
OSCILLATIONS IN AN ARTIFICIAL NEURAL NETWORK CONVERT COMPETING INPUTS INTO A TEMPORAL CODE

Katharina Duecker*

Centre for Human Brain Health
School of Psychology, University of Birmingham
Birmingham, UK
katharina.duecker@gmail.com

Marco Idiart

Institute of Physics
Federal University of Rio Grande do Sul
Porto Alegre, Brazil
marco.idiart@ufrgs.br

Marcel AJ van Gerven

Donders Institute for Brain, Cognition and Behaviour
Radboud University
Nijmegen, the Netherlands
marcel.vangerven@donders.ru.nl

Ole Jensen

Centre for Human Brain Health
School of Psychology, University of Birmingham
Birmingham, UK
ole.jensen@bham.ac.uk

ABSTRACT

Deep convolutional neural networks (CNNs) resemble the hierarchically organised neural representations in the primate visual ventral stream. However, these models typically disregard the temporal dynamics experimentally observed in these areas. For instance, alpha oscillations dominate the dynamics of the human visual cortex, yet the computational relevance of oscillations is rarely considered in artificial neural networks (ANNs). We propose an ANN that embraces oscillatory dynamics with the computational purpose of converting simultaneous inputs, presented at two different locations, into a temporal code. The network was trained to classify three individually presented letters. Post-training, we added semi-realistic temporal dynamics to the hidden layer, introducing relaxation dynamics in the hidden units as well as pulsed inhibition mimicking neuronal alpha oscillations. Without these dynamics, the trained network correctly classified individual letters but produced a mixed output when presented with two letters simultaneously, elucidating a bottleneck problem. When introducing refraction and oscillatory inhibition, the output nodes corresponding to the two stimuli activated sequentially, ordered along the phase of the inhibitory oscillations. Our model provides a novel approach for implementing multiplexing in ANNs. It further produces experimentally testable predictions of how the primate visual system handles competing stimuli.

Keywords Neuronal oscillations · Multiplexing · Computer vision · Dynamical Systems

1 Introduction

The inclusion of convolution in artificial neural networks (ANNs) marked a significant milestone in computer vision [77]. Building on this innovation, deep convolutional neural networks (CNNs) have successfully addressed a wide range of image classification challenges, as demonstrated by Krizhevsky et al [76] in the ImageNet competition [see [114] for review]. Originally inspired by the receptive fields of neurons in visual cortex, the hierarchically organised representations emerging in these networks have been repeatedly shown to map on to those identified from human MEG and fMRI recordings of the visual ventral stream [24, 25, 42] and intracranial recordings from the non-human primate brain [73, 85, 100, 122, 121]. Despite the parallels between CNNs and the primate visual system, there are only few examples of ANNs for computer vision that have drawn inspiration from the temporal dynamics of cortical activity [31, 111]. For instance, alpha oscillations (8-12 Hz) dominate electrophysiological recordings from the human occipital lobe [1, 11], however, their computational benefit has so far not been explicitly explored in ANNs. Here, we show how

*corresponding author




Figure 1: Concept: An interplay between object-based attention and neuronal alpha oscillations implements a pipelining mechanism reflected by a temporal code. **a** Example: A passionfruit and an apple are competing for the processing resources of the visual system. **b** Inhibitory alpha oscillations modulate neuronal firing rhythmically. The alpha oscillations are illustrated by the black line, the aggregate neuronal responses to the passionfruit and the apple are shown as the deep purple and red lines, respectively. As the neurons responding to the passionfruit will receive stronger excitatory inputs, they will activate at an earlier phase of the alpha cycle compared to the apple. Refraction causes a momentary inactivation of the neural representation, allowing the neurons encoding the features of the apple to activate. As the apple is processed in feature-selective cortex (e.g. V4), its representation has already been passed on to the next stage of the hierarchy, the object-selective cortex (e.g. IT).

embracing dynamic activations in the hidden nodes of an ANN allows the network to process competing visual inputs. This work proposes a new framework combining principles from computational neuroscience and machine learning. We show that embedding biologically inspired neural dynamics in an ANN trained for image classification, enables the network to overcome bottleneck problems when presented with multiple stimuli.

Both CNNs and the visual system of primates have a converging architecture, wherein receptive field sizes expand progressively along the hierarchy [29, 34, 37, 51, 104]. Neurons in the early layers of CNNs resemble simple cells in the primary visual cortex and possess small receptive fields that detect edges and contours in the visual input [18, 77]. The receptive fields of neurons in later layers of the CNN have more expansive receptive fields, akin to the inferior temporal (IT) cortex that has been shown to be critically involved in core object recognition [28, 39, 90, 109]. This hierarchical architecture has been argued to give rise to a bottleneck problem [16]: while early visual cortex has been shown to process visual features in parallel [22, 118, 117], object recognition has been argued have a limited capacity, which implies that the semantics of multiple visual stimuli are extracted serially [62, 97]. It has been posited that the visual and auditory systems in human and non-human primates handle simultaneously presented stimuli through “multiplexing” [3], that is, by dynamically switching between the activity patterns associated with each stimulus [19, 79] [also see [96, 95] for review]. Through this mechanism, a single neuron can contribute to the neural code of multiple stimuli [74]. This dynamic interleaving of neural responses to simultaneous stimuli requires precise temporal coordination.

One mechanism through which the neural representations of multiple objects are organised in time is phase coding, which underlies a modulation of spiking activity by ongoing low-frequency neuronal oscillations [80, 93, 103] [also see [96, 95] for review]. For example, place representations, encoded in spatially distributed firing patterns, have been shown to be ordered along the phase of hippocampal theta (4-8 Hz) oscillations [57, 56, 61, 93, 103]. Several conceptual and computational models have extended on these ideas, suggesting that the items in working memory could be segregated into cycles of ongoing gamma oscillations [81, 82, 83]. This represents a multiplexed coding scheme [83].

It has been proposed that visual perception is supported by a similar mechanism, that underlies neuronal alpha oscillations. Alpha oscillations have long been known to reflect functional inhibition [58, 71]. The strength of the inhibition waxes and wanes along the alpha cycle, such that at the peak of the oscillation, strong inhibition reduces the probability of neuronal firing in the population [44, 52]. It has been proposed that only neurons receiving sufficiently strong excitatory inputs will be able to fire at the early phases of the alpha cycle [55, 60, 59]. As the inhibition reduces toward the trough of the cycle, the neurons may fire successively according to their excitability [55, 60, 59]. In this way, the interplay between inhibition by alpha oscillations and neuronal excitability generates a temporal code.

To visualise this process, imagine the following scenario at the supermarket: While searching for the ingredients for a passionfruit martini, your gaze lands on the key element: a passionfruit, placed next to an apple (Figure 1a). With your gaze fixating on both pieces of fruit, your visual system needs to find a way to represent and process them as

coherent but separate objects. As alluded to above, it has been proposed that ongoing inhibitory alpha oscillations in visual cortex serve to organise the representations of simultaneously presented stimuli competing for processing resources. It is well-documented that object-based attention is associated with increased neuronal excitability [8, 21, 66, 86, 92, 101]. Consequently, all neurons responding to the features of the passionfruit receive stronger excitatory inputs than the neurons responding to the apple and will thus overcome the alpha inhibition at an earlier phase (Figure 1b, top panel). Following the burst of activation associated with the passionfruit, refractory dynamics will momentarily deactivate its neural representation. In other words, excitation triggers inhibition, for instance, due to the activation of GABAergic [7, 72] or membrane properties such as the calcium-activated potassium current [108]. As the alpha inhibition decreases further, the neurons attuned to the apple will take over. As Figure 1b illustrates, these dynamics implement a temporal code, whereby the signals corresponding to the competing items activate successively along the alpha cycle. As the passionfruit is processed earlier in the alpha cycle, it will reach the next layer of the hierarchy with a temporal advantage over the apple. As the apple is processed in feature-selective cortex, e.g. V4, the passionfruit has already reached the object-selective cortex, e.g. IT cortex. In this way, the visual system might solve bottleneck problems through pipelining: multiple stimuli are processed in parallel, while their representations are segmented in time.

In the following, we will present how a mechanism inspired by these concepts can be integrated into an ANN, to allow a successive read-out of competing inputs. We will refer to our model as a dynamical artificial neural network (dynamical ANN). This work serves as a proof of principle to demonstrate the basic principles and analyse the ensuing dynamics in the network. Our model relates to previous work investigating emergent dynamical and oscillatory properties in systems for information storage [50] and Recurrent Neural Networks (RNNs) for image classification [31] and sequence learning [80]. In comparison to these works, we tune the dynamics of the system such that the network is able to segment the representations of competing stimuli along the phase of ongoing oscillations in the hidden units.

2 Methods

To implement a dynamical ANN we first trained a two-layer network on a simple image classification task. After training, we added biologically inspired dynamics to the hidden layers motivated by alpha oscillations in the human visual system. The dynamics were not included in the training process and did not change the weights of the network.

2.1 Network architecture

We consider a fully connected ANN with two hidden layers, consisting of 64 and 32 hidden nodes, respectively, and an output layer with three nodes (Figure 3a). A weight matrix of size 28×28 was applied to the input (56×56) with a stride of 28, such that each node in the first layer received $4 \times 28 \times 28$ inputs, ensuring representational invariance across the quadrants in the input. The activation h_j in each unit j of a hidden layer was calculated as:

$$h_j = \sigma(z_j) = \frac{1}{1 + e^{-a(z_j - b)}} \quad (1)$$

with z_j the input to each hidden unit. The input z_j arises from the activation in the previous layer according to $z_j = \sum_i w_{ji} z_i$, with w_{ji} being the weight matrix connecting nodes i and j . The slope of the sigmoid was set to $a = 2$, and the sigmoid was shifted by the bias term $b = 2.5$. These parameters were fixed, such that a small input ($z \approx 0$) would result in an activation h close to 0, while strong inputs well above 2.5, would result in an h close to 1. Consequently, the activations in the hidden layers were approximately binary (“all or none”, see Figure 3d). The activation in each output node was calculated using the softmax function

$$o_j = \text{softmax}(z_j) = \frac{e^{z_j}}{\sum_{j=1}^K e^{z_j}} \quad (2)$$

converting the inputs z_j into K probabilities in the output layer [40]. The weights of the network were initialised according to a uniform distribution within the range $[-x, x]$, where $x = \sqrt{\frac{6}{n_{\text{in}} + n_{\text{out}}}}$ with n_{in} and n_{out} being the number of inputs and outputs to the current layer, respectively (Glorot initialization) [38]. The Adam optimiser was chosen to minimise the mean squared error loss using gradient descent [69]. The network weights were learned by backpropagating the error through the network layers (as mentioned above, the bias term was fixed at $b = -2.5$).

2.2 Network dynamics in the hidden layers

We aimed to implement multiplexing by oscillatory dynamics into a fully connected ANN. To this end, after training the network on an image classification task, we added biologically inspired non-spiking dynamics to each node in the

hidden layer, expressed by non-linear ordinary differential equations (ODEs). The ODEs were solved using the Euler method with a fixed time step of $\Delta t = 0.001$ s. The rate of change in each hidden unit j was defined as:

$$\tau_h \frac{dh_j}{dt} = -h_j + \sigma \left(\frac{z_j - r_j - \alpha(t) + h_j}{s} \right) \quad (3)$$

where τ_h determines the timescale at which h_j approaches the sigmoid activation (see Equation 1). The relaxation term r_j was applied to reduce the activation of each hidden node, ensuring an intermittent activation. The dynamics of the relaxation term are given by

$$\tau_r \frac{dr_j}{dt} = -r_j + c \cdot h_j \quad (4)$$

with time constant τ_r defining the delay with which r approaches h . Rhythmic inhibition, mimicking inhibitory neuronal alpha oscillations in the visual system [60], was integrated into each h_j , implemented as a 10 Hz sine wave

$$\alpha(t) = m \cdot [1 + \sin(2\pi t \cdot 10)] \quad (5)$$

and subtracted from the input z_j ; with amplitude m being adjustable to modulate the strength and offset of the inhibition.

The selection of the parameters was informed by electrophysiological and computational constraints. For instance, the timescale of the activation τ_h was set to 0.01 s, in accordance with the membrane time constant of excitatory neurons of 10-30 ms [17, 46, 113]. The time constant of the refractory term was chosen to be $\tau_r = 0.1$ s akin to afterhyperpolarization effects caused by calcium-activated potassium currents [99, 119]. The effect of parameters c and s on the dynamics will be explored below (see Figure 3).

2.3 Fixed points of the system

The fixed points (steady-state) of h and r are defined by setting Equations 3 and 4 to zero. Solving for h and r yields:

$$h^* = \frac{z}{c-1} + \frac{s}{c-1} \left[\frac{\log\left(\frac{1}{h^*} - 1\right)}{a} - b \right] \quad (6)$$

$$r^* = c \cdot h^* \quad (7)$$

We found the fixed points numerically, using the function `fsolve` implemented in SciPy [112]. For small values of $s \ll 1$, the equilibrium points could be approximated as $h^* = \frac{z}{c-1}$ and $r^* = \frac{c \cdot z}{c-1}$. Unless otherwise specified, the dynamics of all hidden nodes in the full network were initialised at these approximated fixed points with $z = 4$ (the average input size to the first layer).

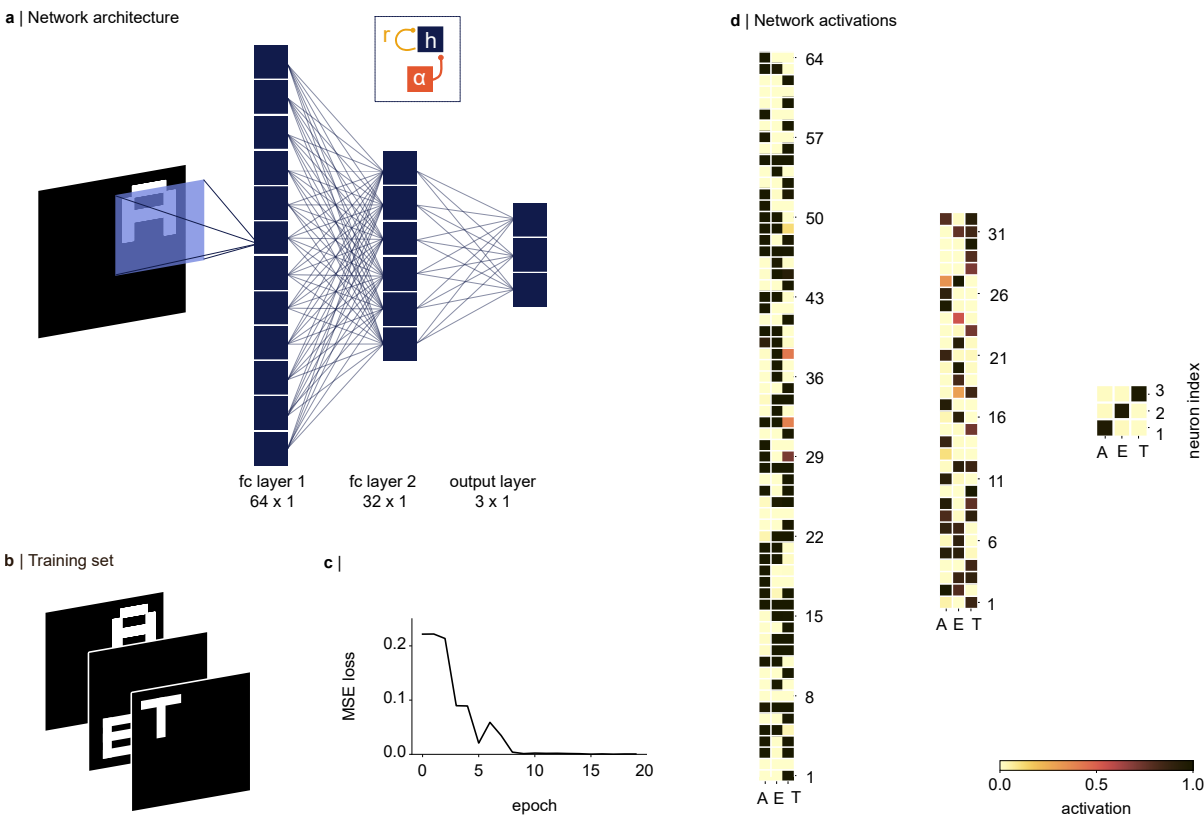


Figure 2: The classification problem and network architecture. **a** A network with two fully connected hidden layers was trained to classify three letters presented on a 56×56 image, in one of four quadrants (28×28). Convergence of the quadrants in the first hidden layer was implemented by sliding a 28×28 weight matrix over the image, with a stride of 28. After the training, we added a refractory term r , and pulses of inhibitory alpha oscillations $\alpha(t)$ to each hidden node h , as shown in the inset, and Equations 3 and 4. **b** Example inputs from the training set. **c** The network learned to classify the three letters within 10 epochs, as indicated by the mean-squared error (MSE) loss approaching 0. **d** Activations in the hidden layers and the output node in response to three inputs, presented in the three columns. The shifted sigmoid (see main text) resulted in approximately binary activations in the hidden layers. The activations in the output node demonstrate that the inputs are classified correctly.

3 Results

The aim of this project was to show that biologically inspired dynamics will allow a neural network to handle competing inputs, despite having been trained on one stimulus at a time. To illustrate these dynamics, we trained the network to classify three letters "A", "E", and "T" (Figure 3a,b). The network learned to read out the images correctly within 10 training epochs (Figure 3c), with the softmax activation in the output layer approaching an activation of 1 in the node of the corresponding image (Figure 3d). Post-training, we integrated the full dynamical system described in Equation 3 and Equation 4, into the hidden layers. We will first investigate the behaviour of the individual hidden nodes.

3.1 Network stability and parameters

Figure 3a demonstrates how the interplay between the input $z = 6.5$ and refraction by r (orange trace) resulted in dynamical activations in h (purple trace) with a period of approximately 100 ms. As h increases, r increases with a delay. As r opposes h , this results in refraction. The adjustable parameters c and s in Equation 3 and Equation 4 were explored to identify values resulting in robust self-sustained dynamics at a frequency of about 10 Hz in the hidden nodes. Figure 3b shows how the frequency of h changes as a function of c , for $s = 0.05$, for different levels of z (ranging from 0.5 to 6.5 as indicated by the yellow to dark brown colour scale). The parameter c influences how strongly r grows depending on h (see Equation 4). Consequently, the frequency of h increases monotonically for increasing values of c . For large values of z , the frequency changes approximately exponentially with c . To induce oscillations in the 10-12 Hz

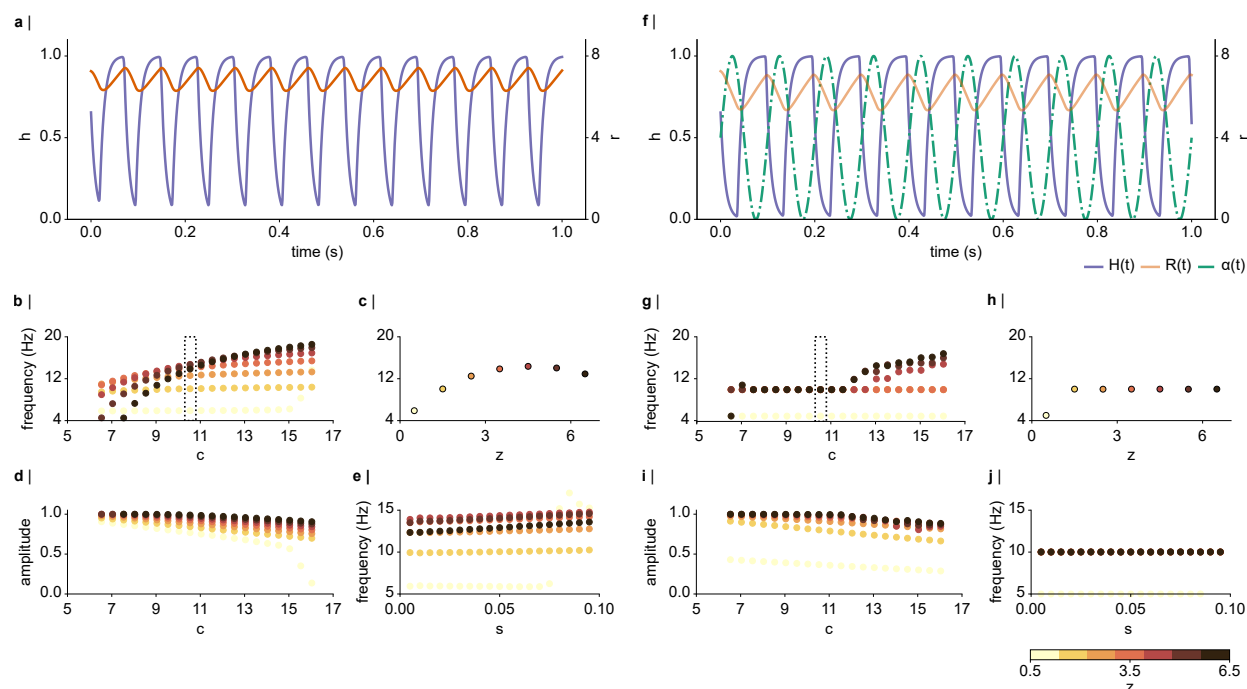


Figure 3: **a** Exemplary dynamics in h for one input $z = 6.5$, with $c = 10$ and $s = 0.05$. The interplay between h (purple trace) and r (orange trace) leads to dynamics with a period of approximately 100 ms. **b** Frequency as a function of c for different input values z (colour coded from 0.5 to 6.5). High values of z require larger values of c to generate oscillatory dynamics in h . **c** A value of $c = 10$ (indicated by the dotted box in **b**) leads to dynamics in the 8-12 Hz range for $z > 0.5$, whereby the frequency of the dynamics increases approximately monotonically with input size z . **d** Amplitude of h as a function of c , showing that for high values of c , h will not reach an activation of 1. Values of $c < 11$ seem appropriate to induce dynamics in the 10-12 Hz range that reach the sigmoid activation of 1. **e** Frequency of h as a function of s , indicating that the dynamics only change marginally for all values of s within a range of 0.005 to 0.1. Only for very small inputs ($z = 0.5$), a large s leads to fast dynamics of up to 15 Hz. **f** The dynamics of h and r shown in **a** are entrained by the 10 Hz alpha inhibition (dotted green line). **g-j** The alpha inhibition stabilises the dynamics in the 10 Hz range. **g** For large values of $c > 11$, the dynamics are able to escape the periodic inhibition and oscillate at a faster frequency. For small values of $z = 0.5$, the dynamics in h skip one alpha cycle, and oscillate at 5 Hz. **h** frequency as a function of z , for $c = 10$, showing that all nodes with inputs $z > 0.5$ are entrained by the 10 Hz rhythm. **i** The amplitude of h as a function of c in the presence of the alpha inhibition. The amplitude tends to be slightly larger than the sigmoid activation for small values of z . **j** In presence of alpha inhibition, h is robustly entrained to the 10 Hz rhythm, for all values of s .

range for a broad range of inputs z , we selected $c = 10$ (indicated by the dotted box). Figure 3c demonstrates how the frequency changes as a function of input size z for $c = 10$, showing that the nodes tend to oscillate at faster rhythms for larger inputs. There is a slight tendency for the frequency to plateau for very large values of z . This is due to r having to oppose a strong input while growing as a function of activation h , which is bounded at 1 (Equation 1). As a result, the frequency does not increase further and eventually decreases for large values of z .

Figure 3d demonstrates that the amplitude of h decreases as a function of c . We aimed to induce dynamics such that h would oscillate between 0 and the sigmoid activation of the current input, which for large values of z is close to 1. This amplitude was achieved for the selected $c = 10$. Note that for $z < 4.5$ the amplitude of h was slightly larger than the sigmoid activation of z . For instance, for $z = 1.5$ the activation is $h = \sigma(-2 \cdot (1.5 - 2.5)) \approx 0.12$, however, the amplitude reached values of about 0.8 (see Figure 3d). We did not find this to cause problems when integrating the dynamics into the network, as most activations in the hidden layers were outside the linear part of the sigmoid, and thus approached an activation of 1 (Figure 3c, layer 1; see below for details).

Another instrumental parameter in Equation 3 is s , which serves to scale the various input parameters. A small s will effectively increase the steepness of the sigmoid, resulting in a more step-like response. Figure 3e depicts the frequency of the node as a function of s and suggests that values ranging from 0.005 to 0.1 result in dynamics of about 10 Hz for a range of inputs z . Based on these observations, we selected $s = 0.05$.

We next explored the effects of adding pulsed inhibition at 10 Hz to the hidden dynamics ($\alpha(t)$ in Equation 3). As depicted in Figure 3f, the alpha inhibition (green dashed line) entrained the dynamics, such that h activated in anti-phase to the 10 Hz rhythm. Another effect of the alpha inhibition is that r oscillates between lower values than before. This is sensible, as the alpha inhibition and refraction r work together to reduce the activations in h . Figure 3g depicts the frequency of h as a function of c , showing that in the presence of the alpha inhibition, all hidden nodes oscillate at a frequency of 10 Hz, when $5 < c \leq 11$, for the current input range of $0.5 < Z \leq 6.5$. For large values of c ($c > 11$) the dynamics escape the alpha rhythm when z is sufficiently large (> 5) and oscillate at a faster rate. For small inputs ($z = 0.5$), h appears to skip one alpha cycle at a time and follows a 5 Hz rhythm. The frequency in h as a function of input size z is shown in the right panel, confirming the 10 Hz entrainment. Figure 3i shows the amplitude of h as a function of c , demonstrating that the amplitude is most stable for $c < 11$. Notably, the amplitude of h again reaches values above the sigmoid activation of z (as described above), however, we did not find this to interfere with the dynamics in the full network. Lastly, the exact value of s within the 0.005 to 0.1 range did not change the frequency, for all $z > 0.5$ (Figure 3j).

We conclude that the entrainment and refractory dynamics of the nodes were stable for a large range of the parameters. Based on the presented simulations, we settled on the parameters $\tau_h = 0.01$, $\tau_y = 0.1$, $c = 10$, $s = 0.05$ in all following simulations, as these produced robust dynamics at approximately 10 Hz and an activation h close to 1 for a wide range of inputs.

3.2 Alpha oscillations stabilise the dynamics in a two-layer Neural Network

After training the network on the “A-E-T” classification problem and exploring the behaviour of the individual nodes shown in Figure 3, we simulated the dynamics in the full network while keeping the weights connecting the layers fixed.

Figure 4a and b show the dynamics in the output and hidden layers of the network, in response to a single input letter, A, with $\tau_h = 0.01$, $\tau_r = 0.1$, $c = 10$, $s = 0.05$, but without any drive from the alpha oscillation ($m = 0$, Equation 5). The output node corresponding to A (blue trace) oscillates at approximately 12 Hz (Figure 4a). The leftmost panel in Figure 4b depicts the activations in h in layer 1 as a function of time, for z 's ranging from 0.5 to 6.5 (only unique values of z are shown). The dynamics demonstrate that the phase of the oscillations depends on the size of input z , leading to inconsistent phase delays between the network nodes. Due to the softmax activation introducing lateral inhibition in the output nodes, these phase delays cause a spurious intermittent activation of the output node corresponding to letters “E” and “T” (orange and green traces). The right panel in Figure 4b, depicting h as a function of r , shows the limit cycle and fixed points (indicated by the diamond-shaped scatters) for each input z . All units demonstrate a limit cycle behaviour but with different amplitudes.

The bottom panel shows the dynamics in the second layer. While the phase-locking between the hidden nodes in the second layer is slightly stronger compared to layer 1, activations in the nodes receiving smaller inputs appear to lag the nodes receiving larger inputs. The right-most plot, showing the relationship between h and r , again suggests a limit cycle behaviour, although the amplitude of each node appears to vary over the course of the simulation. This is likely due to the second layer receiving dynamical, non-phase-locked inputs from the first layer.

Introducing oscillatory entrainment by periodic inhibition into each layer of the network stabilises the dynamics in the entire network. Figure 4c shows the dynamics in the output layer, with the amplitude of the oscillatory drive set to $m = 0.5$, and a phase delay of $\Delta\phi = 0$ between the layers. Comparison to the dynamics without the oscillatory drive shows that the inhibition removes the spurious activation in the output nodes corresponding to “E” and “T” (orange and green trace). This stabilisation of the read-out underlies increased synchrony both within and between the hidden layers, as indicated in Figure 4d. In particular, the phase-locking between the activations in the first layer has been notably increased by the oscillatory drive; as shown in the time course of the activations in h shown in Figure 4d, top left. Comparison of the limit cycles shown in Figure 4b top right, and Figure 4d top right, demonstrates a wider limit cycle in presence of the alpha oscillations, reflecting an increased amplitude of the dynamics in h .

The dynamics in the second layer also show higher synchrony and stabilised trajectories h and r in presence of the alpha oscillations (Figure 4d, bottom right). This is explained by the more synchronous inputs from the first layer, and the alpha oscillations applied to the second layer. Comparison of the plots at the bottom right in Figure 4b and d reveals that the amplitude of the activation in the second layer varies less in presence of the alpha inhibition. These simulations show that the dynamics of the multi-layer network are dramatically stabilised by the oscillatory alpha inhibition.

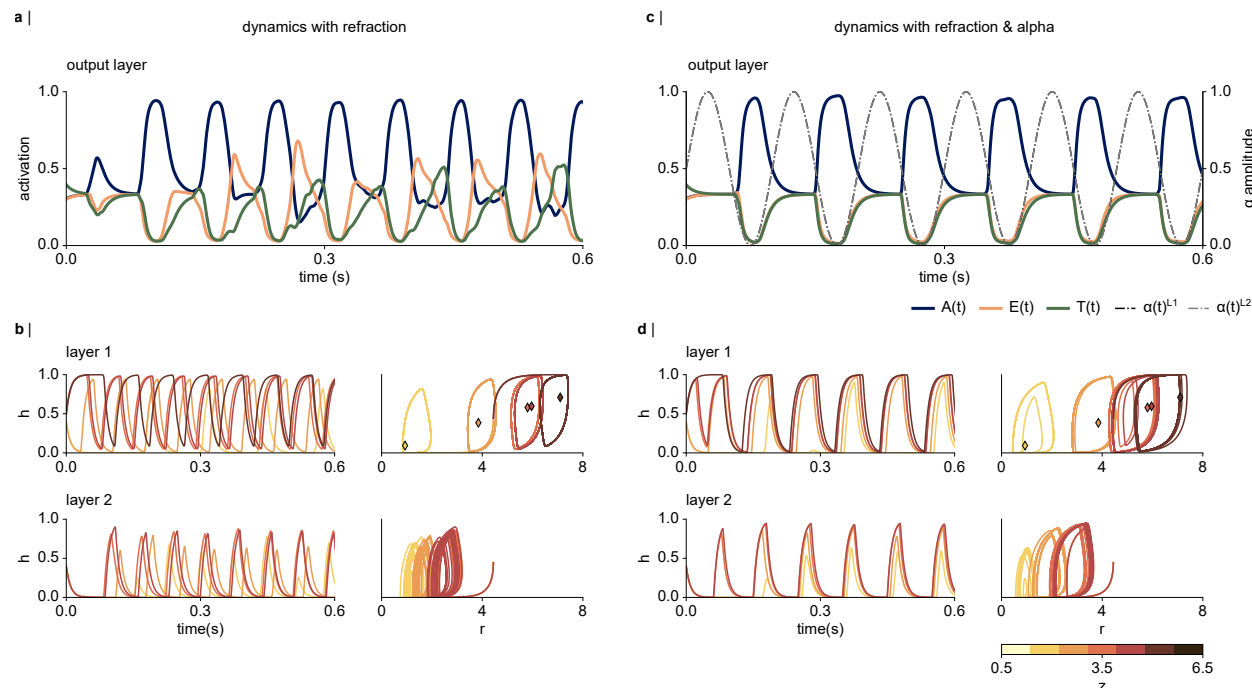


Figure 4: **a-b** Dynamics underlying the interplay between h and refraction r (Equation 3 and 4, without the alpha inhibition). **a** The output node corresponding to the letter A oscillates at about 12 Hz. The system spuriously activates the output node E and T. **b** (top left) The activations in the first hidden layer show that the frequency of the oscillation depends on the input size z , as suggested in Figure 3c. (top right) Trajectory of h as a function of r , with the fixed points indicated by the coloured diamonds. The dynamics are attracted towards a stable limit cycle, depending on the magnitudes of the input, and orbits around the fixed point. (bottom left) Activation of h in layer 2 as a function of time. There is a notable phase lag between the network nodes based on the input size z . (bottom right) The amplitude of the activation in each node appears to vary over time. **c-d** Periodic inhibition stabilises the dynamics in the system. **c** Periodic inhibition at 10 Hz with an amplitude of $m = 0.5$ was added to each layer. The read-out of the presented letter oscillates in anti-phase to the alpha inhibition. **d** Alpha oscillations notably stabilise the dynamics within and between the network layers, as indicated by the phase-locking between the nodes, and the stabilised H-R trajectories for the activations in layer 2.

3.3 Simultaneous presentation of two inputs produces a temporal code

3.3.1 The bottleneck problem in absence of the dynamics

The network correctly classified individually presented stimuli after the training, as demonstrated in Figure 3c. However, when presented with two stimuli simultaneously (Figure 5a), the network produced a mixed output. The right panel in 5a shows the network activations in response to all possible inputs of stimuli. Figure 6b shows the corresponding time course to these simultaneous inputs, which was achieved using Equation 3, with $s = 1$, $r = 0$, $\alpha(t) = 0$, with the network dynamics initialised at $h = r = 0$ in all hidden nodes. The network distributes the activations in the output layer over the nodes corresponding to the respective inputs. This suggests that the output layer produces a weighted average of the activations to both stimuli. This indicates a bottleneck problem when the network needs to classify two stimuli. In contrast, the abilities of our visual systems to recognise a stimulus do not decrease with the number of objects. As such, the visual system must be able to segment the representations of the different stimuli. As outlined above, it has been proposed that multiplexing in the visual system may underly oscillatory dynamics. In the following, we will explore how the complete dynamical network presented in Figure 4 and described by Equations 3 and 4 responds to simultaneous stimuli.

3.3.2 Refraction and alpha inhibition allow read-out of competing stimuli as a temporal code

Figure 6b shows how the output layer of the dynamical network responds to two simultaneous inputs (“A” and “E”, and “T” and “E”). To mimic spatial attention to letter E, we multiplied the pixels defining the letter E with 1.2, and




Figure 5: The bottleneck problem when the network receives two inputs simultaneously. **a** (left column) Exemplary input combinations of the two letters. (right) Activations in the output layer are distributed approximately evenly over the respective nodes. **b** Time course of the activation in the output layer, for $s = 1$, $r = 0$, and $\alpha(t) = 0$. The shared output is a consequence of the two items competing.

the pixels defining the letters A and T with 0.8 (Figure 6a). Figure 6b shows the dynamics in the output layer, based on the simulations with parameters $\tau_h = 0.01$, $\tau_r = 0.1$, $c = 10$, $S = 0.05$, without any oscillatory drive. The node corresponding to “E” activates first and dominates the dynamics in the first 150ms (orange trace). Subsequently, the network starts to alternate between the two inputs, whereby the nodes corresponding to the presented letters reach activation levels between 0.9 and 1. As such, the interplay between activation and refraction implements a multiplexing mechanism. However, the dynamics of the multiplexed representations are unstable. For instance, in Figure 6b, top panel, the node corresponding to the letter E (orange trace) activates over longer periods of time and reaches higher activations than the node representing A (blue).

Introducing the oscillatory alpha inhibition results in stabilised dynamics, whereby the simultaneous inputs are read out as a temporal code, organised along the alpha phase (Figure 6c). Notably, the attended stimulus “E” activates first, within the first alpha cycle. In the second cycle, the system starts to generate the phase code, which further stabilises with time. We replicated these dynamics for all combinations of simultaneous inputs (two letters at a time), as shown in Figure S1.

A more detailed investigation of the multiplexing dynamics in the different layers is shown in Figure 7; for the exemplary simultaneous input E and T (as presented in Figure 6a, bottom panel). The top panels in Figure 7a and b indicate how strongly the hidden representations to the combined input correspond to the activations to the individually presented inputs. For instance, for letter “E” (orange trace) this measure of similarity was calculated as:

$$s_E(t) = \frac{\sum_{j=1}^K (h_j^{ET}(t) \odot h_j^E)}{\sum_{j=1}^K (h_j^E \odot h_j^E)} \quad (8)$$

with h_j^{ET} being the activation in hidden node j at time point t , to the simultaneously presented letters T and E, h_j^E being the activation in hidden node j to the letter E in the trained, non-dynamical network (see Figure 3c), and \odot being the Hadamard operator (element-wise multiplication). This measure can be interpreted as a normalised dot product.

The time course of the normalized dot product indicates that the similarity between the activations in the first layer, and the representations to both letters E and T oscillates in antiphase to the alpha inhibition (Figure 7a). Notably, the similarities to both letters follow the same time course, indicating that the first layer does not segment the competing inputs. Moreover, the activations to the attended letter extend over the entire half-cycle, while the activations of the unattended letter are more short-lived. The bottom panel shows a raster plot of the magnitude of the activation in each neuron over time. The simultaneous presentation of two letters activates a large fraction of the hidden nodes in the first

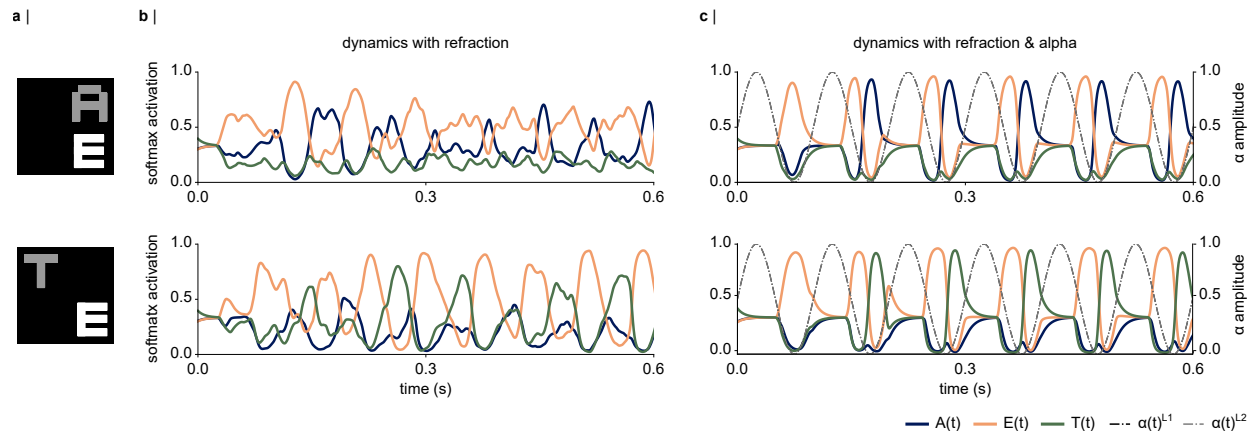


Figure 6: The dynamical artificial neural network (dynamical ANN) multiplexes simultaneously presented stimuli. **a** Examples of the simultaneous inputs. Attention towards the letter E was mimicked by increasing the input of all pixels belonging to E while reducing the inputs of the letters A and T. **b** The interplay between the excitation and refraction results in dynamic activations in the output layer, whereby one of the two letters is read-out at a time. However, the dynamics show notable instability, and the node corresponding to the “attended” input E activates over longer periods of time. **c** Introducing pulses of inhibition (alpha oscillations) into the network generates a temporal code in the output layer. The system first activates the node corresponding to the attended letter E. Following that, the letters are read out as a code, ordered along the phase of the alpha oscillation according to the input gain.

layer (Figure 7a, bottom). Note the rhythmic silencing of the network activations by the alpha inhibition, that affects most, but not all nodes. This implies that the first layer activates to both stimuli in parallel.

In comparison, the activations in the second layer demonstrate that the nodes responding to each letter are activated in succession: the normalised dot product between the current representations and the activations to an individual letter “E” (orange trace) precede the ones corresponding to letter T (green trace, Figure 7b). The bottom panel in Figure 7b indicates that a smaller fraction of the network is activated at each time point, and the successive activation of the hidden nodes can be observed. Finally, Figure 7c shows the read-out in the output layer, confirming that the representations of E and T are fully separated after the second cycle of the alpha inhibition (also see Figure 6c).

In sum, our simulations show how integrating dynamics driven by excitation and refraction enables a fully connected neural network to multiplex simultaneous inputs – a task it has not been trained on explicitly. This mechanism is further stabilised by pulses of inhibition, akin to alpha oscillations in the human visual system.

3.4 Making and breaking the temporal code: the effect of phase delay between the layers

Previous research on neural codes in the visual and auditory system has suggested that the phase of low-frequency oscillations in electrophysiological recordings carries information about the sensory input [67, 84, 91]. The communication through coherence (CTC) theory predicts that the phase relationship between two populations is critical for their communication [35, 36] (also see [2, 87] for computational implementations). While CTC was initially proposed for oscillations in the gamma-band [35, 36], related ideas have been explored for the alpha-band [14]. Based on these concepts, we next tested how different phase delays between the network layers impact the temporal code. Figure 8 shows the temporal codes for different phase delays between the alpha oscillations. Visual inspection suggests that the most robust temporal code of the two competing letters is achieved when the oscillations are synchronous in the two layers (top left), or when the second layer lags or precedes the first layer by 10 ms (top, second from left and bottom right). Notably, the temporal code breaks for larger phase delays between the layers. One might have expected an anti-phase delay between layers 1 and 2 to cut off all communication, however, the network is still able to identify at least one of the presented stimuli, albeit with high instability.

In sum, these simulations show that the shape of the temporal code, and the number of items that can be read-out, strongly varies based on the phase delay between the layers and is most robust for synchronous oscillations in both layers.

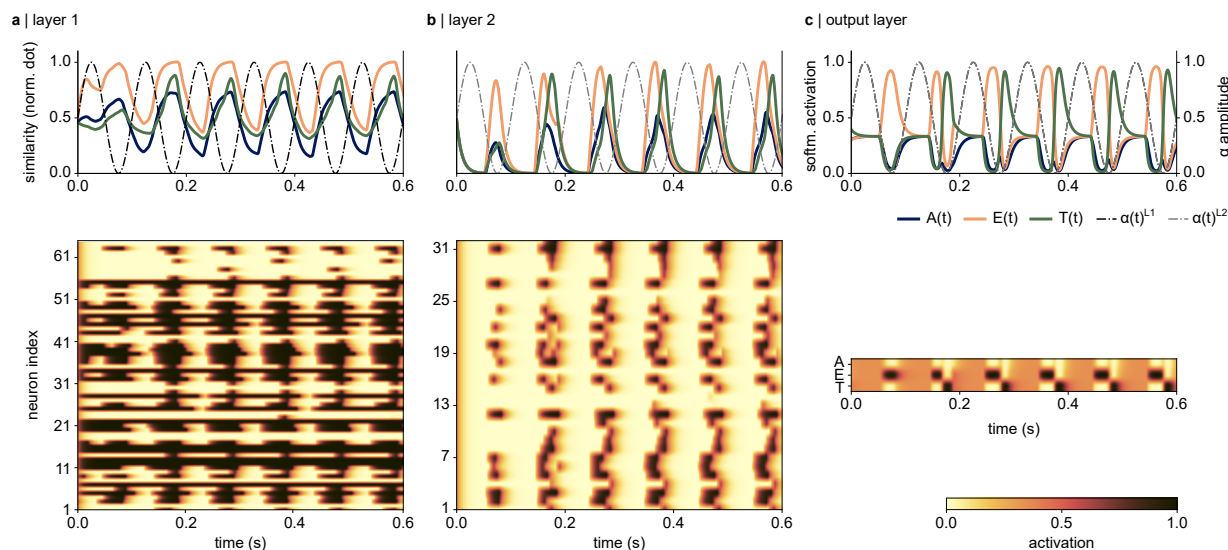


Figure 7: The representations of the competing inputs are segmented in the second layer of the dynamical Neural Network in response to an input image containing the letters E and T, as shown in Figure 5a, bottom. **a** (top) Layer 1: The nodes corresponding to both the letter E and the letter T activate in antiphase to the alpha inhibition, with no segregation in time. Notably, the network activates the representations to the attended letter E more strongly, and over longer periods of time. (bottom) The network activations over time suggest that a large portion of the hidden nodes activates in antiphase to the alpha inhibition. **b** Segregation of the simultaneous stimuli can be observed in the second hidden layer. (top) The normalised dot product suggests that the nodes corresponding to each input activate at different phases relative to the alpha inhibition. (bottom) The temporal segmentation observed in the top panel can also be observed as a successive activation of the nodes in the network. The alpha inhibition silences the activation in the entire layer. **c** The softmax activation and the output layer reflects a temporal code, whereby the inputs are read out along the phase of the alpha inhibition, ordered according to the magnitude of the input.

4 Discussion

We here demonstrate how integrating semi-realistic neuronal dynamics into an ANN enables multiplexing during image classification. The network was first trained to classify individual letters presented in quadrants. Post-training, the network correctly recognised the individual inputs. When presented with two letters simultaneously, however, the output layer produced a mixed representation of both stimuli due to the bottleneck in the network. Adding refractory dynamics to the nodes of the network resulted in alternating activations to the simultaneously presented inputs, suggesting that the network was able to segregate the representations. However, the dynamics expressed notable instability and asynchrony within and between layers. Adding oscillatory inhibition to the layers, akin to alpha oscillations in the visual system, stabilised the dynamics in the hidden layers. When two inputs were presented simultaneously, the interplay between activation, refraction, and pulses of inhibition resulted in a stable multiplexed code, whereby the output nodes were activated sequentially based on the strength of the input.

Our simulations provide an implementation of the idea that inhibitory alpha oscillations in visual cortex serve to support the processing of simultaneously presented stimuli, by segmenting them into a temporal code [60, 59]. As the inhibition is strongest at the peak of the oscillations, only neurons receiving sufficiently strong excitatory inputs will be able to activate [55, 44, 52]. In turn, the neural representations associated with the attended stimulus will overcome the inhibition at an earlier phase of the alpha cycle than the unattended object. This allows temporal segmentation and thus multiplexing of simultaneously presented stimuli [60].

While it has been shown that simple visual features can largely be processed in parallel [22, 118, 117], object recognition has been demonstrated to be supported by serial processes [63, 97]. This indicates a bottleneck problem which has been argued to arise from the converging hierarchical structure of the visual system [16, 59, 97, 117]. By segmenting individual object representations in time, alpha oscillations have been suggested to allow the attended stimulus to pass through the network layers with a temporal advantage over the unattended stimulus (Figure 1, also see [59]). These parallel and serial properties resulting from multiplexing can be observed in our network simulation as well. In the first layer, the stimuli are represented in parallel, while the second layer begins to segment them along the phase of the

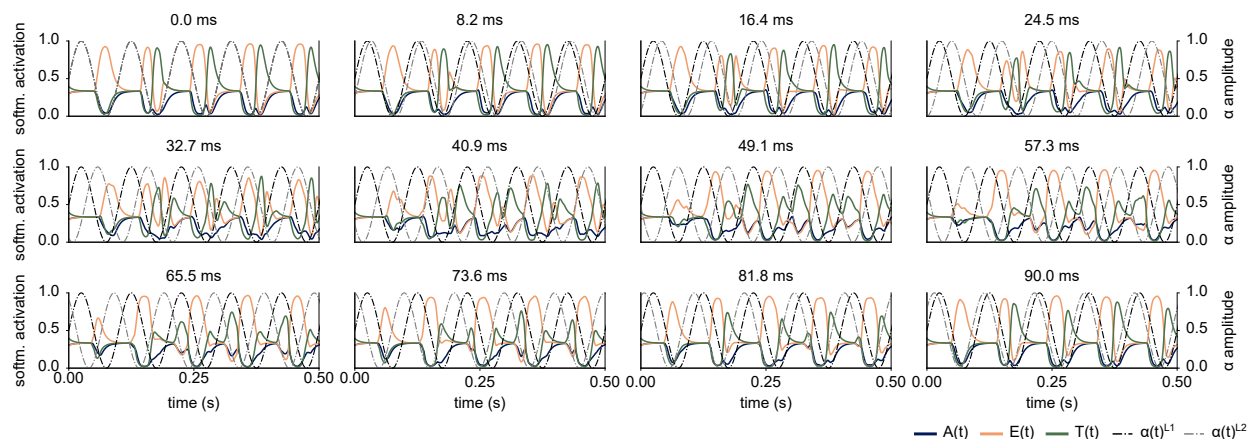


Figure 8: The temporal code as a function of phase delay between the layers, indicated in the title (in ms). The cleanest read-out of the two competing stimuli as a temporal code is achieved by a phase lag of zero (top left) and if the second layer precedes or lags the first layer by 10 ms.

oscillation (Figure 7). These simulations show how integrating biologically inspired dynamics into a neural network results in a system that can multiplex simultaneous stimuli while embracing the parallel and serial properties of visual object recognition.

Evidence for a phase-coding mechanism similar to the one presented here has been robustly observed in recordings from the rodent and human hippocampus, whereby spiking activity has been shown to be modulated along the phase of ongoing theta oscillations (4-8 Hz [61, 64, 93, 98, 103]). The order in which a sequence of inputs has been experienced, has further been proposed to be preserved in the spiking activity [61, 93] but see [80].

Intracranial recordings from the visual cortex in non-human primates have revealed that the phase of spontaneous alpha oscillations modulates spiking activity [12, 13, 30, 44] and neuronal gamma oscillations [10, 15, 88, 106]. In this context, alpha oscillations have been argued to organize visual processing in a top-down manner [68]. Intracranial and MEG recordings from the human brain have replicated the observed phase-amplitude coupling between gamma and alpha oscillations during visual processing [94, 115]. Recent human EEG recordings have additionally posited that alpha waves travelling from occipital to frontal areas are actively involved in visual processing [4, 5] (but see [125] for a critical perspective). Similarly, ECoG recordings from marmoset visual cortex have revealed travelling waves in the dorsal and ventral stream that were linked to visual performance [27] and a modulation of neural processing along the visual hierarchy following saccade initiation [65]. These reports show that alpha oscillations may propagate over cortical areas involved in visual processing to coordinate neural activity. It has so far not been determined, however, if this modulation results in successive activations of the neural representations in line with our model simulations.

Our simulations result in two testable predictions. First, we propose that neural representations activate along the phase of spontaneous alpha oscillations, ordered according to attention or salience [60, 55]. This prediction can be tested using electrophysiological recordings during visual tasks with more than one stimulus. The neural representations of each stimulus could be extracted from these data using decoding methods such as multivariate pattern analysis [49] or linear discriminant analysis [43]. For instance, using MEG, van Es et al. [32] have recently investigated the effect of ongoing alpha oscillations on the decoding accuracy of visual stimuli in a spatial attention task. The authors have demonstrated that the phase of alpha oscillations in the frontal eye field and parietal cortex of the human brain modulated the performance of the decoder. However, a phase delay between the attended and unattended stimuli has not explicitly been reported. Alternatively, these predictions could be tested based on intracranial recordings from the mouse brain. Using Neuropixels probes, spiking activity and local field potentials can be simultaneously recorded intracranially from several cortical areas in mice (and non-human primates, [53, 107]). It is well-established that the mouse visual system exhibits a hierarchical structure similar to the one observed in primates [116, 102]. As such, these data could be used to test whether spiking activity is segmented along the phase of ongoing alpha oscillations, for instance, to distinguish a figure from a background [70].

The second prediction of our model is that a stable temporal code emerges from approximately synchronous oscillations in the consecutive network layers. Recent work using concurrent iEEG and MEG recordings has suggested that interactions between alpha oscillations in prefrontal cortex and mediodorsal thalamus mediate visual performance [41]. In light of the literature on travelling alpha waves [6, 4, 5, 9, 27, 65, 124] this begs the question of whether

neuronal processing along the visual hierarchy is controlled by one driving force as suggested by our simulations, e.g. the thalamus or prefrontal cortex, or a travelling wave propagating forward or backwards along the visual hierarchy. With recent advances in brain-wide recordings in mice using Neuropixels probes [53], it may be possible to investigate whether the driving force of these travelling waves can be established.

Our network relates to previous computational models that have explored the role of biologically plausible dynamics for multiplexing and inter-areal communication [3, 54, 82, 87, 110]. We expand on this work by demonstrating how multiplexing and communication through synchronous oscillatory activity can enhance the computational versatility of neural network in the context of multi-item image classification. As we aimed to provide a proof-of-principle, we trained the network on a comparably simple classification problem involving a small training set. The non-monotonic loss function in Figure 3b implies that the network might have learned to solve the problem by memorizing the inputs and could thus show low generalisability to new inputs [23]. Our aim is to expand the presented principles to CNNs with a deeper architecture that can solve benchmark image classification problems such as (E)MNIST [26, 78], CIFAR-10 [75], and ImageNet [120]. One technical detail to consider is that modern DNNs typically implement non-linearities using the rectified linear unit (ReLU) function which reduces the vanishing gradient problem in deep architectures and speeds up learning [23, 40]. Since ReLUs are not bind the activations between 0 and 1, a different set of ODEs will be needed to describe the dynamics in future versions of this model.

Alternatively, the dynamics could be integrated in additional layers with sigmoid-like activation functions, between the trained network layers, as conventionally done in spiking neural networks [105, 123]. For instance, Sørensen et al. [105] integrated spiking dynamics into a pre-trained CNN, which allowed the network to find a target stimulus in a complex natural image, an ability the model did not exhibit without the spiking dynamics. As alpha oscillations have been shown to modulate spiking activity [12, 13, 44], it would be interesting to understand to what extent oscillatory dynamics could serve to modulate activations in spiking neural networks. By incorporating spiking or non-spiking dynamics into extra layers with activations constrained between 0 and 1, the concepts presented here could be explored in pre-existing deep neural networks. In sum, future versions of this network will expand to deeper architectures and modern image classification benchmarks.

The rate of change in the network nodes was defined by a set of ODEs, following conventional practice in computational neuroscience [89]. ODEs have also found applications in the development of RNNs. For instance, in the form of neural ODEs [20], liquid-time constant neural networks [47, 48], and RNNs consisting of damped oscillators [31]. These networks had great success in learning long-range dependencies in time series data [20, 47, 48], sequences of images [80], and image classification when the pixels of the input are transformed into a time series (sequential MNIST) [31, 47]. The goal of our work, in comparison, was to convert spatially presented inputs into a time series. The dynamics in our model therefore do not provide information about the dependencies between inputs but rather serve to segment them in time. However, integrating biologically plausible dynamics into the training process would undoubtedly be interesting. Oscillatory activity in the theta, alpha, and gamma-band has been repeatedly shown to support learning and memory [for review see [33, 45]]. In line with this, Effenberger, Carvalho, Dubinin, and Singer [31] have shown that an RNN whose nodes reflect damped oscillators outperform the performance of non-oscillatory RNNs when classifying the sequential MNIST data set. For future extensions of the presented algorithm, it would be interesting to explore if and how biologically plausible dynamics could be used to support the training process.

5 Conclusion

We here present a proof-of-concept showing that integrating oscillatory dynamics based on excitation, refraction, and pulses of inhibition into the hidden nodes of an ANN enables multiplexing of competing stimuli, even though the network was only trained to classify individual inputs. Our simulations predict that the visual system of humans and non-human primates handles the processing of multiple stimuli by organising their neural representations along the phase of inhibitory alpha oscillations. These predictions can be experimentally tested using simple attention paradigms and electrophysiological recordings in humans, non-human primates, and rodents. Future versions of the network will include extensions to deeper architectures and modern image classification benchmarks.

6 Supplementary Material

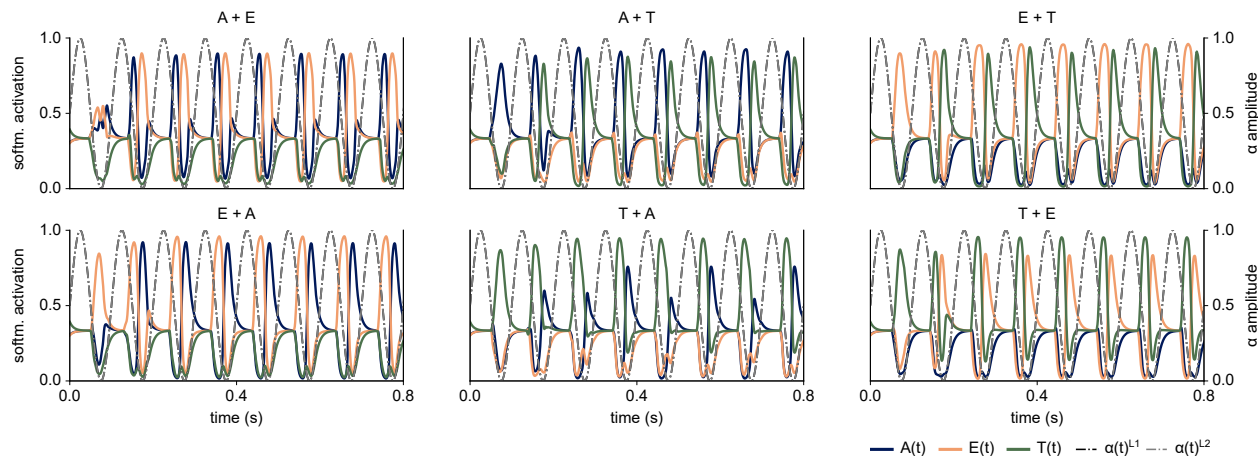


Figure S1: Examples of the temporal code for all input combinations. The first letter in the title is the one for which input gain has been increased.

References

- [1] Edgar D. Adrian and Bryan H. C. Matthews. “The Berger Rhythm: Potential Changes from the Occipital Lobes in Man”. In: *Brain* 57.4 (Dec. 1, 1934), pp. 355–385. ISSN: 0006-8950. DOI: 10.1093/brain/57.4.355. URL: <https://doi.org/10.1093/brain/57.4.355> (visited on 07/10/2023).
- [2] Thomas E. Akam and Dimitri M. Kullmann. “Efficient “communication through coherence” requires oscillations structured to minimize interference between signals”. In: *PLoS computational biology* 8.11 (2012), e1002760. ISSN: 1553-7358. DOI: 10.1371/journal.pcbi.1002760.
- [3] Thomas E. Akam and Dimitri M. Kullmann. “Oscillatory multiplexing of population codes for selective communication in the mammalian brain”. In: *Nature Reviews. Neuroscience* 15.2 (Feb. 2014), pp. 111–122. ISSN: 1471-0048. DOI: 10.1038/nrn3668.
- [4] Andrea Alamia, Lucie Terral, Malo R. D’Ambra, and Rufin VanRullen. “Distinct roles of forward and backward alpha-band waves in spatial visual attention”. In: *eLife* 12 (Mar. 6, 2023), e85035. ISSN: 2050-084X. DOI: 10.7554/eLife.85035. URL: <https://elifesciences.org/articles/85035> (visited on 06/20/2023).
- [5] Andrea Alamia and Rufin VanRullen. “A Traveling Waves Perspective on Temporal Binding”. In: *Journal of Cognitive Neuroscience* (May 8, 2023), pp. 1–9. ISSN: 0898-929X. DOI: 10.1162/jocn_a_02004. URL: https://doi.org/10.1162/jocn_a_02004 (visited on 06/20/2023).
- [6] Andrea Alamia and Rufin VanRullen. “Alpha oscillations and traveling waves: Signatures of predictive coding?”. In: *PLOS Biology* 17.10 (Oct. 3, 2019). Publisher: Public Library of Science, e3000487. ISSN: 1545-7885. DOI: 10.1371/journal.pbio.3000487. URL: <https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.3000487> (visited on 10/18/2023).
- [7] Jeffrey Anderson et al. “Stimulus dependence of two-state fluctuations of membrane potential in cat visual cortex”. In: *Nature Neuroscience* 3.6 (June 2000), pp. 617–621. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/75797. URL: http://www.nature.com/articles/nn0600_617 (visited on 04/29/2020).
- [8] Lawrence G. Appelbaum and Anthony M. Norcia. “Attentive and pre-attentive aspects of figural processing”. In: *Journal of Vision* 9.11 (Oct. 19, 2009), p. 18. ISSN: 1534-7362. DOI: 10.1167/9.11.18. URL: <https://doi.org/10.1167/9.11.18> (visited on 10/18/2023).
- [9] Ali Bahramisharif et al. “Propagating Neocortical Gamma Bursts Are Coordinated by Traveling Alpha Waves”. In: *Journal of Neuroscience* 33.48 (Nov. 27, 2013), pp. 18849–18854. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.2455-13.2013. URL: <https://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.2455-13.2013> (visited on 05/30/2022).

- [10] Andre M Bastos, Julien Vezoli, and Pascal Fries. “Communication through coherence with inter-areal delays”. In: *Current Opinion in Neurobiology*. SI: Brain rhythms and dynamic coordination 31 (Apr. 1, 2015), pp. 173–180. ISSN: 0959-4388. DOI: 10.1016/j.conb.2014.11.001. URL: <http://www.sciencedirect.com/science/article/pii/S0959438814002165> (visited on 05/02/2020).
- [11] Hans Berger. “Über das Elektrenkephalogramm des Menschen”. In: *Archiv für Psychiatrie und Nervenkrankheiten* 87.1 (Dec. 1, 1929), pp. 527–570. ISSN: 1433-8491. DOI: 10.1007/BF01797193. URL: <https://doi.org/10.1007/BF01797193> (visited on 06/22/2023).
- [12] Anil Bollimunta, Yonghong Chen, Charles E. Schroeder, and Mingzhou Ding. “Neuronal mechanisms of cortical alpha oscillations in awake-behaving macaques”. In: *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 28.40 (Oct. 1, 2008), pp. 9976–9988. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.2699-08.2008.
- [13] Anil Bollimunta, Jue Mo, Charles E. Schroeder, and Mingzhou Ding. “Neuronal mechanisms and attentional modulation of corticothalamic α oscillations”. In: *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 31.13 (Mar. 30, 2011), pp. 4935–4943. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.5580-10.2011.
- [14] Mathilde Bonnefond, Sabine Kastner, and Ole Jensen. “Communication between Brain Areas Based on Nested Oscillations”. In: *eNeuro* 4.2 (Mar. 1, 2017). Publisher: Society for Neuroscience Section: Theory/New Concepts. ISSN: 2373-2822. DOI: 10.1523/ENEURO.0153-16.2017. URL: <https://www.eneuro.org/content/4/2/ENEURO.0153-16.2017> (visited on 08/17/2023).
- [15] Conrado A. Bosman et al. “Attentional stimulus selection through selective synchronization between monkey visual areas”. In: *Neuron* 75.5 (Sept. 6, 2012), pp. 875–888. ISSN: 1097-4199. DOI: 10.1016/j.neuron.2012.06.037.
- [16] Donald E. Broadbent. *Perception and communication*. Perception and communication. Pages: v, 340. Elmsford, NY, US: Pergamon Press, 1958. v, 340. DOI: 10.1037/10037-000.
- [17] György Buzsáki. “Neural Syntax: Cell Assemblies, Synapsembles, and Readers”. In: *Neuron* 68.3 (Nov. 4, 2010), pp. 362–385. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2010.09.023. URL: <https://www.sciencedirect.com/science/article/pii/S0896627310007658> (visited on 02/28/2023).
- [18] Matteo Carandini. “Do We Know What the Early Visual System Does?” In: *Journal of Neuroscience* 25.46 (Nov. 16, 2005), pp. 10577–10597. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.3726-05.2005. URL: <http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.3726-05.2005> (visited on 02/12/2021).
- [19] Valeria C. Caruso et al. “Single neurons may encode simultaneous stimuli by switching between activity patterns”. In: *Nature Communications* 9.1 (July 13, 2018). Number: 1 Publisher: Nature Publishing Group, p. 2715. ISSN: 2041-1723. DOI: 10.1038/s41467-018-05121-8. URL: <https://www.nature.com/articles/s41467-018-05121-8> (visited on 08/24/2022).
- [20] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. “Neural Ordinary Differential Equations”. In: *Advances in Neural Information Processing Systems*. Vol. 31. Curran Associates, Inc., 2018. URL: <https://proceedings.neurips.cc/paper/2018/hash/69386f6bb1dfed68692a24c8686939b9-Abstract.html> (visited on 04/07/2022).
- [21] Yanqing Chen, Anil K. Seth, Joseph A. Gally, and Gerald M. Edelman. “The power of human brain magnetoencephalographic signals can be modulated up or down by changes in an attentive visual task”. In: *Proceedings of the National Academy of Sciences* 100.6 (Mar. 18, 2003). Publisher: Proceedings of the National Academy of Sciences, pp. 3501–3506. DOI: 10.1073/pnas.0337630100. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.0337630100> (visited on 10/18/2023).
- [22] Yuzhi Chen and Eyal Seidemann. “Attentional Modulations Related to Spatial Gating but Not to Allocation of Limited Resources in Primate V1”. In: *Neuron* 74.3 (May 10, 2012). Publisher: Elsevier, pp. 557–566. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2012.03.033. URL: [https://www.cell.com/neuron/abstract/S0896-6273\(12\)00344-3](https://www.cell.com/neuron/abstract/S0896-6273(12)00344-3) (visited on 08/17/2023).
- [23] Francois Chollet. *Deep learning with Python*. Simon and Schuster, 2021.
- [24] Radoslaw M. Cichy, Aditya Khosla, Dimitrios Pantazis, and Aude Oliva. “Dynamics of scene representations in the human brain revealed by magnetoencephalography and deep neural networks”. In: *NeuroImage* 153 (June 2017), pp. 346–358. ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2016.03.063. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5542416/> (visited on 03/22/2023).

- [25] Radoslaw M. Cichy et al. “Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence”. In: *Scientific Reports* 6.1 (June 10, 2016). Number: 1 Publisher: Nature Publishing Group, p. 27755. ISSN: 2045-2322. DOI: 10.1038/srep27755. URL: <https://www.nature.com/articles/srep27755> (visited on 10/08/2020).
- [26] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and André van Schaik. “EMNIST: Extending MNIST to handwritten letters”. In: *2017 International Joint Conference on Neural Networks (IJCNN)*. 2017 International Joint Conference on Neural Networks (IJCNN). ISSN: 2161-4407. May 2017, pp. 2921–2926. DOI: 10.1109/IJCNN.2017.7966217. URL: https://ieeexplore.ieee.org/abstract/document/7966217?casa_token=VYVpqIorsuMAAAAA:DUdR0y9jm9Sstzh4r8C25IBSwKWxf5RYq-OuMcNUPhh8o4pgjBvBY9JBtWGCLirImhtzy_Wd (visited on 10/23/2023).
- [27] Zachary W. Davis et al. “Spontaneous Traveling Cortical Waves Gate Perception in Behaving Primates”. In: *Nature* 587.7834 (Nov. 2020), pp. 432–436. ISSN: 0028-0836. DOI: 10.1038/s41586-020-2802-y. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7677221/> (visited on 10/23/2023).
- [28] James J. DiCarlo and David D. Cox. “Untangling invariant object recognition”. In: *Trends in Cognitive Sciences* 11.8 (Aug. 1, 2007), pp. 333–341. ISSN: 1364-6613. DOI: 10.1016/j.tics.2007.06.010. URL: <https://www.sciencedirect.com/science/article/pii/S1364661307001593> (visited on 01/20/2022).
- [29] James J. DiCarlo, Davide Zoccolan, and Nicole C. Rust. “How does the brain solve visual object recognition?”. In: *Neuron* 73.3 (Feb. 9, 2012), pp. 415–434. ISSN: 1097-4199. DOI: 10.1016/j.neuron.2012.01.010.
- [30] Kacie Dougherty et al. “Ongoing alpha activity in V1 regulates visually driven spiking responses”. In: *Cerebral Cortex* 27 (2017). Publisher: Oxford University Press, pp. 1113–1124. DOI: 10.1093/cercor/bhv304.
- [31] Felix Effenberger, Pedro Carvalho, Igor Dubinin, and Wolf Singer. *A biology-inspired recurrent oscillator network for computations in high-dimensional state space*. Pages: 2022.11.29.518360 Section: New Results. Nov. 29, 2022. DOI: 10.1101/2022.11.29.518360. URL: <https://www.biorxiv.org/content/10.1101/2022.11.29.518360v1> (visited on 07/06/2023).
- [32] Mats W. J. van Es et al. “Phasic modulation of visual representations during sustained attention”. In: *European Journal of Neuroscience* 55.11 (2022). _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/ejn.15084>, pp. 3191–3208. ISSN: 1460-9568. DOI: 10.1111/ejn.15084. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/ejn.15084> (visited on 08/17/2023).
- [33] Juergen Fell and Nikolai Axmacher. “The role of phase synchronization in memory processes”. In: *Nature Reviews Neuroscience* 12.2 (Feb. 2011). Number: 2 Publisher: Nature Publishing Group, pp. 105–118. ISSN: 1471-0048. DOI: 10.1038/nrn2979. URL: <https://www.nature.com/articles/nrn2979> (visited on 09/28/2023).
- [34] Daniel J. Felleman and David C. Van Essen. “Distributed hierarchical processing in the primate cerebral cortex”. In: *Cerebral Cortex (New York, N.Y.: 1991)* 1.1 (Feb. 1991), pp. 1–47. ISSN: 1047-3211. DOI: 10.1093/cercor/1.1.1-a.
- [35] Pascal Fries. “A mechanism for cognitive dynamics: neuronal communication through neuronal coherence”. In: *Trends in cognitive sciences* 9.10 (Oct. 2005). Publisher: Trends Cogn Sci. ISSN: 1364-6613. DOI: 10.1016/j.tics.2005.08.011. URL: <https://pubmed.ncbi.nlm.nih.gov/16150631/> (visited on 09/30/2023).
- [36] Pascal Fries. “Rhythms For Cognition: Communication Through Coherence”. In: *Neuron* 88.1 (Oct. 7, 2015), pp. 220–235. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2015.09.034. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4605134/> (visited on 07/17/2023).
- [37] Ricardo Gattass et al. “Cortical visual areas in monkeys: location, topography, connections, columns, plasticity and cortical dynamics”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 360.1456 (Apr. 29, 2005). Publisher: Royal Society, pp. 709–731. DOI: 10.1098/rstb.2005.1629. URL: <https://royalsocietypublishing.org/doi/full/10.1098/rstb.2005.1629> (visited on 09/22/2023).
- [38] Xavier Glorot and Yoshua Bengio. “Understanding the difficulty of training deep feedforward neural networks”. In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics. ISSN: 1938-7228. JMLR Workshop and Conference Proceedings, Mar. 31, 2010, pp. 249–256. URL: <https://proceedings.mlr.press/v9/glorot10a.html> (visited on 09/23/2023).
- [39] Melvyn A. Goodale and A. David Milner. “Separate visual pathways for perception and action”. In: *Trends in Neurosciences* 15.1 (Jan. 1992), pp. 20–25. ISSN: 0166-2236. DOI: 10.1016/0166-2236(92)90344-8.
- [40] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.

- [41] Benjamin J. Griffiths et al. “Rhythmic interactions between the mediodorsal thalamus and prefrontal cortex precede human visual perception”. In: *Nature Communications* 13.1 (June 29, 2022). Number: 1 Publisher: Nature Publishing Group, p. 3736. ISSN: 2041-1723. DOI: 10.1038/s41467-022-31407-z. URL: <https://www.nature.com/articles/s41467-022-31407-z> (visited on 10/05/2023).
- [42] Umut Güçlü and Marcel A. J. van Gerven. “Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream”. In: *Journal of Neuroscience* 35.27 (July 8, 2015), pp. 10005–10014. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.5023-14.2015. URL: <http://www.jneurosci.org/cgi/doi/10.1523/JNEUROSCI.5023-14.2015> (visited on 10/06/2020).
- [43] Matthias Guggenmos, Philipp Sterzer, and Radoslaw M. Cichy. “Multivariate pattern analysis for MEG: A comparison of dissimilarity measures”. In: *NeuroImage* 173 (June 1, 2018), pp. 434–447. ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2018.02.044. URL: <https://www.sciencedirect.com/science/article/pii/S1053811918301411> (visited on 10/11/2023).
- [44] Saskia Haegens et al. “ α -Oscillations in the monkey sensorimotor network influence discrimination performance by rhythmical inhibition of neuronal spiking”. In: *Proceedings of the National Academy of Sciences* 108.48 (2011), pp. 19377–19382.
- [45] Simon Hanslmayr, Bernhard P. Staresina, and Howard Bowman. “Oscillations and Episodic Memory: Addressing the Synchronization/Desynchronization Conundrum”. In: *Trends in Neurosciences* 39.1 (Jan. 1, 2016). Publisher: Elsevier, pp. 16–25. ISSN: 0166-2236, 1878-108X. DOI: 10.1016/j.tins.2015.11.004. URL: [https://www.cell.com/trends/neurosciences/abstract/S0166-2236\(15\)00259-3](https://www.cell.com/trends/neurosciences/abstract/S0166-2236(15)00259-3) (visited on 09/28/2023).
- [46] Kenneth D. Harris et al. “Organization of cell assemblies in the hippocampus”. In: *Nature* 424.6948 (July 2003). Number: 6948 Publisher: Nature Publishing Group, pp. 552–556. ISSN: 1476-4687. DOI: 10.1038/nature01834. URL: <https://www.nature.com/articles/nature01834> (visited on 08/16/2023).
- [47] Ramin Hasani et al. “Closed-form continuous-time neural networks”. In: *Nature Machine Intelligence* 4.11 (Nov. 2022). Number: 11 Publisher: Nature Publishing Group, pp. 992–1003. ISSN: 2522-5839. DOI: 10.1038/s42256-022-00556-7. URL: <https://www.nature.com/articles/s42256-022-00556-7> (visited on 12/07/2022).
- [48] Ramin Hasani et al. *Liquid Time-constant Networks*. Dec. 14, 2020. arXiv: 2006.04439[cs, stat]. URL: <http://arxiv.org/abs/2006.04439> (visited on 12/07/2022).
- [49] James V. Haxby, Andrew C. Connolly, and J. Swaroop Guntupalli. “Decoding neural representational spaces using multivariate pattern analysis”. In: *Annual Review of Neuroscience* 37 (2014), pp. 435–456. ISSN: 1545-4126. DOI: 10.1146/annurev-neuro-062012-170325.
- [50] John J. Hopfield. “Neural networks and physical systems with emergent collective computational abilities.” In: *Proceedings of the National Academy of Sciences* 79.8 (Apr. 1982). Publisher: Proceedings of the National Academy of Sciences, pp. 2554–2558. DOI: 10.1073/pnas.79.8.2554. URL: <https://www.pnas.org/doi/10.1073/pnas.79.8.2554> (visited on 10/27/2023).
- [51] David H. Hubel and T. N. Wiesel. “Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex”. In: *The Journal of Physiology* 160.1 (Jan. 1962), pp. 106–154.2. ISSN: 0022-3751. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1359523/> (visited on 07/12/2020).
- [52] Luca Iemi et al. “Ongoing neural oscillations influence behavior and sensory representations by suppressing neuronal excitability”. In: *NeuroImage* 247 (Feb. 15, 2022), p. 118746. ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2021.118746. URL: <https://www.sciencedirect.com/science/article/pii/S1053811921010181> (visited on 05/30/2022).
- [53] International Brain Laboratory et al. *A Brain-Wide Map of Neural Activity during Complex Behaviour*. preprint. Neuroscience, July 4, 2023. DOI: 10.1101/2023.07.04.547681. URL: <http://biorxiv.org/lookup/doi/10.1101/2023.07.04.547681> (visited on 10/23/2023).
- [54] Ole Jensen. “Information Transfer Between Rhythmically Coupled Networks: Reading the Hippocampal Phase Code”. In: *Neural Computation* 13.12 (Dec. 2001). Conference Name: Neural Computation, pp. 2743–2761. ISSN: 0899-7667. DOI: 10.1162/089976601317098510.
- [55] Ole Jensen, Mathilde Bonnefond, and Rufin VanRullen. “An oscillatory mechanism for prioritizing salient unattended stimuli”. In: *Trends in Cognitive Sciences* 16.4 (Apr. 2012), pp. 200–206. ISSN: 1879-307X. DOI: 10.1016/j.tics.2012.03.002.
- [56] Ole Jensen and John E. Lisman. “Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells”. In: *Learning & Memory (Cold Spring Harbor, N.Y.)* 3.2 (1996), pp. 279–287. ISSN: 1072-0502. DOI: 10.1101/lm.3.2-3.279.

- [57] Ole Jensen and John E. Lisman. “Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase coding”. In: *Journal of Neurophysiology* 83.5 (May 2000), pp. 2602–2609. ISSN: 0022-3077. DOI: 10.1152/jn.2000.83.5.2602.
- [58] Ole Jensen and Ali Mazaheri. “Shaping Functional Architecture by Oscillatory Alpha Activity: Gating by Inhibition”. In: *Frontiers in Human Neuroscience* 4 (2010). ISSN: 1662-5161. URL: <https://www.frontiersin.org/article/10.3389/fnhum.2010.00186> (visited on 05/05/2022).
- [59] Ole Jensen, Yali Pan, Steven Frisson, and Lin Wang. “An oscillatory pipelining mechanism supporting previewing during visual exploration and reading”. In: *Trends in Cognitive Sciences* 25.12 (Dec. 1, 2021). Publisher: Elsevier, pp. 1033–1044. ISSN: 1364-6613, 1879-307X. DOI: 10.1016/j.tics.2021.08.008. URL: [https://www.cell.com/trends/cognitive-sciences/abstract/S1364-6613\(21\)00225-4](https://www.cell.com/trends/cognitive-sciences/abstract/S1364-6613(21)00225-4) (visited on 08/17/2023).
- [60] Ole Jensen, Eelke Spaak, and Johanna M. Zumer. “Human brain oscillations: From physiological mechanisms to analysis and cognition”. In: *Magnetoencephalography*. Vol. 9783642330452. Springer, July 1, 2014, pp. 359–403. ISBN: 978-3-642-33044-5. DOI: 10.1007/978-3-642-33045-2_17. URL: <http://www.scopus.com/inward/record.url?scp=84949175823&partnerID=8YFLogxK> (visited on 10/08/2023).
- [61] Karel Jezek et al. “Theta-paced flickering between place-cell maps in the hippocampus”. In: *Nature* 478.7368 (Sept. 28, 2011), pp. 246–249. ISSN: 1476-4687. DOI: 10.1038/nature10439.
- [62] Daniel Kahneman. *Attention and effort*. Prentice-Hall series in experimental psychology. Englewood Cliffs, N.J: Prentice-Hall, 1973. 246 pp. ISBN: 978-0-13-050518-7.
- [63] Daniel Kahneman, Anne Treisman, and Brian J. Gibbs. “The reviewing of object files: object-specific integration of information”. In: *Cognitive Psychology* 24.2 (Apr. 1992), pp. 175–219. ISSN: 0010-0285. DOI: 10.1016/0010-0285(92)90007-o.
- [64] Jan Kamiński, Aneta Brzezicka, Adam N. Mamelak, and Ueli Rutishauser. “Combined Phase-Rate Coding by Persistently Active Neurons as a Mechanism for Maintaining Multiple Items in Working Memory in Humans”. In: *Neuron* 106.2 (Apr. 22, 2020), 256–264.e3. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2020.01.032. URL: <https://www.sciencedirect.com/science/article/pii/S0896627320300581> (visited on 10/23/2023).
- [65] Takaaki Kaneko et al. “Cortical neural dynamics unveil the rhythm of natural visual behavior in marmosets”. In: *Communications Biology* 5.1 (Feb. 3, 2022). Number: 1 Publisher: Nature Publishing Group, pp. 1–13. ISSN: 2399-3642. DOI: 10.1038/s42003-022-03052-1. URL: <https://www.nature.com/articles/s42003-022-03052-1> (visited on 10/23/2023).
- [66] Sabine Kastner et al. “Increased Activity in Human Visual Cortex during Directed Attention in the Absence of Visual Stimulation”. In: *Neuron* 22.4 (Apr. 1, 1999). Publisher: Elsevier, pp. 751–761. ISSN: 0896-6273. DOI: 10.1016/S0896-6273(00)80734-5. URL: [https://www.cell.com/neuron/abstract/S0896-6273\(00\)80734-5](https://www.cell.com/neuron/abstract/S0896-6273(00)80734-5) (visited on 10/18/2023).
- [67] Christoph Kayser. “Phase resetting as a mechanism for supramodal attentional control”. In: *Neuron* 64.3 (Nov. 12, 2009), pp. 300–302. ISSN: 1097-4199. DOI: 10.1016/j.neuron.2009.10.022.
- [68] Timo van Kerkoerle et al. “Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex”. In: *Proceedings of the National Academy of Sciences* 111.40 (Oct. 7, 2014), pp. 14332–14341. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.1402773111. URL: <http://www.pnas.org/lookup/doi/10.1073/pnas.1402773111> (visited on 05/20/2020).
- [69] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. Jan. 29, 2017. DOI: 10.48550/arXiv.1412.6980. arXiv: 1412.6980 [cs]. URL: <http://arxiv.org/abs/1412.6980> (visited on 09/23/2023).
- [70] Lisa Kirchberger et al. “The essential role of recurrent processing for figure-ground perception in mice”. In: *Science Advances* 7.27 (June 1, 2021). Publisher: American Association for the Advancement of Science Section: Research Article, eabe1833. ISSN: 2375-2548. DOI: 10.1126/sciadv.abe1833. URL: <https://advances.sciencemag.org/content/7/27/eabe1833> (visited on 07/02/2021).
- [71] Wolfgang Klimesch, Paul Sauseng, and Simon Hanslmayr. “EEG alpha oscillations: The inhibition–timing hypothesis”. In: *Brain Research Reviews* 53.1 (Jan. 1, 2007), pp. 63–88. ISSN: 0165-0173. DOI: 10.1016/j.brainresrev.2006.06.003. URL: <https://www.sciencedirect.com/science/article/pii/S016501730600083X> (visited on 04/14/2021).

- [72] Nancy Kopell and Bard Ermentrout. “Chemical and electrical synapses perform complementary roles in the synchronization of interneuronal networks”. In: *Proceedings of the National Academy of Sciences of the United States of America* 101.43 (Oct. 26, 2004), pp. 15482–15487. ISSN: 0027-8424. DOI: 10.1073/pnas.0406343101.
- [73] Nikolaus Kriegeskorte. “Deep Neural Networks: A New Framework for Modeling Biological Vision and Brain Information Processing”. In: *Annual Review of Vision Science* 1 (Nov. 24, 2015), pp. 417–446. ISSN: 2374-4650. DOI: 10.1146/annurev-vision-082114-035447.
- [74] William B Kristan and Brian K Shaw. “Population coding and behavioral choice”. In: *Current Opinion in Neurobiology* 7.6 (Dec. 1, 1997), pp. 826–831. ISSN: 0959-4388. DOI: 10.1016/S0959-4388(97)80142-0. URL: <https://www.sciencedirect.com/science/article/pii/S0959438897801420> (visited on 09/20/2023).
- [75] Alex Krizhevsky. “Learning Multiple Layers of Features from Tiny Images”. In: 2009. URL: <https://www.semanticscholar.org/paper/Learning-Multiple-Layers-of-Features-from-Tiny-Krizhevsky/5d90f06bb70a0a3dced62413346235c02b1aa086> (visited on 10/23/2023).
- [76] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. “ImageNet Classification with Deep Convolutional Neural Networks”. In: *Advances in Neural Information Processing Systems*. Vol. 25. Curran Associates, Inc., 2012. URL: https://papers.nips.cc/paper_files/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html (visited on 03/22/2023).
- [77] Yann LeCun and Yoshua Bengio. “Convolutional networks for images, speech, and time-series”. In: *The handbook of brain theory and neural networks*. Ed. by M.A. Arbib. MIT Press, 1995.
- [78] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. “Gradient-based learning applied to document recognition”. In: *Proceedings of the IEEE* 86.11 (Nov. 1998). Conference Name: Proceedings of the IEEE, pp. 2278–2324. ISSN: 1558-2256. DOI: 10.1109/5.726791. URL: <https://ieeexplore.ieee.org/document/726791> (visited on 10/23/2023).
- [79] Kang Li et al. “Neurons in Primate Visual Cortex Alternate between Responses to Multiple Stimuli in Their Receptive Field”. In: *Frontiers in Computational Neuroscience* 10 (2016). ISSN: 1662-5188. URL: <https://www.frontiersin.org/articles/10.3389/fncom.2016.00141> (visited on 09/21/2023).
- [80] Stefanie Liebe et al. *Phase of firing does not reflect temporal order in sequence memory of humans and recurrent neural networks*. Pages: 2022.09.25.509370 Section: New Results. Sept. 27, 2022. DOI: 10.1101/2022.09.25.509370. URL: <https://www.biorxiv.org/content/10.1101/2022.09.25.509370v1> (visited on 10/07/2022).
- [81] John Lisman and György Buzsáki. “A Neural Coding Scheme Formed by the Combined Function of Gamma and Theta Oscillations”. In: *Schizophrenia Bulletin* 34.5 (Sept. 1, 2008), pp. 974–980. ISSN: 0586-7614. DOI: 10.1093/schbul/sbn060. URL: <https://doi.org/10.1093/schbul/sbn060> (visited on 09/21/2023).
- [82] John E. Lisman and M. A. Idiart. “Storage of 7 +/- 2 short-term memories in oscillatory subcycles”. In: *Science (New York, N.Y.)* 267.5203 (Mar. 10, 1995), pp. 1512–1515. ISSN: 0036-8075. DOI: 10.1126/science.7878473.
- [83] John E. Lisman and Ole Jensen. “The Theta-Gamma Neural Code”. In: *Neuron* 77.6 (Mar. 2013), pp. 1002–1016. ISSN: 08966273. DOI: 10.1016/j.neuron.2013.03.007. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0896627313002316> (visited on 08/17/2023).
- [84] Beth A. Lopour, Abtine Tavassoli, Itzhak Fried, and Dario L. Ringach. “Coding of Information in the Phase of Local Field Potentials within Human Medial Temporal Lobe”. In: *Neuron* 79.3 (Aug. 2013), pp. 594–606. ISSN: 08966273. DOI: 10.1016/j.neuron.2013.06.001. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0896627313004923> (visited on 09/22/2023).
- [85] Tiago Marques, Martin Schrimpf, and James J. DiCarlo. *Multi-scale hierarchical neural network models that bridge from single neurons in the primate primary visual cortex to object recognition behavior*. preprint. Neuroscience, Mar. 2, 2021. DOI: 10.1101/2021.03.01.433495. URL: <http://biorxiv.org/lookup/doi/10.1101/2021.03.01.433495> (visited on 11/22/2021).
- [86] Carrie J. McAdams and John H. Maunsell. “Attention to both space and feature modulates neuronal responses in macaque area V4”. In: *Journal of Neurophysiology* 83.3 (Mar. 2000), pp. 1751–1755. ISSN: 0022-3077. DOI: 10.1152/jn.2000.83.3.1751.

- [87] Douglas McLelland and Rufin VanRullen. “Theta-Gamma Coding Meets Communication-through-Coherence: Neuronal Oscillatory Multiplexing Theories Reconciled”. In: *PLOS Computational Biology* 12.10 (Oct. 14, 2016). Publisher: Public Library of Science, e1005162. ISSN: 1553-7358. DOI: 10.1371/journal.pcbi.1005162. URL: <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1005162> (visited on 09/21/2023).
- [88] Georgios Michalareas et al. “Alpha-beta and gamma rhythms subserve feedback and feedforward influences among human visual cortical areas”. In: *Neuron* 89.2 (Jan. 20, 2016), pp. 384–397. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2015.12.018. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4871751/> (visited on 11/24/2020).
- [89] Paul Miller. *An introductory course in computational neuroscience*. MIT Press, 2018.
- [90] A. David Milner and Melvyn A. Goodale. “Two visual systems re-viewed”. In: *Neuropsychologia. Consciousness and Perception: Insights and Hindsight* 46.3 (Jan. 1, 2008), pp. 774–785. ISSN: 0028-3932. DOI: 10.1016/j.neuropsychologia.2007.10.005. URL: <https://www.sciencedirect.com/science/article/pii/S0028393207003545> (visited on 09/27/2021).
- [91] Marcelo A. Montemurro et al. “Phase-of-Firing Coding of Natural Visual Stimuli in Primary Visual Cortex”. In: *Current Biology* 18.5 (Mar. 2008), pp. 375–380. ISSN: 09609822. DOI: 10.1016/j.cub.2008.02.023. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0960982208001681> (visited on 08/17/2023).
- [92] Jeffrey Moran and Robert Desimone. “Selective attention gates visual processing in the extrastriate cortex”. In: *Science (New York, N.Y.)* 229.4715 (Aug. 23, 1985), pp. 782–784. ISSN: 0036-8075. DOI: 10.1126/science.4023713.
- [93] John O’Keefe and Michael L. Recce. “Phase relationship between hippocampal place units and the EEG theta rhythm”. In: *Hippocampus* 3.3 (July 1993), pp. 317–330. ISSN: 1050-9631, 1098-1063. DOI: 10.1002/hipo.450030307. URL: <https://onlinelibrary.wiley.com/doi/10.1002/hipo.450030307> (visited on 04/10/2023).
- [94] Daria Osipova, Dora Hermes, and Ole Jensen. “Gamma Power Is Phase-Locked to Posterior Alpha Activity”. In: *PLOS ONE* 3.12 (Dec. 22, 2008). Publisher: Public Library of Science, e3990. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0003990. URL: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0003990> (visited on 10/23/2023).
- [95] Stefano Panzeri, Nicolas Brunel, Nikos K. Logothetis, and Christoph Kayser. “Sensory neural codes using multiplexed temporal scales”. In: *Trends in Neurosciences* 33.3 (Mar. 1, 2010), pp. 111–120. ISSN: 0166-2236. DOI: 10.1016/j.tins.2009.12.001. URL: <https://www.sciencedirect.com/science/article/pii/S0166223609002008> (visited on 07/06/2023).
- [96] Stefano Panzeri, Jakob H. Macke, Joachim Gross, and Christoph Kayser. “Neural population coding: combining insights from microscopic and mass signals”. In: *Trends in Cognitive Sciences* 19.3 (Mar. 1, 2015), pp. 162–172. ISSN: 1364-6613. DOI: 10.1016/j.tics.2015.01.002. URL: <https://www.sciencedirect.com/science/article/pii/S1364661315000030> (visited on 09/21/2023).
- [97] Dina V. Popovkina, John Palmer, Cathleen M. Moore, and Geoffrey M. Boynton. “Is there a serial bottleneck in visual object recognition?” In: *Journal of Vision* 21.3 (Mar. 11, 2021), p. 15. ISSN: 1534-7362. DOI: 10.1167/jov.21.3.15. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7961120/> (visited on 08/17/2023).
- [98] Ueli Rutishauser, Ian B. Ross, Adam N. Mamelak, and Erin M. Schuman. “Human memory strength is predicted by theta-frequency phase-locking of single neurons”. In: *Nature* 464.7290 (Apr. 2010). Number: 7290 Publisher: Nature Publishing Group, pp. 903–907. ISSN: 1476-4687. DOI: 10.1038/nature08860. URL: <https://www.nature.com/articles/nature08860> (visited on 10/23/2023).
- [99] Pankaj Sah and E. S. Louise Faber. “Channels underlying neuronal calcium-activated potassium currents”. In: *Progress in Neurobiology* 66.5 (Apr. 1, 2002), pp. 345–353. ISSN: 0301-0082. DOI: 10.1016/S0301-0082(02)00004-7. URL: <https://www.sciencedirect.com/science/article/pii/S0301008202000047> (visited on 09/25/2023).
- [100] Martin Schrimpf et al. “Brain-Score: Which Artificial Neural Network for Object Recognition is most Brain-Like?” In: *bioRxiv* (Jan. 2, 2020). Publisher: Cold Spring Harbor Laboratory Section: New Results, p. 407007. DOI: 10.1101/407007. URL: <https://www.biorxiv.org/content/10.1101/407007v2> (visited on 10/26/2020).

- [101] Eyal Seidemann and William T. Newsome. “Effect of Spatial Attention on the Responses of Area MT Neurons”. In: *Journal of Neurophysiology* 81.4 (Apr. 1, 1999), pp. 1783–1794. ISSN: 0022-3077, 1522-1598. DOI: 10.1152/jn.1999.81.4.1783. URL: <https://www.physiology.org/doi/10.1152/jn.1999.81.4.1783> (visited on 10/18/2023).
- [102] Joshua H. Siegle et al. “Survey of spiking in the mouse visual system reveals functional hierarchy”. In: *Nature* 592.7852 (Apr. 2021). Number: 7852 Publisher: Nature Publishing Group, pp. 86–92. ISSN: 1476-4687. DOI: 10.1038/s41586-020-03171-x. URL: <https://www.nature.com/articles/s41586-020-03171-x> (visited on 10/11/2023).
- [103] William E. Skaggs, Bruce L. McNaughton, Matthew A. Wilson, and Carol A. Barnes. “Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences”. In: *Hippocampus* 6.2 (1996). _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/%28SICI%291098-1063%281996%296%3A2%3C149%3A%3AAID-HIPO6%3E3.0.CO%3B2-K>, pp. 149–172. ISSN: 1098-1063. DOI: 10.1002/(SICI)1098-1063(1996)6:2<149::AID-HIPO6>3.0.CO;2-K. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/%28SICI%291098-1063%281996%296%3A2%3C149%3A%3AAID-HIPO6%3E3.0.CO%3B2-K> (visited on 09/25/2023).
- [104] Andrew T. Smith, Krish D. Singh, Adrian L. Williams, and Mark W. Greenlee. “Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex”. In: *Cerebral Cortex (New York, N.Y.: 1991)* 11.12 (Dec. 2001), pp. 1182–1190. ISSN: 1047-3211. DOI: 10.1093/cercor/11.12.1182.
- [105] Lynn K. A. Sörensen et al. “Leveraging Spiking Deep Neural Networks to Understand the Neural Mechanisms Underlying Selective Attention”. In: *Journal of Cognitive Neuroscience* 34.4 (Mar. 5, 2022), pp. 655–674. ISSN: 0898-929X. DOI: 10.1162/jocn_a_01819. URL: https://doi.org/10.1162/jocn_a_01819 (visited on 09/27/2023).
- [106] Eelke Spaak et al. “Layer-specific entrainment of gamma-band neural activity by the alpha rhythm in monkey visual cortex”. In: *Current Biology* 22 (2012). Publisher: Elsevier, pp. 2313–2318. DOI: 10.1016/j.cub.2012.10.020.
- [107] Nicholas A. Steinmetz et al. “Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings”. In: *Science* 372.6539 (Apr. 16, 2021). Publisher: American Association for the Advancement of Science, eabf4588. DOI: 10.1126/science.abf4588. URL: <https://www.science.org/doi/10.1126/science.abf4588> (visited on 10/11/2023).
- [108] Johan F. Storm. “Potassium currents in hippocampal pyramidal cells”. In: *Progress in Brain Research* 83 (1990), pp. 161–187. ISSN: 0079-6123. DOI: 10.1016/s0079-6123(08)61248-0.
- [109] Leslie G. Ungerleider and James V. Haxby. “‘What’ and ‘where’ in the human brain”. In: *Current Opinion in Neurobiology* 4.2 (Apr. 1994), pp. 157–165. ISSN: 0959-4388. DOI: 10.1016/0959-4388(94)90066-3.
- [110] Rufin VanRullen and Simon J Thorpe. “Is it a Bird? Is it a Plane? Ultra-Rapid Visual Categorisation of Natural and Artificial Objects”. In: *Perception* 30.6 (June 1, 2001). Publisher: SAGE Publications Ltd STM, pp. 655–668. ISSN: 0301-0066. DOI: 10.1068/p3029. URL: <https://doi.org/10.1068/p3029> (visited on 03/04/2021).
- [111] Kasper Vincken, Xavier Boix, and Gabriel Kreiman. “Incorporating intrinsic suppression in deep neural networks captures dynamics of adaptation in neurophysiology and perception”. In: *Science Advances* 6.42 (Oct. 1, 2020). Publisher: American Association for the Advancement of Science Section: Research Article, eabd4205. ISSN: 2375-2548. DOI: 10.1126/sciadv.abd4205. URL: <https://advances.sciencemag.org/content/6/42/eabd4205> (visited on 02/01/2021).
- [112] Pauli Virtanen et al. “SciPy 1.0: fundamental algorithms for scientific computing in Python”. In: *Nature Methods* 17.3 (Mar. 2020). Number: 3 Publisher: Nature Publishing Group, pp. 261–272. ISSN: 1548-7105. DOI: 10.1038/s41592-019-0686-2. URL: <https://www.nature.com/articles/s41592-019-0686-2> (visited on 09/26/2023).
- [113] Tim P. Vogels and L. F. Abbott. “Gating multiple signals through detailed balance of excitation and inhibition in spiking networks”. In: *Nature Neuroscience* 12.4 (Apr. 2009). Number: 4 Publisher: Nature Publishing Group, pp. 483–491. ISSN: 1546-1726. DOI: 10.1038/nn.2276. URL: <https://www.nature.com/articles/nn.2276> (visited on 09/20/2022).
- [114] Athanasios Voulodimos, Nikolaos Doulamis, Anastasios Doulamis, and Eftychios Protopapadakis. “Deep Learning for Computer Vision: A Brief Review”. In: *Computational Intelligence and Neuroscience* 2018 (Feb. 1, 2018). Publisher: Hindawi, e7068349. ISSN: 1687-5265. DOI: 10.1155/2018/7068349. URL: <https://www.hindawi.com/journals/cin/2018/7068349/> (visited on 03/22/2023).

- [115] Bradley Voytek et al. “Shifts in Gamma Phase–Amplitude Coupling Frequency from Theta to Alpha Over Posterior Cortex During Visual Tasks”. In: *Frontiers in Human Neuroscience* 4 (2010). ISSN: 1662-5161. URL: <https://www.frontiersin.org/articles/10.3389/fnhum.2010.00191> (visited on 10/11/2023).
- [116] Saskia E. J. de Vries et al. “A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex”. In: *Nature Neuroscience* 23.1 (Jan. 2020). Number: 1 Publisher: Nature Publishing Group, pp. 138–151. ISSN: 1546-1726. DOI: 10.1038/s41593-019-0550-9. URL: <https://www.nature.com/articles/s41593-019-0550-9> (visited on 10/11/2023).
- [117] Alex L. White, John Palmer, Geoffrey M. Boynton, and Jason D. Yeatman. “Parallel spatial channels converge at a bottleneck in anterior word-selective cortex”. In: *Proceedings of the National Academy of Sciences* 116.20 (May 14, 2019). Publisher: Proceedings of the National Academy of Sciences, pp. 10087–10096. DOI: 10.1073/pnas.1822137116. URL: <https://www.pnas.org/doi/10.1073/pnas.1822137116> (visited on 10/02/2023).
- [118] Alex L. White et al. “Evidence for unlimited capacity processing of simple features in visual cortex”. In: *Journal of Vision* 17.6 (June 1, 2017), p. 19. ISSN: 1534-7362. DOI: 10.1167/17.6.19.
- [119] Robert K. Wong, David A. Prince, and Allan I. Basbaum. “Intradendritic recordings from hippocampal neurons”. In: *Proceedings of the National Academy of Sciences of the United States of America* 76.2 (Feb. 1979), pp. 986–990. ISSN: 0027-8424. DOI: 10.1073/pnas.76.2.986.
- [120] Ren Wu et al. *Deep Image: Scaling up Image Recognition*. July 5, 2015. DOI: 10.48550/arXiv.1501.02876. arXiv: 1501.02876[cs]. URL: <http://arxiv.org/abs/1501.02876> (visited on 10/23/2023).
- [121] Daniel L. K. Yamins and James J. DiCarlo. “Using goal-driven deep learning models to understand sensory cortex”. In: *Nature Neuroscience* 19.3 (Mar. 2016). Number: 3 Publisher: Nature Publishing Group, pp. 356–365. ISSN: 1546-1726. DOI: 10.1038/nn.4244. URL: <https://www.nature.com/articles/nn.4244> (visited on 03/22/2023).
- [122] Daniel L. K. Yamins et al. “Performance-optimized hierarchical models predict neural responses in higher visual cortex”. In: *Proceedings of the National Academy of Sciences* 111.23 (June 10, 2014). Publisher: National Academy of Sciences Section: Biological Sciences, pp. 8619–8624. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.1403112111. URL: <https://www.pnas.org/content/111/23/8619> (visited on 01/20/2022).
- [123] Davide Zambrano, Roeland Nusselder, H. Steven Scholte, and Sander M. Bohtë. “Sparse Computation in Adaptive Spiking Neural Networks”. In: *Frontiers in Neuroscience* 12 (2019). Publisher: Frontiers. ISSN: 1662-453X. DOI: 10.3389/fnins.2018.00987. URL: <https://www.frontiersin.org/articles/10.3389/fnins.2018.00987/full> (visited on 02/08/2021).
- [124] Honghui Zhang, Andrew J. Watrous, Ansh Patel, and Joshua Jacobs. “Theta and Alpha Oscillations Are Traveling Waves in the Human Neocortex”. In: *Neuron* 98.6 (June 27, 2018), 1269–1281.e4. ISSN: 1097-4199. DOI: 10.1016/j.neuron.2018.05.019.
- [125] Alexander Zhigalov and Ole Jensen. “Perceptual echoes as travelling waves may arise from two discrete neuronal sources”. In: *NeuroImage* 272 (May 15, 2023), p. 120047. ISSN: 1095-9572. DOI: 10.1016/j.neuroimage.2023.120047.