20170258. [from Paper II].
- Adelson, E. H. and Bergen, J. R. (1991). The plenoptic function and the elements of early vision. In *Computational Models of Visual Processing*, pages 3–20. MIT Press.
- Baldini, G. and Dimc, F. (2025). Transient-based radio frequency fingerprinting using singular superlet transform and convolutional neural network. *Internet Technol. Lett*, 8. [from Paper II].
- Bar, M. (2007). The proactive brain: using analogies and associations to generate predictions. *Trends in Cognitive Sciences*, 11:280–289.
- Belitski, A., Gretton, A., Magri, C., Murayama, Y., Montemurro, M., Logothetis, N., and Panzeri, S. (2008). Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *J. Neurosci*, 28:5696–5709. [from Paper II].
- Bengtsson, F., Brasselet, R., Johansson, R. S., Arleo, A., and Jörntell, H. (2013). Integration of sensory quanta in cuneate nucleus neurons in vivo. *PLoS ONE*, 8:e56630.
- Bensmaïa, S. J. and Hollins, M. (2003). The vibrations of texture. *Somatosensory & Motor Research*, 20(1):33–43.
- Bhattacharjee, M., Prasanna, S., and Guha, P. (2020). Speech/music classification using features from spectral peaks. *IEEE/ACM Trans. Audio Speech. Lang. Process*, 28:1549–1559. [from Paper II].
- Bimbard, C., Sit, T. P. H., Lebedeva, A., Reddy, C. B., Harris, K. D., and Carandini, M. (2023). Behavioral origin of sound-evoked activity in mouse visual cortex. *Nature Neuroscience*, 26:251–258.
- Buzsáki, G. (2004). Large-scale recording of neuronal ensembles. *Nat. Neurosci*, 7:446–451. [from Paper II].
- Cavanaugh, J. T., Kelty-Stephen, D. G., and Stergiou, N. (2017). Multifractality, interactivity, and the adaptive capacity of the human movement system: a perspective for

- advancing the conceptual basis of neurologic physical therapy. *J. Neurol. Phys. Ther.*, 41:245–251. [from Paper I].
- Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A., and Fiete, I. (2019). The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nat. Neurosci.*, 22:1512–1520. [from Paper I].
- Chikkerur, S., Cartwright, A., and Govindaraju, V. (2007). Fingerprint enhancement using stft analysis. *Pattern Recognit.*, 40:198–211. [from Paper II].
- Choi, H. and Williams, W. (1989). Improved time-frequency representation of multicomponent signals using exponential kernels. *IEEE Trans. Acoust.*, 37:862–871. [from Paper II].
- Chung, J., Joo, H., Fan, J., Liu, D., Barnett, A., Chen, S., Geaghan-Breiner, C., Karlsson, M., Karlsson, M., Lee, K., and Liang, H. (2019). High-density, long-lasting, and multiregion electrophysiological recordings using polymer electrode arrays. *Neuron*, 101:21–31. [from Paper II].
- Chung, J., Magland, J., Barnett, A., Tolosa, V., Tooker, A., Lee, K., Shah, K., Felix, S., Frank, L., and Greengard, L. (2017). A fully automated approach to spike sorting. *Neuron*, 95:1381–1394. [from Paper II].
- Clark, A. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36:181–204.
- Cohen, L. (1966). Generalized phase-space distribution functions. *J. Math. Phys.*, 7:781–786. [from Paper II].
- Craig, A. D. (2003). Pain mechanisms: labeled lines versus convergence in central processing. *Annual Review of Neuroscience*, 26:1–30.
- de Lafuente, V. and Romo, R. (2006). Neural correlate of subjective sensory experience gradually builds up across cortical areas. *Proceedings of the National Academy of Sciences*, 103:14266–14271.
- Dum, R. P., Levinthal, D. J., and Strick, P. L. (2009). The spinothalamic system targets motor and sensory areas in the cerebral cortex of monkeys. *Journal of Neuroscience*, 29:14223–14235.
- Emanuel, A. J., Lehnert, B. P., Panzeri, S., Harvey, C. D., and Ginty, D. D. (2021). Cortical responses to touch reflect subcortical integration of LTMR signals. *Nature*, 600:680–685.
- Enander, J. M. D. and Jörntell, H. (2019). Somatosensory cortical neurons decode tactile input patterns and location from both dominant and non-dominant digits. *Cell Reports*, 26:3551–3560.

- Enander, J. M. D., Spanne, A., Mazzoni, A., Bengtsson, F., Oddo, C. M., and Jörntell, H. (2019). Ubiquitous neocortical decoding of tactile input patterns. *Frontiers in Cellular Neuroscience*, 13:140.
- Etemadi, L., Enander, J. M. D., and Jörntell, H. (2022). Remote cortical perturbation dynamically changes the network solutions to given tactile inputs in neocortical neurons. *iScience*, 25:103557.
- Etemadi, L., Enander, J. M. D., and Jörntell, H. (2023). Hippocampal output profoundly impacts the interpretation of tactile input patterns in SI cortical neurons. *iScience*, 26:106885.
- Faisal, A. A., Selen, L. P. J., and Wolpert, D. M. (2008). Noise in the nervous system. *Nat. Rev. Neurosci.*, 9:292–303. [from Paper I].
- Funahashi, S., Bruce, C. J., and Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey’s dorsolateral prefrontal cortex. *J. Neurophysiol.*, 61:331–349. [from Paper I].
- Gentry, C., Liao, C., You, W., Ryan, S., Varner, B., Shi, X., Guan, M., Gray, T., Temple, D., Meng, S., and Raschke, M. (2022). Super-resolved time-frequency measurements of coupled phonon dynamics in a 2d quantum material. *Sci. Rep.*, 12:19734. [from Paper II].
- Ghazanfar, A. A. and Schroeder, C. E. (2006). Is neocortex essentially multisensory? *Trends in Cognitive Sciences*, 10:278–285.
- Goldman-Rakic, P. S. (1996). Colloquium paper this paper was presented at a colloquium entitled memory: recording experience in cells and circuits. *Proc. Natl Acad. Sci.*, 93:13473–13480. [from Paper I].
- Haque, Z. Z., Samandra, R., and Mansouri, F. A. (2021). Neural substrate and underlying mechanisms of working memory: insights from brain stimulation studies. *J. Neurophysiol.*, 126:2038–2053. [from Paper I].
- Hashemi, H. (2012). Fuzzy clustering of seismic sequences: segmentation of time-frequency representations. *IEEE Signal Process. Mag.*, 29:82–87. [from Paper II].
- Hayward, V. (2011). Is there a ‘plenhaptic’ function? *Philosophical Transactions of the Royal Society B*, 366:3115–3122.
- Hayward, V., Terekhov, A. V., Wong, S.-C., Geborek, P., Bengtsson, F., and Jörntell, H. (2014). Spatio-temporal skin strain distributions evoke low variability spike responses in cuneate neurons. *Journal of the Royal Society Interface*, 11:20131015.

- He, L., Fu, J., Li, Y., Xiong, X., and Zhang, J. (2022). Wnsa-net: an axial-attention-based network for schizophrenia detection using wideband and narrowband spectrograms. *IEEE/ACM Trans. Audio Speech Lang. Process*, 31:721–733. [from Paper II].
- Jazayeri, M. and Ostojic, S. (2021). Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity. *Curr. Op. Neurobiol*, 70:113–120. [from Paper I].
- Johansson, R. S. and Flanagan, J. R. (2009). Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience*, 10(5):345–359.
- Johansson, R. S., Landström, U., and Lundström, R. (1982). Responses of mechanoreceptive afferent units in the glabrous skin of the human hand to sinusoidal skin displacements. *Brain Research*, 244:17–25.
- Jörntell, H., Bengtsson, F., Geborek, P., Spanne, A., Terekhov, A. V., and Hayward, V. (2014). Segregation of tactile input features in neurons of the cuneate nucleus. *Neuron*, 83:1444–1452.
- Jörntell, H., Devillard, A., and Burdet, E. (2026). The identification of haptic invariants in humans and their applications to robotics. *Multisensory Research*, 39(3–5).
- Jun, J. J., Steinmetz, N. A., Siegle, J. H., et al. (2017). Fully integrated silicon probes for high-density recording of neural activity. *Nature*, 551:232–236.
- Khilkevich, A., Lohse, M., Low, R., Orsolic, I., Bozic, T., Windmill, P., and Mrsic-Flogel, T. D. (2024). Brain-wide dynamics linking sensation to action during decision-making. *Nature*, 634:890–900.
- Khona, M. and Fiete, I. R. (2022). Attractor and integrator networks in the brain. *Nat. Rev. Neurosci*, 23:744–766. [from Paper I].
- Kristensen, S. S. and Jörntell, H. (2023). Differential encoding of temporally evolving color patterns across nearby V1 neurons. *Frontiers in Cellular Neuroscience*, 17:1249522.
- Kristensen, S. S., Kesgin, K., and Jörntell, H. (2024). High-dimensional cortical signals reveal rich bimodal and working memory-like representations among S1 neuron populations. *Communications Biology*, 7(1):1043.
- Loeb, G. E. and Fishel, J. A. (2014). Bayesian action & perception: representing the world in the brain. *Front. Neurosci*, 8:112655. [from Paper I].
- Loughlin, P., Pitton, J., and Hannaford, B. (1994). Approximating time-frequency density functions via optimal combinations of spectrograms. *IEEE Signal Process. Lett*, 1:199–202. [from Paper II].

- Luczak, A. and Barthó, P. (2012). Consistent sequential activity across diverse forms of up states under ketamine anesthesia. *Eur. J. Neurosci*, 36:2830–2838. [from Paper I].
- Luo, Y. et al. (2025). Tfdnet: time-frequency enhanced decomposed network for long-term time series forecasting. *Pattern Recognit*, 162:111412. [from Paper II].
- Manfredi, L. R., Baker, A. T., Themba, D. O., Yoshioka, T., and Bensmaia, S. J. (2012). The effect of surface wave propagation on neural responses to vibration in primate glabrous skin. *PLoS ONE*, 7(2):e31203.
- Marks, T. D. and Goard, M. J. (2021). Stimulus-dependent representational drift in primary visual cortex. *Nat. Commun*, 12. [from Paper I].
- Matteucci, G. et al. (2022). Cortical sensory processing across motivational states during goal-directed behavior. *Neuron*, 110:4176–4193. [from Paper I].
- Mimica, B., Tombaz, T., Battistin, C., Fuglstad, J. G., Dunn, B. A., and Whitlock, J. R. (2023). Behavioral decomposition reveals rich encoding structure employed across neocortex in rats. *Nature Communications*, 14:3947.
- Moca, V. V., Bârzan, H., Nagy-Dăbâcan, A., and Mureşan, R. C. (2021). Time-frequency super-resolution with superlets. *Nature Communications*, 12:337. [from Paper II].
- Mogensen, H., Norrlid, J., Enander, J. M. D., Wahlbom, A., and Jörntell, H. (2019). Absence of repetitive correlation patterns between pairs of adjacent neocortical neurons in vivo. *Frontiers in Neural Circuits*, 13:48.
- Morlet, J., Arens, G., Fourgeau, E., and Glard, D. (1982). Wave propagation and sampling theory—part i: complex signal and scattering in multilayered media. *Geophysics*, 47:203–221. [from Paper II].
- Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S., and Churchland, A. K. (2019). Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience*, 22:1677–1686.
- Nagamori, A., Laine, C. M., Loeb, G. E., and Valero-Cuevas, F. J. (2021). Force variability is mostly not motor noise: theoretical implications for motor communications biology | (2024) 7:1043 control. *PLoS Comput. Biol*, 17. [from Paper I].
- Niedermeyer, E. and Lopes da Silva, F. H. (2005). *Electroencephalography: basic principles, clinical applications, and related fields*. Lippincott Williams & Wilkins.
- Nilsson, M. N. P. and Jörntell, H. (2021). Channel current fluctuations conclusively explain neuronal encoding of internal potential into spike trains. *Phys. Rev. E*, 103. [from Paper I].

- Norrlid, J., Enander, J. M. D., Mogensen, H., and Jörntell, H. (2021). Multi-structure cortical states deduced from intracellular representations of fixed tactile input patterns. *Frontiers in Cellular Neuroscience*, 15:677568.
- Oddo, C. M., Mazzoni, A., Spanne, A., Enander, J. M. D., Mogensen, H., Bengtsson, F., Camboni, D., Micera, S., and Jörntell, H. (2017). Artificial spatiotemporal touch inputs reveal complementary decoding in neocortical neurons. *Scientific Reports*, 8:45898.
- Olcese, U., Iurilli, G., and Medini, P. (2013). Cellular and synaptic architecture of multi-sensory integration in the mouse neocortex. *Neuron*, 79:579–593.
- Pachitariu, M., Sridhar, S., Pennington, J., and Stringer, C. (2024). Spike sorting with Kilosort4. *Nature Methods*, 21:914–921.
- Prsa, M., Morandell, K., Cuenu, G., and Huber, D. (2019). Feature-selective encoding of substrate vibrations in the forelimb somatosensory cortex. *Nature*, 567(7748):384–388.
- Ramasamy, A. and Jörntell, H. (2026). Cross-modal invariants as the foundational basis of emerging brain circuitry representations. *Physics of Life Reviews*, 58:223–240.
- Rey, H., Pedreira, C., and Quiroga, R. (2015). Past, present and future of spike sorting techniques. *Brain Res. Bull.*, 119:106–117. [from Paper II].
- Roland, P. E. (2023). How far neuroscience is from understanding brains. *Frontiers in Systems Neuroscience*, 17:1147896.
- Rongala, U. B. and Jörntell, H. (2026). Autonomous learning from intrinsic sensory dependencies yields generalizable representations in cortical-like networks. *PRX Life*, 4(2):023030.
- Rongala, U. B., Seyfarth, A., Hayward, V., and Jörntell, H. (2024). The import of skin tissue dynamics in tactile sensing. *Cell Reports Physical Science*, 5(5):101943.
- Rongala, U. B., Spanne, A., Mazzoni, A., Bengtsson, F., Oddo, C. M., and Jörntell, H. (2018). Intracellular dynamics in cuneate nucleus neurons support self-stabilizing learning of generalizable tactile representations. *Frontiers in Cellular Neuroscience*, 12:210.
- Rossant, C., Kadir, S., Goodman, D., Schulman, J., Hunter, M., Saleem, A., Grosmark, A., Belluscio, M., Denfield, G., Ecker, A., and Tolias, A. (2016). Spike sorting for large, dense electrode arrays. *Nat. Neurosci.*, 19:634–641. [from Paper II].
- Shao, Y., Hayward, V., and Visell, Y. (2016). Spatial patterns of cutaneous vibration during whole-hand haptic interactions. *Proceedings of the National Academy of Sciences*, 113(15):4188–4193.
- Spanne, A. and Jörntell, H. (2015). Questioning the role of sparse coding in the brain. *Trends in Neurosciences*, 38(7):417–427.

- Steinmetz, N., Aydin, C., Lebedeva, A., Okun, M., Pachitariu, M., Bauza, M., Beau, M., Bhagat, J., Böhm, C., Broux, M., and Chen, S. (2021). Neuropixels 2.0: a miniaturized high-density probe for stable, long-term brain recordings. *Science*, 372. [from Paper II].
- Stergiou, N. and Decker, L. M. (2011). Human movement variability, nonlinear dynamics, and pathology: is there a connection? *hum. Mov. Sci.*, 30:869–888. [from Paper I].
- Stringer, C., Michaelos, M., Tsybouski, D., Lindo, S. E., and Pachitariu, M. (2021). High-precision coding in visual cortex. *Cell*, 184:2767–2778.
- Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., and Harris, K. D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364:255.
- Szeier, S. and Jörntell, H. (2026). Brain circuitry behavioral control emerging from complexity of nested recurrent loops extending into the body. *bioRxiv*. Preprint.
- Szeier, S., Svensson, E., Wangel, R., and Jörntell, H. (2025). Individual cortical neurons innervating large proportions of the neocortical areas in the MouseLight database. *bioRxiv*.
- Tang, D. L., Parrell, B., and Niziolek, C. A. (2022). Movement variability can be modulated in speech production. *J. Neurophysiol.*, 128:1469–1482. [from Paper I].
- Trainito, C., von Nicolai, C., Miller, E. K., and Siegel, M. (2019). Extracellular spike waveform dissociates four functionally distinct cell classes in primate cortex. *Current Biology*, 29. [from Paper II].
- Tripathi, P., Kumar, A., Kumar, M., and Komaragiri, R. (2022). Multilevel classification and detection of cardiac arrhythmias with high-resolution superlet transform and deep convolution neural network. *IEEE Trans. Instrum. Meas.*, 71:1–13. [from Paper II].
- Tripathi, P., Kumar, A., Kumar, M., and Komaragiri, R. (2023). Automatic seizure detection and classification using super-resolution superlet transform and deep neural network – a preprocessing-less method. *Comput. Methods Programs Biomed.*, 240:107680. [from Paper II].
- Turecek, J. and Ginty, D. D. (2024). Coding of self and environment by Pacinian neurons in freely moving animals. *Neuron*, 112:3267–3277.
- Wahlbom, A., Enander, J. M. D., Bengtsson, F., and Jörntell, H. (2019). Focal neocortical lesions impair distant neuronal information processing. *Journal of Physiology*, 597:4357–4371.
- Wahlbom, A., Enander, J. M. D., and Jörntell, H. (2021). Widespread decoding of tactile input patterns among thalamic neurons. *Frontiers in Systems Neuroscience*, 15:640085.

- Wallace, M. T., Ramachandran, R., Stein, B. E., and Raichle, M. E. (2004). A revised view of sensory cortical parcellation. *PNAS*, 101:2167–2172. [from Paper I].
- Weber, A. I., Saal, H. P., Lieber, J. D., Cheng, J.-W., Manfredi, L. R., Dammann, J. F., and Bensmaia, S. J. (2013). Spatial and temporal codes mediate the tactile perception of natural textures. *Proceedings of the National Academy of Sciences*, 110(42):17107–17112.
- Wessberg, J., Stambaugh, C. R., Kralik, J. D., Beck, P. D., Laubach, M., Chapin, J. K., Kim, J., Biggs, S. J., Srinivasan, M. A., and Nicolelis, M. A. L. (2000). Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, 408:361–365.
- Wiertelowski, M. and Hayward, V. (2012). Mechanical behavior of the fingertip in the range of frequencies and displacements relevant to touch. *Journal of Biomechanics*, 45(11):1869–1874.
- Wiertelowski, M., Hudin, C., and Hayward, V. (2011a). On the $1/f$ noise and non-integer harmonic decay of the interaction of a finger sliding on flat and sinusoidal surfaces. In *2011 IEEE World Haptics Conference*, pages 25–30.
- Wiertelowski, M., Lozada, J., and Hayward, V. (2011b). The spatial spectrum of tangential skin displacement can encode tactual texture. *IEEE Transactions on Robotics*, 27(3):461–472.
- Wilson, R., Calway, A., and Pearson, E. (1992). A generalized wavelet transform for fourier analysis: the multiresolution fourier transform and its application to image and audio signal analysis. *IEEE Trans. Inf. Theory*, 38:674–690. [from Paper II].
- Wright, N., Borden, P., Liew, Y., Bolus, M., Stoy, W., Forest, C., and Stanley, G. (2021). Rapid cortical adaptation and the role of thalamic synchrony during wakefulness. *J. Neurosci*, 41:5421–5439. [from Paper II].
- Wu, H. G., Miyamoto, Y. R., Castro, L. N. G., Ölveczky, B. P., and Smith, M. A. (2014). Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat. Neurosci*, 17:312–321. [from Paper I].
- Wu, M., Wan, T., Wan, X., Fang, Z., and Du, Y. (2021). A new localization method for epileptic seizure onset zones based on time-frequency and clustering analysis. *Pattern Recognit*, 111:107687. [from Paper II].
- Xia, J., Marks, T. D., Goard, M. J., and Wessel, R. (2021). Stable representation of a naturalistic movie emerges from episodic activity with gain variability. *Nat. Commun*, 12. [from Paper I].
- Zhao, T., Fang, L., Ma, X., Li, X., and Zhang, C. (2025). Tfformer: A time-frequency domain bidirectional sequence-level attention based transformer for interpretable long-term sequence forecasting. *Pattern Recognit*, 158:110994. [from Paper II].

Zhou, A. Y., Fuster, J. M., Zhou, Y., and Fuster, J. M. (2020). Visuo-tactile crossmodal associations in cortical somatosensory cells. *PNAS*, 97:9777–9782. [from Paper I].