UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Coding cognition

Representing and communicating information in perceptual and memory systems

Dorman, R.

Publication date
2025

Link to publication

Citation for published version (APA):
Dorman, R. (2025). *Coding cognition: Representing and communicating information in perceptual and memory systems*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:09 Nov 2026

Chapter 5

General Discussion

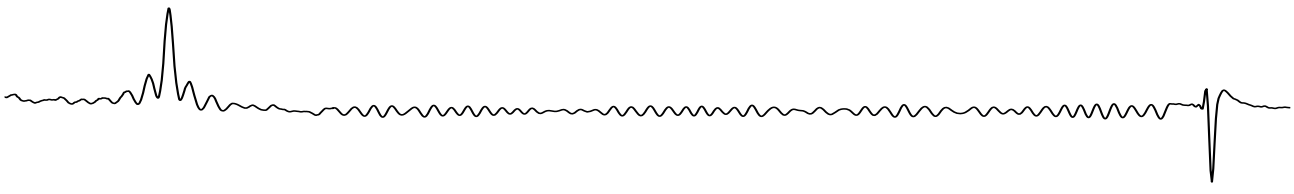

5.1 Summary of results

In this thesis we present three studies investigating how neuronal activity relates to perception and cognition. Specifically, we have investigated correlations between spiking activity within and between populations of different sensory and mnemonic brain areas (Chapter 2); spiking of visual cortex neurons in relation to the oscillations of the local LFPs (Chapter 3); and the correlation of spiking in V1 to motor behaviour during a task with a forced response delay.

Spike-based coupling between single neurons and populations across rat sensory cortices, perirhinal cortex and hippocampus. In chapter 2, we correlated single-neuron spiking with summed population activity within the same area and between different areas. While earlier studies investigated such Single-Neuron to Population Coupling (SNPC) within the same area (Okun et al., 2015), this analysis has not been performed before between various cortical areas.

While investigating the SNPC as an aspect of neural communication, we concluded the following: 1) within-area coupling was found in all areas, with a noticeably strong within-area correlation in the CA1 region of the hippocampus; 2) between-area coupling shows a broader temporal profile compared to within-area coupling, presumably due to the predominantly monosynaptic connectivity within areas, and variable (and likely polysynaptic) pathways between the different areas. The variability in these profiles might hide complex patterns of neuronal connectivity not easily discernible here; 3) generally, a positive correlation was found between within-area and between-area coupling, suggesting that neurons strongly connected within their local population are also important for relaying the local information to other areas; 4) Despite point 3, we have identified cells strongly 'broadcasting' to different areas while being weakly correlated to local populations. When relating this quantification of single-neuron coupling to spatial selectivity of neurons in the four recorded areas, we report a negative correlation between spatial selectivity and coupling strength, notably in hippocampal CA1 neurons, corresponding with the specific information spatially selective cells would require to generate their highly specific place-field responses (Barnes et al., 1990; Sosa, Gillespie, & Frank, 2018) Lastly, we find within-area coupling to vary between active task performance, quiet wakefulness (i.e. during the inter-trial interval) and non-REM sleep. Between areas, we only saw noticeable effects of these brain states on the correlations between single cells of CA1 with the other three populations.

Feature and modality specific phase-coding in V1 during multisensory detection. In chapter 3, we turn our attention not only to the way spiking activity is employed to communicate across the brain, but also to the question how spiking activity and LFPs might form a coherent, relevant code suitable for representing the identity of different modalities and features within the visual cortex. To this end, we correlated spiking activity with 10 Hz (theta) oscillations in V1 and investigated spiking coherent with particular phases of the LFP. We found that phase-locking occurs during unimodal

conditions (i.e. upon auditory and visual stimuli) and increases with multisensory stimulation. Within the subset of neurons that show phase-locking across multiple stimulus conditions, we see that neurons can shift their preferred phase-of-firing to the theta oscillations in two conditions: with respect to the orientation of the visual stimulus, and between the visual-only and multisensory conditions.

Low impact of orofacial motion on selective primary visual cortex responses to visual stimulation with delayed introduction of reward apparatus. In chapter 4, we aimed to investigate the "contamination" of stimulus-related activity in V1 with non-sensory signals (Bimbard et al., 2023; Oude Lohuis et al., 2024, but see later). To this end, we made use of a forced delay in the same task as used in Chapter 3: While the stimulus was presented, the animal was unable to physically reach the reward apparatus until a delay period had passed. During stimulus presentation, but before the licking response was initiated, orofacial motion was generally low. When animals were able to respond, V1 activity increased, correlated with the observed orofacial motion of the animal. Secondly, auditory stimulation led to noticeable orofacial movement, even in the delay period, which correlated with V1 spiking activity. V1 cells showed only a minor orofacial component in their responses to visual stimuli and thus generated a clearer representation of visual information during the delay period, before behavioural responses started.

Comparing single-neuron to population coupling with phase-coding
The studies presented here have shown several types of dynamics between neurons and population activity: in chapter 2 we studied spike-based dynamics by way of the SNPC, and in chapter 3 we used spikes and the LFP. Both metrics probe the relation between neuron activity and broader population and network activity; the LFP, as used in conjunction with spiking activity in phase-coding, arises primarily from mass synaptic input to the recorded area ((Buzsáki, Anastassiou, & Koch, 2012), while the SNPC is a direct spike-based correlation between population activity and single neurons.

Population coupling, or the SNPC, is a measure directly based on spiking data. In its raw form, cross-correlation of neuron and population spiking activity, this provides a straightforward measure of the temporal alignment of neurons and populations. In chapter 2, the metric was calculated on large, singular sections of data (ranging up to several hours). It provides a generalized metric of neuronal coupling as it is not explicitly calculated on stimuli epochs or locked to stimuli. From our data it is not known how much specific stimulus-related activity would influence the generalization of this metric, but Okun et al (2015) showed it to be invariant to stimuli (however, this applies to within-area coupling only). The metric reflects functional connectivity, which is related to anatomical connectivity, but with the added benefit of including dynamics in the spiking signals that pertain to different sections of the network (i.e., brain areas) under scrutiny.

The most relevant unknown factor influencing population coupling would be the convergence of neuronal activity outside the scope of the recording electrodes to the neural activity that we actually record. For example, many sensory cortical streams converge on the perirhinal cortex and indirectly on the hippocampus (Agster & Burwell, 2009; Burwell & Amaral, 1998; Naber, Witter, & Lopes Da Silva, 2000; Zakiewicz, Bjaalie, & Leergaard, 2014). Convergence can also be considered in a top-down fashion: hippocampal neurons correlating in their activity with visual cortex populations may engage in relevant information exchange with visual areas outside of the reach of the electrodes. At present, it is technically unfeasible to investigate all (or even a significant portion of) convergence of spiking signals down the cortico-hippocampal hierarchy, so the SNPC will be limited in explaining functional connectivity down this hierarchy. To specifically investigate convergence one could supplement the SNCP with optogenetics to identify relevant target populations to relate with other regions in the cortico-hippocampal hierarchy.

In comparison, the metric calculated in phase-coding is computed from two types of physiological data - spiking data and LFPs - which both have to be scrutinized for appropriate interpretation. The LFP is largely made up of postsynaptic potentials generated by dendrites oriented in parallel (Buzsáki, 2004; Buzsáki, Anastassiou, & Koch, 2012; Einevoll et al., 2013). To get noticeable voltage deflections in the LFP, such dendritic inputs have to be nearly synchronous. In other words: strong, negative deflections in the LFP would correspond to a measurement close to a voltage sink, and would theoretically correlate with synchronized inward currents in the dendrite, e.g. from excitatory inputs. In the same vein, positive deflections (when measured close to a source) would be caused by either positive return currents or synchronized inhibition. However, the biological and measurement conditions are rarely as simple as this (see also Box 3 in Einevoll et al., 2013). For instance, the causal contributions of the LFP to the generation of spiking activity are not always clear, and not all input in a certain area will be impacting on parallel dendrites conducive to the LFP. Furthermore, spike-spike correlations between areas need not directly correspond with spike-LFP correlations due to filtering properties in dendrites, synaptic inhibition, dynamic spiking thresholds (Buzsáki, Anastassiou, & Koch, 2012; Buzsáki & Schomburg, 2015; Currier & Clandinin, 2025; Platkiewicz & Brette, 2010; Yi et al., 2015). Additionally, probe recordings provide little information of spatially significant properties of LFP generation. Discussions are still ongoing on the hippocampal origin of cortical theta oscillations (Kragel et al., 2020; Nuñez & Buño, 2021; Zimmerman et al., 2025; see also Chapter 3), but volume-conducted transients from other regions than the recorded area cannot be readily discarded without direct measurements from those regions. Lastly, however the LFP signal is generated exactly, it also suffers from the convergence problem as applies to the SNPC. Even if we assume that the LFP is strongly correlated with input to an area, we are unsure about all the sources that drive this input unless we extensively probe all possible sources.

Additionally, although phase-coding is technically interpretable without knowing about the origins of the LFP, knowing the exact relation between the LFP and phase-

coding in neurons would ideally provide us with relevant insights in how the brain came to use, and is using, such a mode of neural coding in processing information. To further investigate the precise relationship between LFPs and spikes, analytical techniques are necessary that might not entirely align with the variable nature of neural tissue. In Chapter 3, although periods of similar oscillatory power have been selected, LFP signals on which phase-of-firing is based are susceptible to trial-by-trial variation, noise and volume conduction, and we cannot necessarily assume the recorded LFP signal is the same as the biophysically relevant signal for any phase-coding neuron. Moreover, to get comparable phase-of-firing distributions from individual neurons to which we can apply statistics with enough power, sufficient collections of spikes have to be pooled. Being fairly conservative, we opted for a Bonferroni-corrected threshold for statistical significance to distinguish phase-locked neurons, minimizing type II errors. While this procedure aims to get a statistically valid interpretation, there is a body of data which falls in a grey area: are we correct in assuming there are neurons that are just not phase-locking? Potentially, there could be an unknown physiological or behavioral factor that would explain a trial-by-trial or spike-by-spike variability in the phase-of-firing. Lastly, the choice of the frequency of interest, 10 Hz, although based on empirical data and literature review, is still constrained by the computational choices we made. It is unlikely that neurobiology would work with exact integers or oscillations fixed from moment to moment. These points highlight the complexity of studying phase-coding as a component of neural coding in general.

Although we would ideally establish neuroscientific models and theories that match the data perfectly, the collected data versus the entirety of the neurobiological system leave a large gap that is at present technically unfeasible to cross except in a rough, statistically verifiable way. The rapid developments in recording techniques and methods are bridging this gap with more neurons to record simultaneously, and more data being gathered. Additionally, much progress can be made in sophisticated behavioral and intervention paradigms. Alternatively, this bridging can be extended from the other side by testing neuroscientific predictions with computational methods and vice versa.

What can we learn about the networks we investigate with both population-coupling and phase-coding? Whereas single-neuron to population coupling is a fairly straightforward metric of functional connectivity between neurons and a population, phase-coding reflects rhythmic coordination of neuronal signals and input synchrony (although, as mentioned above, the relation of LFPs with synaptic input has its complexities). Both metrics can be interpreted, neurophysiologically, in the context of how single neurons are embedded in a particular network, and how their activity coheres with synaptic relationships. As we show in Chapter 3, phase-locking and phase-of-firing are affected by what the network is processing, but at present this has not yet been compared with the coupling levels of its neurons. Changes in the phase-of-firing of multiple neurons were not found to consistently move in one direction, so these would likely be averaged out when calculating levels of population coupling. Possibly, they may be visible as changes in levels of coupling due to changes in temporal

correlations between single-neuron spiking and population activity. Combining both methods in overlapping datasets may reveal how these metrics compare and show the overlap of hidden population structures between structural population coupling and phase-coding in the temporal domain.

Technical considerations

The experiments performed in the field of cognitive neurophysiology come with many parameters (many dependent on the research questions, and random effects) which could influence the results. Examples of these parameters comprise behavioral factors (e.g. task design and task rules, stimuli, timing, method of training), recording factors (electrode layout and placement, insertion speed [see Fiáth et al. (2019)], physical and electrical stability, etc.) and neurobiological factors (animal model, the individual subject's anatomy, tissue-specific volume conduction, effects of stress or other non-ecological conditions, etc.). Here I will consider the two recording techniques we used and their consequences, and the consideration of appropriate sensory stimuli in cognitive neurophysiology.

Recording techniques. For chapter 2, sizeable investments have been made in the development of the Quad-drive. This device increased the amount of tetrodes from the typically used Hyperdrives (Lansink et al., 2007; Wilson & McNaughton, 1993) from 14 to 36, allowing our current study to target four different brain areas with 9 tetrodes each. However, such devices come with technical challenges. Implantation with four craniotomies was necessary, and for chronic recordings in behaving animals such a device needs to be implanted stably enough to prevent additional noise. Moreover, although the placement of the drive and the direction of the individual tetrodes has been modelled beforehand, slight deviations during implantation could lead to less-than-ideal placement of the tetrodes during the final recording sessions. Lastly, the abundance of tetrode leads means a larger scale to work with, in terms of preparing the implants and tetrodes, and inserting individual tetrodes into the respective areas (which is crucial to be done slowly for stability and good quality recordings). Nevertheless, the device allows for the study of multiple areas while the animal is allowed to freely move around, which is invaluable for answering certain questions in cognitive neuroscience such as those concerning spatial representations.

Contrasting to the Quad-drive, silicon probe recordings as performed in chapter 3 and 4 offer their own benefits and disadvantages. The difference between the ability to record chronically from multiple areas during free motion of the animal, using either tetrodes or silicon probes has diminished with recent advances in Neuropixel-drives (Bimbard et al., 2025; Melin et al., 2024). A strong benefit of a head-fixed setup is the ability to provide a more stable stimulation environment compared to freely moving rats or mice, offering a significant benefit in providing a temporally and spatially specific stimulation relevant for the study in question. While eye movements could be problematic for the retinotopic specificity of a visual stimulus in head-fixed

conditions, this can be appropriately countered by eye-tracking, or by presenting large enough stimuli, appropriate training or task-dependent fixation (McFarland, Cumming, & Butts, 2016, but note this study was in primates).

A noticeable disadvantage of head-fixed recordings may be the actual fixation of the animal: this prevents normal roaming behaviour. This may especially be the case when untrained, head-fixed animals try to perform movements which, in the head-fixed setup, do not elicit the normal sensorimotor responses in the brain. Well-habituated and trained animals may appear more relaxed during fixation, but an innate tendency to move would be not controlled for. Given the relevance of motor-related neural activity within sensory cortical areas (Chapter 3; Bimbard et al., 2023; Oude Lohuis et al., 2024), the non-ecological conditions of head-fixation can be improved to ensure the absence of signals not relevant for the hypothesis in question. Advancements in head-fixation setups with running capabilities, such as linear treadmills (Dombeck et al., 2007), Styrofoam balls (Kirkels et al., 2018), or flat-floored, air-lifted platforms (Kislin et al., 2014) address these issues, although such setups are still inhibiting small, specific movements, such as neck/head movements, which have been shown to modulate visual cortical activity (Bouvier, Senzai, & Scanziani, 2020).

In the context of motor-related signals in sensory areas, it is important to remember that, although these signals are sometimes called confounding, they may be alternatively interpreted as signals relevant for visual processing. Motor efference copies provide information to the visual system for appropriate processing (Leinweber et al., 2017; Pennartz et al., 2019). Nevertheless, the problem of erroneously interpreting this as cross-modal signals is significant, and future studies should aim to control for such activity.

Considerations on stimuli. Lastly, a consideration in the study of perception concerns the stimuli used with respect to the neural signals of interest. Neuroscientific studies often opt for low-dimensional stimuli: simple, non-natural stimuli with few parameters and low statistical complexity, that have been designed to elicit strong responses in the brain area in question. A great example in the visual domain is the use of the Gabor-patch (or related, moving gratings), which can be fine-tuned to hit the receptive fields of V1 neurons (Ringach, 2004). In the context of spatiotemporal control over the visual stimulus, this type of stimulus is suitable for eliciting strong responses in V1 and provides room to test specific hypotheses on sensory processing steps. Yet, a simple Gabor-patch is highly artificial and ecologically virtually irrelevant for mammals, and foregoes many ancillary and higher-order effects elicited by the entire visual scene that affect the activity of the neurons in question (David, Vinje, & Gallant, 2004). Conversely, complex, natural scenes are rich, with variable statistical features and complex non-linearities. Moreover, the proportion of neurons that can be recorded from is still miniscule compared to the total population processing such a stimulus, making it extremely difficult to process the data on natural scenes in such a way that valid inferences can be drawn on the functional processing being done by the recorded neurons. Both simple, low-dimensional and complex, natural stimuli provide

us with a different window on cortical processing, and each have their drawbacks, and it is understandable for specific research questions to resort to one type of stimulus over the other. However, to prevent the pitfalls of oversimplified hypotheses (for simple stimuli), or missing interesting components in cortical processing (for complex naturalistic stimuli), it would be advisable to investigate both stimuli (e.g., see Montijn et al., 2016).

5.2 A perspective on neural coding, consciousness and dynamic multimodal attractor states

In the final parts of this thesis, I would like to discuss one of the most intriguing subjects of cognitive neuroscience: consciousness. Neural coding phenomena in the brain as investigated by fundamental neurophysiology present us with parameters and/or functions which provide us with clues on the neurobiological structure necessary for consciousness. In the following section I will discuss the hypothetical implications of the results of Chapter 3 on phase-coding, continuing from the work of Pennartz (2009) where phase-coding was proposed to be used as a parameter for appropriately representing a multimodal, unified, stable conscious percept. The concept of consciousness as we will use it here will focus not so much on the "hard problem" (Chalmers, 1995), but on the way that an information processing system such as the brain can generate the unified, multimodal phenomenal experiences that are apparent in healthy, awake and behaving humans.

Consciousness is generally envisioned to be a unified experience, integrating multiple sensory modalities, stable but able to be dynamic in a changing world (Pennartz, 2015, 2022). If we want to describe the way information is represented in a system correlating with these features, we can conceptualize this as a set of attractor states. An attractor in dynamical systems theory is described as a state towards which a system tends to evolve over time, or where it stably settles. The simplest example would be a point-attractor in parameter space where the values of these parameters would tend to move towards one single value or range of values. A multimodal, unified, stable conscious percept could be described as a set of local attractor states converging to a relatively stable overall (global) state for the appropriate conscious experience at a given moment. An example to understand the stability, unification and dynamical properties of attractor states can be found in the intriguing case of binocular rivalry (Giles, Lau, & Odegaard, 2016): instead of overlapping or piecemeal images, healthy humans perceive only one stable image, until the parameters of the perceiving network shift enough such that the hypothesized attractor state can shift from one stable position to another one (likely this is facilitated by neurobiological constraints in the neural network representing the stimulus (see for instance Noest et al., 2007). Multimodal attractor states (figure 1) for the representation of conscious perception would be possible with Dynamic Associative Networks (DAN) which share features with neural networks, but also have the ability to generate attractor states at varying timescales (Pennartz, 2009). Pennartz (2009) stipulates that an attractor state representing a conscious percept would have to represent three types of information: 1) single object features, 2) position information, and 3) relations between features and modalities. If a consciousness experience can be represented as an attractor state based on these three types of information, which neural parameters contribute to its structure?

Parameter space for modalities, features and spatial representations

To understand the parameters of an attractor state, we need a clear understanding of what information needs to be represented in the brain to enable conscious perception.

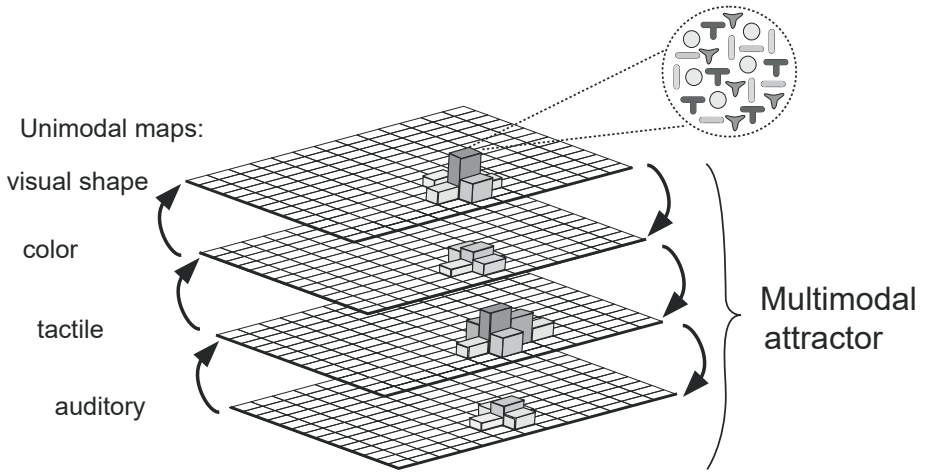

Figure 1: hypothetical scheme of a dynamic associative multimodal network built up from interconnected unimodal sensory maps, where a conscious state is interpreted as a multimodal attractor state across these maps. For convenience, this illustration does not show higher-order or subcortical inputs. Each layer represents a modality, from top to bottom: visual shape and color (two visual submodalities), the tactile modality (touch) and the auditory modality. Each map has a spatial structure, where each square in said spatial structure contains feature detectors relevant for said modality (inset). An attractor state as presented here can be visualized as a bump in activity (darker grey blocks; darker grey features in the inset) on a unimodal map (i.e. a unimodal attractor), whereas the overall state of the combined modalities is the global-multimodal attractor. From Pennartz, 2009.

Neuronal representations of specific sensory modalities like vision and hearing originate from distinct physical phenomena (photons vs. sound waves) with different formats, i.e. distinct physical phenomena that impinge on the sense organs in specific ways. However, in the brain, both collapse into specific electrical activity orchestrated by specifically connected neurons. This raises a question: if neurons integrate inputs from various origins, can a neuron's representation of a sensory modality ever be discrete and specific? A cell primarily tuned to certain visual orientations might still receive auditory signals (Ghazanfar & Schroeder, 2006; Ibrahim et al., 2016; Iurilli et al., 2012; Meijer et al., 2017; Oude Lohuis et al., 2024), reshaping its activity even if it might be insignificant to the current attractor state.

Let us envision modality specificity of a given neuron as a continuum in modality space (i.e. one dimension for each modality, ranging from zero influence from modality A on a given neuron until 100% influence from modality A). In this continuum, neurons cluster around the dominant sensory input they correlate with (such as neurons in the visual cortex receiving relatively direct photic information from the eyes, via the lateral geniculate nucleus of the thalamus) but remain influenced by peripheral signals originating from other sense organs. In this view, discrete modal categories (e.g., "vision") may reflect peaks in the spectrum of a modality-space, anchored by survival-

critical alignment with external reality, but they are not separate entities. Yet, in our subjective experience of reality, sensory modalities are generally experienced as clearly separate and distinct (although there are specific cases such as synesthesia where modality-specific representations reveal strong cross-modal influences).

In sensory cortical areas, we not only find cells generally responding to the appropriate sensory input, but also to specific features within a given sensory modality. Feature sensitivity can be interpreted as a computed sub-modality: such features are not directly derived from peripheral sensory receptors responsive to such a specific feature in the outside world but are computed to be statistically strongly correlated to the outside reality. For example, our photoreceptors respond to a "pixel" of retinotopic space, but through computations between nearby neurons in the visual cortex (and in earlier retinal and thalamic stages), these neurons can be specifically responsive to edges, orientations, etc. (Hubel & Wiesel, 1959; 1962; although how such computations have arisen over the course of evolution and in neurodevelopment remains a question for other scientific disciplines such as developmental neurobiology).

Neuroscience has discovered that features (in the broad sense of the term) can be strongly represented in the brain, ranging from the orientation of bars in the visual field, represented by cells in visual cortex (Hubel & Wiesel, 1959, 1962) to object-specific cells in inferotemporal cortex (Tanaka, 1996). While these neuronal responses are generally discrete when analysed experimentally (i.e. responding to one of the stimuli tested in a study, but not the other one), neural representations of such a wide breadth of features can again be conceptualized to be laid out in a feature-space, comparable to a modality-space for modalities: a given neuron can be said to have a certain level of responsivity to a particular orientation of an object, and to other features of that object.

Furthermore, spatial coding inherits a structure from physical reality if the sensory organs are equipped for it: the retina is organized in a 2D map corresponding directly to the spatial organization of photons impinging from the outside world, albeit with discontinuities (e.g., visual field edges), which is used throughout visual processing to form coherent spatial representations of reality (be it allocentric or egocentric). Three-dimensional space can apply to the senses in the tactile, proprioceptive and vestibular domain. 3D auditory space can be modelled by the shape of the pinna, ear movements in some animals, and by the use of inter-aural time and intensity differences, even though the cochlea only directly registers amplitudes of various frequencies in sound pressure waves, which lack an inherent spatial structure. Lastly, the chemoreceptors subserving olfaction and gustation do not have clear spatially organized sensors, yet even here we find ways to compute spatial locations of chemical sources with the spare spatial information fed into the sense organs (Bao et al., 2019; Porter et al., 2005; with a particularly interesting case for stereo-vomeronasal olfaction in snakes: Schwenk, 1994). Gustation would be difficult to envision spatially, it being concentrated within the oral cavity (although, as surprising as with olfaction, here too we find spatial information, see Spence, 2022). Based on these senses, we find many forms of spatial representations in the brain: for instance, the retinotopic organization in V1 correlating to space in

our visual field - at least if the eyes and head are fixated (Pennartz, Oude Lohuis, & Olcese, 2023), a spatial map of the body in somatosensory cortex, and place cells in the hippocampus correlating to the current subject's position in the environment (Herweg & Kahana, 2018). An interesting challenge for the brain (and for neuroscience) would be to relate this multitude of spatial representations to a coherent framework (cf. Pouget, Deneve, & Duhamel, 2002), yet this goes beyond the scope of this discussion.

Modality, feature and spatial information from the outside world are three dimensions that the brain must appropriately and dynamically represent. The challenge of an efficient neural code being used to represent a complete conscious experience would be to compress all perceptual parameters in a stable, specific way, with room for dynamics should the sensory system require this. In the model by Pennartz (2009), feature encoding is assigned to the firing rate activity of feature detectors, in relation to the location on spatial brain maps whereon these features are situated. In this model, the relations between modalities or features were proposed to be represented by phase coding. In the case of communication by firing rate alone, we encounter the significant caveat that unimodal representations might be disturbed if cross-modal interactions are communicated through firing rates as well (Pennartz 2015). As a visual experience normally stays the same and remains stable, in a visual-only and an audiovisual condition, the usage of phase-coding might provide the necessary part of the neural code to represent the linkage between the auditory and visual modality and thus serves to identify the modality by its pattern of connectivity in a multimodal topology (Pennartz, 2009, 2015). In Chapter 3, we investigated phase-coding as a potential mechanism for multisensory integration and asked whether it could be interpreted as an additional neural coding parameter establishing representational relations, leaving features of the stimulus to be coded by firing activity of specific neurons, and the location-bound parameters of stimuli by the spatial organization in the network. To briefly reiterate, we find stable firing rates but changing phases-of-firing between uni- and multimodal conditions, but interestingly, also between different visual orientations. Thus, firing rate may not be the only mode used for feature coding, and phase coding may co-occur with firing-rate coding in a multiplexed fashion. This needs to be examined more deeply in future work.

Phase-of-firing as a continuous, circular parameter for relational representations in conjunction with rate-based feature representations

Here I ask whether phase-of-firing could represent a continuous parameter used for representing relations between modalities. Phase-of-firing is a continuous parameter in a circular, bounded space (e.g., 0–360°). Some experimental results indicate there might be a physiological constraint to certain phases of LFP oscillations, for instance where neurons in the CA1 region of the hippocampus show a preference for firing just after the negative peak of the theta oscillation in the LFP (Kamondi et al., 1998, recorded from the CA1 pyramidal layer). Kamondi et al. discussed whether physiological mechanisms such as the interactions between strong dendritic inputs and synaptic inhibition can influence the spiking activity of a neuron, which implies limitations of the phase preferences to a particular phase range of the LFP. However, the results in Chapter 3

show the phase-of-firing of V1 neurons under visual and multisensory conditions to be evenly distributed across all phases. Results in auditory and visual cortex of the monkey by Kayser, Ince, and Panzeri (2012) seem to show spikes across all phases for all neurons with a slight (untested) tendency towards a bimodal distribution (although the reported distribution consisted of all recorded cells, without selection for phase-locking). This might indicate that, dependent on the area, the phase-of-firing can be constrained by different mechanisms. Whether this is related to the respective area's representations is not yet known.

If phase-of-firing is not strictly constrained to certain parts of the oscillation, we can ask whether the phase-of-firing is indeed a relevant parameter in representing the (co-)presence of modalities across areas. It is tempting to think that phase-coding as a mechanism to code for relations between modalities is as simple as assigning a specific phase to a specific modality. However, in Chapter 3 we find overlapping phase preferences between various sensory modalities. Instead, we may hypothesize that phase coding is more likely to be used as a comparative relational mechanism between sensory modalities, i.e. coding for relationships between modalities. Our results describe a subset of neurons changing phase-of-firing when an auditory stimulus is added while the visual stimulus and the related firing rates remain essentially the same. Should this same principle of phase change occur in the area representing the other modality ('other' as referenced to vision - for instance, in the auditory domain if we consider an audiovisual stimulus), this interplay on the level of primary sensory areas might lead to specific synchronous spiking in downstream areas, facilitating the concurrent representation of information in two modalities at a higher, multimodal level in said downstream areas. Although we can only speculate on the exact mechanism, the temporal dimension in which phase-coding operates would fit well in a mechanism that facilitates synchronization of the appropriately veridical, and behaviourally relevant information from different modalities.

Why then, did we find changes in phase-of-firing to visual orientations, as well as to different modalities? A straightforward explanation to this would be that, to the brain, features need to be related to each other as much as modalities. Considering the same problem of modality identification: any neuron in the brain would not "know" if the spike it receives would be related to visual or auditory information (Pennartz, 2009, 2015). Similarly, neurons in secondary visual areas, receiving information from orientation selective units in primary visual cortex, would not "know" if any spike is related to a particular orientation. This suggests that there might be a spectrum of relational representations, spanning from distinct sensory modalities for which information arrives in the brain through specific sensory organs, up to specific features that have been computed from the input from one such organ. A downstream neuron would not receive specific information on this distinction: each spike is simply an unlabelled spike. The anatomical distinctions are not to be ignored here: strong connectivity between certain sets of cortical areas implies exchange of information. Future studies will have to address whether this mechanism of phase-coding for constructing relational representations is indeed similar for neural coding of modalities and features.

Caveats

Here I would like to acknowledge that the framework sketched above is speculative and requires more research. Systems with multimodal attractor states that use rate, phase and spatial codes come with caveats. Four crucial neurobiological complexities need to be addressed:

1. As we found in chapter 3, a minority of neurons participates in phase coding. A sizeable portion of neurons shows phase-locking in certain conditions, but we find a small subpopulation of neurons (after selecting for appropriate phase-locking and stimulus responsivity) to show alterations in phase-of-firing when stimulus properties change. This might be related to the low dimensionality of the stimulus used, and how the neural networks we investigate have to be able to represent a wide variety of relations between modalities and/or features present in the natural world.

2. Secondly, we must consider the evolution of systems for perception and action across geological time. Natural selection does not act in a directed way. Individuals with traits that are sufficient for appropriate survival, but not necessarily perfect, tend to pass on said traits. In the context of the brain and neural coding, one can state that evolutionary pressure optimized our cognitive machinery to produce "good enough" representations. Mathematical purity is not a given, so this is not what we should be looking for.

3. Thirdly, our data from Chapter 3 has focussed exclusively on a strict and invariant definition of theta oscillations (specifically in the 10 Hz band). Oscillations in both lower (alpha, beta, theta) and higher (gamma) frequency ranges have been correlated with sensory stimulation, where higher frequencies carry less LFP power (i.e. considering the 1/f relation) and are typically considered signs of local synchronization, whereas lower frequencies are thought to have a larger spatial reach and are implicated in integrating information (Mazzoni, Logothetis, & Panzeri, 2012; see also the surprising results explaining several aspects of gamma with a simpel model: Spyropoulos et al., 2022). The spatially localized nature of gamma oscillations might support local synchronization, but does not necessarily support this across multiple areas, as would be the case for lower frequencies such as theta (Torres et al., 2019). Theta oscillations are suited ideally to support integrative mechanisms such as place cell sequences and means for cortico-hippocampal communication (Buzsáki, 2006; Wang, 2010; Womelsdorf et al., 2010), fitting with a role as 'clock' for providing a crucial phase-coding parameter.

   The range of theta oscillations is not necessarily stable (and the definition of 'theta' can vary between studies). Realistically, in a biological system variations in an oscillatory mechanism are to be expected, but are these variations related to each other, or do they simply reflect a noisy system? Shifts in the optimal frequency of theta have been related to task demand and stimulus input (Lowet et al., 2015; Senoussi et al., 2022). If theta frequency can be "shifted" to a more optimal

frequency (depending on cognitive requirements), this might indicate that the neurobiological system supporting the attractor state pushes the theta oscillation into a range ideally suited for representing the information in a suitable way. As the tasks in Senoussi et al. (2022) differed in difficulty, this might indicate the brain is actively trying to create a stable, unified representation ideally suited to complete the task. For researching phase-coding, we thus need a better understanding of neurophysiological mechanisms of LFP oscillations and phase-coding, in order to appropriately make claims on the functionality of LFP oscillations.

4. Lastly, as investigated in chapter 4 and by other colleagues (Bimbard et al., 2023; Kaplan & Zimmer, 2020; Oude Lohuis et al., 2024), the "contamination" of perceptual processing by behaviour-related activity will have to be addressed in future work investigating neural signals. This "contamination" may actually fit in a predictive processing framework, as motor commands can be used to probe the outside world, establish whether the animal's current world model is effective in guiding behavior and help create stable world representations in the face of eye and body movements. As mentioned before, this framework is well compatible with the concept of multimodal attractor states. Unfortunately, however, motor-related activity contributes another layer of complexity to the sensorimotor system of the brain that must be controlled for by the researcher.

Assuming that conscious representations have evolved in evolution to provide an animal with benefits for survival, we must be vigilant for other factors that we might be overlooking and might be necessary for constructing a neural code that is so strongly correlated with external reality (i.e. near-veridical). Importantly, besides the input from the sense organs, the brain has no access to the ground-truth of reality. Predictive Coding (Dayan et al., 1995; Friston, 2003; Pennartz, 2015, 2022; Rao & Ballard, 1999; Srinivasan, Laughlin, & Dubs, 1982) is a concept of interest for understanding sensory processing. Through interactions with the world, the process of minimizing prediction errors would help the brain establish an accurate (or "good enough") representation of reality. Predictive coding uses prediction errors (Keller & Mrsic-Flogel, 2018) which could pose a problem by adding another parameter into the multimodal attractor model. However, if neural representations are to be interpreted not as a direct representation of the outside world (which is impossible due to the outside world being unknown) but as corresponding to the subject's interpretation of the world, this would be compatible with the multimodal attractor model as mentioned above. Notably, efforts are ongoing to produce computational models in which predictive coding, neural oscillations and concepts like convergence into attractor states are being integrated (Lee, Pennartz, & Mejias, 2024).

Future research to establish the framework of multimodal attractor states can be taken in several directions: 1) acquire multi-area, large scale datasets to probe if phase-coding is a relevant parameter in other brain areas than V1, generating other representations (i.e. of more complex features); 2) assess the relevance of theta oscillations for phase- and temporal coding in general, which can be investigated by comparison of different

frequency bands. Additionally, intervention studies should be able to demarcate if theta acts as a necessary central clock for integrating or relating information from the sensory systems. This would also provide interesting clues as to whether the hippocampus is a crucial oscillator for cortical processing or if local cortical oscillators can be sufficient for providing a scaffold for phase-coding.

Concluding remarks

Consciousness is not static. The attractor states hypothesized as underlying stable perception of the world, will shift with the flow of sensory input and with changing internal goals, coupled to top-down prioritization of relevant information. Here, I would speculate a system having a "high information mass", or "density" if sizeable amounts of information about the outside (and inside, i.e. intra-body) world is encoded via firing rate, phase-of-firing, and spatial mapping. Such a "high information mass"' may be used to conceptualize levels of consciousness throughout the animal kingdom and the multimodally rich but unified conscious experience of humans.

Viewing conscious experience as a generative outcome from a multimodal set of attractor networks, engaging in predictive-coding driven representations, reconciles discrete conscious perception (i.e. experience of discrete inputs, objects, internal goals) with the underlying continuity of sensory signals, as it can define stable, unified states within a continuous sensory space. It positions the brain as a pragmatic sculptor of representations of reality, carving out useful discrete clusters of neuronal representations corresponding with percepts while retaining flexibility to switch from one stable state to the next. Lastly, the parameters of neural coding will likely interact with each other, as each of the parameters is in some way involved in setting the electrical activity of neurons and the way this activity is distributed across a network. These interactions would make for a complex, partially auto-correlated state (i.e., arising from recurrent interactions between such parameters) which might obfuscate the importance of individual parameters in the biological neural network. The code of cognition will likely be found in a conglomerate of neural activity parameters intrinsic to the system.