The cortical microcircuitry of predictions and context – a multi-scale perspective

Lars Muckli[1,2,*], Lucy S. Petro[1,2,*], Clement Abbatecola[1,2], Ahsan Adeel[3], Johanna Bergmann[4], Nicolas Deperrois[5], Alain Destexhe[6], Nikolaus Kriegeskorte[7], Christiaan N. Levelt[8,9], Wolfgang Maass[10], A. Tyler Morgan[11], Paolo Papale[12], Cyriel M. A. Pennartz[13,14], Benjamin Peters[1], Mihai A. Petrovici[5], William. A. Phillips[15], Pieter R. Roelfsema[16], Robert N.S. Sachdev[17], Koen Seignette[8], Matthew W. Self[12], Fraser W. Smith[18], Johan F. Storm[19], Michele Svanera[1,2], Wim Vanduffel[20], Walter Senn[5]*, Matthew E. Larkum[17,21]*

[1]Centre for Cognitive Neuroimaging, School of Psychology and Neuroscience, College of Medical, Veterinary and Life Sciences, University of Glasgow, Glasgow, United Kingdom, G12 8QB.
[2]Imaging Centre for Excellence, Queen Elizabeth University Hospital, University of Glasgow, G51 4LB.
[3]CMI Lab, Department of Computing Science and Mathematics, University of Stirling, FK9 4LA, Stirling, UK, Oxford Computational Neuroscience Lab, Nuffield Department of Surgical Sciences, University of Oxford, Oxford, UK.
[4]Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany.
[5]Department of Physiology, University of Bern, Switzerland
[6]Paris-Saclay University, CNRS, Saclay, France
[7]Zuckerman Mind Brain Behavior Institute, Columbia University, 3227 Broadway, New York, NY, 10027, USA
[8]Netherlands Institute for Neuroscience, Molecular Visual Plasticity Group, Royal Netherlands Academy of Arts and Sciences, Meibergdreef 47, 1105 BA Amsterdam, the Netherlands.
[9]Department of Molecular and Cellular Neurobiology, Center for Neurogenomics and Cognitive Research, VU University Amsterdam, de Boelelaan 1085, 1081 HV Amsterdam, the Netherlands.
[10]Institute of Theoretical Computer Science, Graz University of Technology, Graz, Austria.
[11]Section on Functional Imaging Methods, National Institute of Mental Health, Bethesda, MD 20817, USA.
[12]Department of Vision & Cognition, Netherlands Institute for Neuroscience (KNAW), 1105 BA Amsterdam, the Netherlands.
[13]Cognitive and Systems Neuroscience Group, Swammerdam Institute for Life Sciences, University of Amsterdam, Science Park 904, 1098XH Amsterdam, The Netherlands.
[14]Amsterdam Brain and Cognition, University of Amsterdam, Science Park 904, 1098XH Amsterdam, The Netherlands.
[15]Department of Psychology, University of Stirling, Stirling, FK9 4LA, UK.
[16]Department of Vision & Cognition, Netherlands Institute for Neuroscience (KNAW), 1105 BA Amsterdam, the Netherlands; Department of Integrative Neurophysiology, VU University, De Boelelaan 1085, 1081 HV Amsterdam, the Netherlands; Department of Neurosurgery, Academic Medical Centre, Postbus 22660, 1100 DD Amsterdam, the Netherlands; Laboratory of Visual Brain Therapy, Sorbonne Université, INSERM, CNRS, Institut de la Vision, 17 rue Moreau, 75012 Paris, France.
[17]Institute of Biology, Humboldt Universität zu Berlin, Virchowweg 6, 10117 Berlin, Germany.
[18]School of Psychology, University of East Anglia, Norwich Research Park, Norwich NR4 7TJ, UK.
[19]Brain Signalling Group, Department of Physiology, Institute of Basic Medicine, University of Oslo, Oslo, Norway.
[20]Laboratory for Neuro- and Psychophysiology, Department of Neurosciences, KU Leuven Medical School, Leuven, 3000, Belgium. Leuven Brain Institute, KU Leuven, Leuven, 3000, Belgium. Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, MA, 02129, USA. Department of Radiology, Harvard Medical School, Boston, MA, 02144, USA.
[21]Neurocure Center for Excellence, Charité Universitätsmedizin Berlin, Germany
*Equal contribution

Abstract
Conscious cognition depends on the ability of the neocortex to generate internal models of the outside world. During wakefulness, the neocortex maintains and updates knowledge of the world and uses this knowledge through top-down projections to make predictions, test hypotheses, and/or contextualise input from the senses. How are these information streams combined in cortical microcircuits? Is their computational function to test internal models on the basis of their predictions or to contextualise sensory signals, or both? In addition to their somatic integration zones, many pyramidal neurons have a site of top-down and other contextual information integration near the top of the apical dendrite's trunk. This architecture enables top-down contextualisation of bottom-up information, amplifying or attenuating sensory responses depending on prior knowledge and current context. However, current deep neural network models of sensory processing lack such a mechanism, and cognitive theories do not reach the explanatory level of intracellular two-compartment integration. In this interactive 'live' paper, we envision how a continued synthesis of multi-scale, multi-species experimental data and theoretical and data-driven models will drive further insights into the biophysics, microcircuitry and dynamics of context-sensitive two-compartment neurons, and their role in predictive cognition.

Introduction - mapping contextual information onto microcircuits that incorporate dendritic computations

The integration of external evidence with internal models is crucial for intelligent behaviour. Our ability to think and move depends on quickly understanding information from the world around us, using what we already know, and adjusting our mental models. All this is achieved apparently effortlessly by the brain, while continuously taking many things into account, such as prior knowledge, current demands and future goals. Converging evidence suggests that cortical microcircuits and intraneuronal mechanisms combine top-down contextual information with current sensory information about the world (Larkum, 2013, Schuman et al., 2021; Pennartz et al. 2019), and that this integration occurs already at the level of single neurons. Many studies suggest that this operation is performed by specialised functions within the main cells of the cerebral cortex that can associate top-down and bottom-up information streams arriving at apical and basal sites respectively (Larkum, 2022; Larkum et al., 2022). Understanding the details of this operation has required a paradigm shift to view cortical neurons as multi-compartment processors performing nonlinear computations. Input to pyramidal neuron dendrites in cortical layer 1 is crucial for cognitive function, as this is an important layer for corticocortical feedback, and receives the strongest neuromodulatory input (Schuman et al., 2021, see also Ledderose et al, 2023).

We propose that this paradigm shift, applicable to both biological and artificial neurons, has the potential to transform the field of machine learning (Beniaguev et al., 2021; Haider et al., 2021; Archarya et al., 2022; Iyer et al., 2022, Max et al., 2023; Wybo et al., 2023), revolutionise neuromorphic computing (Adeel, et al., 2022; Khacef et al., 2022; Stöckel and Eliasmith, 2022), and stimulate novel mechanistic explanations of the cellular and microcircuit foundations of computation and cognition (Kreutzer et al., 2022; Phillips, 2023, Figure 1). However, investigating the integration of top-down and bottom-up information in individual neurons and cortical microcircuits is not trivial (Pardi et al., 2023), requiring theoretical

perspectives combined with multi-species approaches, using paradigms that can disentangle information streams. Here we review recent studies using a behavioural paradigm (visual occlusion) designed to isolate the influence of top-down information. This approach deliberately obscures a portion of an image (the "occluded" region). This region of the visual field in the primary visual cortex of humans, monkeys and mice does not receive bottom-up receptive field stimulation, and recordings therefore measure functional and structural principles of top-down processing. We then describe how the spatial separation of top-down and bottom-up inputs within single pyramidal neurons and the thalamocortical control of their integration might have causal implications for perception, learning, attention, working memory, and consciousness. Lastly, we evaluate which biological abstraction and computational motifs could inspire novel search spaces for artificial intelligence architectures. In summary, we offer a new perspective on the function of the cerebral cortex that emphasises the cooperative nature of its components and architecture (Phillips, 2023).

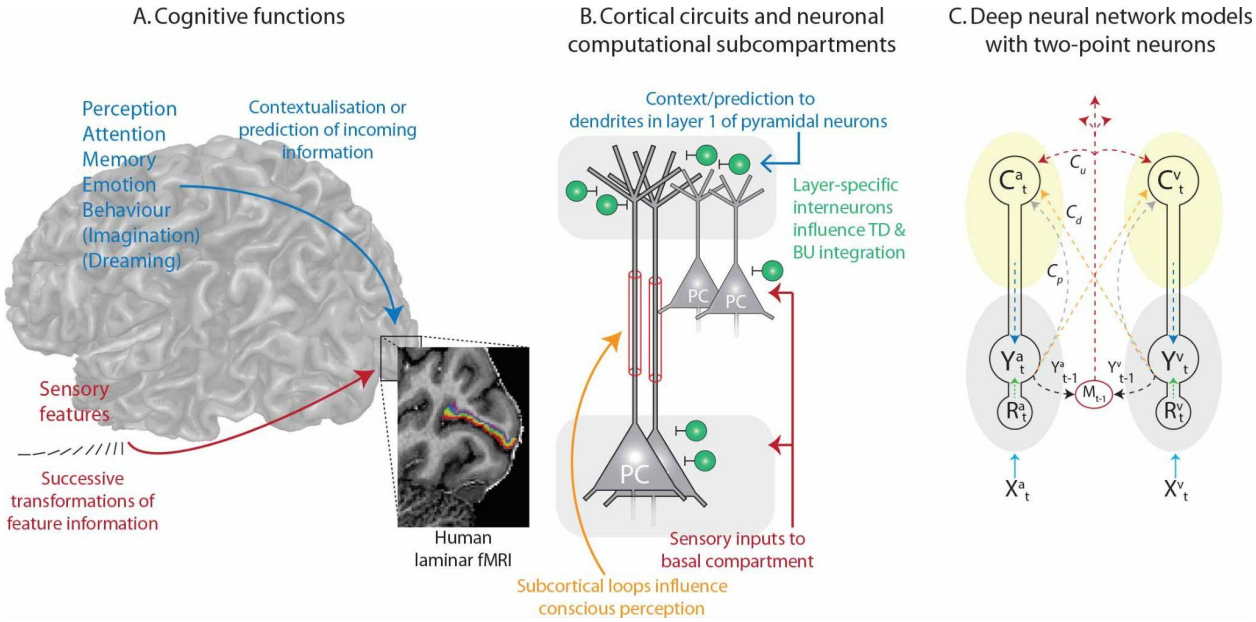

Figure 1: Exploring multiscale empirical frameworks and machine learning to elucidate mechanisms of cognition based on context-dependent processing as a computational goal. We reframe sensory processing from a bottom-up (BU), feature tuning model to a cooperative process in which top-down (TD) connections regulate the bottom-up flow of information via contextual inputs to dendrites of pyramidal neurons in the neocortex. (A) High-resolution human brain imaging (and sophisticated imaging and electrophysiological techniques in animals) can probe layer-specific function during cognitive states revealing the combination of top-down with bottom-up sensory information. (B) The assimilation of these distinct sources of information can be understood with neural models that include dendritic function. Interneuron subtypes and their functional roles are described in more detail in the text. (C) The characterisation of dendritic mechanisms, in turn, inspire advancements in machine learning (see Beniaguev et al., 2021; Pagalos et al., 2023a). Example DNN with a context-sensitive (CS) two-point neuron (TPN) inspired architecture processes real-world audio-visual data more efficiently (modified from Adeel et al., 2023a, and see later section 'Context-dependent neural computation in next-generation AI and neuromorphic computing paradigms').

[An extended description of the DNN is as as follows: Individual CS-TPNs receive audio receptive field (R_a^t) and video receptive field (R_v^t) and cooperate moment-by-moment via proximal, local, and universal contextual fields (C_P, C_D, and C_U) to conditionally segregate the coherent and conflicting FF streams and then recombine only the coherent multistream to extract synergistic FF components (brief memory): M. The extracted brief memory components are

broadcasted and received by neighbouring processors along with the current C_P and C_D. The integrated context (C) and integrated R are used by the asynchronous modulatory transfer function (AMTF) to split the signal into coherent and incoherent signals. The magic for selecting the relevant FF information for the ongoing transmission lies in the new kind of AMTF that overrules the typical dominance of receptive field (R) (outside world) being the driving force behind neural output, and awards more authority to the contextual information coming from the neighbouring neurons (inside world) (Adeel et al., 2020, 2022, 2023a, 2023b, 2023c). Here, C_P defines the modulatory signal coming from the neighbouring cell of the same network or the cell's output at time t-1, C_D defines the modulatory sensory signal coming from some other parts of the brain (in principle from anywhere in space-time), and C_U defines the outside environment and anticipated behaviour or widely distributed general context (based on past learning and reasoning). For simplicity, C_U is linked to brief memory formation and retrieval which is universal in that its modulatory effect is broadcasted to all sensory modalities. In general, C_U could explicitly be extended to the sources of inputs to include general information about the target domain acquired from prior experiences, emotional states, intentions, cognitive load, and semantic knowledge. This architecture with CS-TPNs permits a reduction of irrelevant feedforward information flow, disambiguating feedforward signals, making learning easier, and reducing energy demand during training and inference compared to point neurons. See demo: cmilab.org/research. The spiking simulations demonstrate an efficient solution to the credit assignment problem through apical dendrites, using CS-TPNs-driven burst-dependent synaptic plasticity (BDSP). It is shown how CS-TPNs tend to remain largely silent when information is less relevant but become active (bursting) when information is relevant. Additionally, it is notable that TPNs fire more frequently than CS-TPNs. For a simple XOR problem, CS-TPNs learn faster compared to context-insensitive TPNs-driven BDSP (Payeur, et al., 2021)].

How does the cortex do it? A summary of current knowledge about how the cortex represents top-down information

A prevailing view is that the architecture of the cerebral cortex is structured to support the integration of sensory inputs with current demands, expectations and prior knowledge, in distinct processing streams flowing in bottom up and top down directions, embedded in laminar microcircuits. Within this neocortical architecture, stored information modulates the processing of incoming information through top-down circuits. Cognition, therefore, works in a predictive or generative way such that the cortex internally generates a representation of the outside world that is fed back to sensory areas and compared to external inputs, for example, to complete inference and improve internal models by learning (Dayan et al., 1995; Rao and Ballard, 1999; Lee and Mumford, 2003; Friston, 2005; Clark, 2015). Similarly, evidence suggests that top-down circuitry with lateral corticocortical connections in both superficial and deep cortical layers (Markov and Kennedy, 2013), conveys crossmodal sensory information supporting the construction of coherent multimodal world models (Garner and Keller 2022; Pearson et al. 2021; Pennartz et al. 2023). However, the content of internal representations and the functional roles of top-down processing remain often theoretical or lack mechanistic detail, in part due to the empirical challenge of disentangling feedback inputs from strong driving feedforward processing (Keller et al., 2020), especially at intracellular levels. It is therefore essential to establish parsimonious approaches to isolating top-down influences.

Visual occlusion paradigms for isolating the influence top-down information

Visual occlusion is a paradigmatic example in which top-down processing provides input layers with additional contextual information beyond what is available from bottom-up thalamic input (e.g. Smith and Muckli, 2010). By partially occluding visual scenes, cortical retinotopic regions or neurons that receive only lateral and cortical top-down input can be isolated. Here, the remaining signals relate to contextual information or to internally-generated model-based

predictions of the missing stimulus features. Applying this occlusion paradigm in a multispecies approach establishes the beginnings of an empirical framework parameterising contextual top-down processes in vision, and the functional relationships between internal representations and sensory signals.

The content of top-down information at occluded locations

Using this partial visual occlusion paradigm has allowed the field to make more precise statements about the content and effects of feedback information. When humans are presented with partially occluded images during fMRI, the surrounding scene content creates a mental model prediction that can be read out from the retinotopic V1 regions processing only the occluded part of the image. The fMRI signal in this situation is consistent with the interpretation that feedback information from higher cortical areas (rather than local recurrent or lateral connections, Papale et al., 2023) is responsible for activity in V1. However, beyond contextual modulation, this top-down input might code for higher-level predictions about the incoming stimulus content, which is missing in the case of occlusion (Morgan et. al., 2019; Pennartz et al., 2019). In line with this interpretation, top-down scene-specific information in the non-stimulated retinotopic region of V1 correlates with orientation information found in internal model predictions of scenes, sampled in humans as behavioural line drawings of missing scene features. Therefore, simplified representations of visual scenes, at least in this case, provide a good description of an internal model of the scene and conceivably the contents of feedback information to V1 (Morgan et al., 2019, Figure 2A).

Surprise if what should be there is occluded

One possibility is that, during amodal completion, where the brain uses prior knowledge to interpret what is hidden even though the occluded content is not consciously perceived (Peters and Kriegeskorte, 2021), an internal generative model is fed back to V1 that matches the expected scene features. The absence of expected features drives a negative prediction error, in L2/3 neurons encoding errors (Seignette et al., 2023a). The negative prediction error might also generate an fMRI signal in more superficial layers for amodal (Muckli et al. 2015) and modal completion (Bergmann et al. 2019), due to the absence of expected sensory input. We propose that what becomes conscious in modal and amodal completion is the figure that is in front of the occluded background. The activity in layer 2/3 in relation to amodal and modal completion is detectable in human fMRI in superficial cortical depth layers and in mouse data showing stronger illusory modulation in the superficial layers of mouse V1 (Pak et al., 2020). The conscious readout of the foreground figure might activate L5 pyramidal neurons, (Aru et al. 2020a,b) but might only be detectable in human fMRI downstream (fMRI signal measures more input and less output signal, Logothetis, 2008). Modal and amodal image completion are distinguished from imagery by the ability to simply describe missing low-level features versus the subjective sense of imagined content. This raises the question as to whether sensory completion versus imagery reflects differences in the contents of feedback, or some other process. We hypothesize that imagery involves the activation of thalamo-cortical loops involving deep thalamic-projecting pyramidal neurons that creates the sensation of imagined content whereas sensory completion (or rather also *sensory absence* during modal or amodal conditions) also activates pyramidal neurons in superficial cortical layers related to error

signalling of missing expected features (Seignette et al., 2023a). In this scenario, feedback would represent a back-projected template that need not precisely fill in retinotopic features but can generalise over a few degrees of visual angle (Petro et al., 2023), and spatial frequencies (Revina et al., 2018). This suggestion has been further supported by a multi-layer computational model of predictive coding that combines specific object representation with view-invariant recognition and is capable of pattern completion in the occlusion paradigm (Brucklacher et al., 2023).

Layer fMRI

Laminar imaging localises this top-down contextual information to superficial layers of V1 (Muckli et al., 2015, Figure 2A). This is enabled by high-resolution imaging with 7T fMRI, that has sufficient spatial resolution to approximate energy consumption in cortical layers and reflects locations of high synaptic activity (Uludag and Blinder, 2018). This fMRI data might be detecting top-down inputs to the apical dendrites in L1 of L2/3 and/or L5 pyramidal neurons encoding perceptual internal models in humans, bridging human laminar fMRI with animal neuronal models (Larkum et al., 2018). Top-down signals in the absence of bottom-up inputs can also be investigated in humans with imagery, attention, or cross-modal stimulation tasks (Lawrence et al., 2019; Gao et al., 2020; Marquardt et al., 2020). fMRI studies reveal activity or information in upper and lower cortical depth compartments containing feedback information, in line with anatomical projection profiles. Imagining visual content can be detected in elevated information in the BOLD signal restricted to the deep layers of human V1, whereas seemingly 'real' illusory content is mainly located in superficial layers (but see Kok et al., 2016). This suggests that early visual microcircuits receive distinct types of feedback information that reflect the mode of conscious perception (Bergmann et al., 2019). As yet, it is not clear how these different classes of internally generated signals are re-integrated into cortical columns in primary sensory regions and further evidence is required to interpret the specific influence of feedback processing (Vezoli et al., 2021).

Monkeys

Accessing the neuronal physiology implied by these human visual microcircuits is necessary to understand the cellular mechanisms of these top-down inputs. We employed the partial visual occlusion paradigm in a cross-species approach, presenting monkeys and mice with the identical partially occluded stimuli shown to humans (Papale et al., 2023 and Seignette et al., 2023a, Figure 2B,C). These studies demonstrate again that, counter to classical theories of vision, neurons deprived of bottom-up sensory stimulation nevertheless encode stimulus-selective information. This was revealed by decoding image identity based on responses to the occluded portion of the stimuli. The dynamics of this contextual feedback information, studied using electrophysiology in monkeys, shows that V1 neurons in the occluded region (that don't receive feed-forward information) process scene-specific contextual information less than 20 ms later than neurons whose receptive fields are presented with the scene stimulus (Papale at al., 2023, Figure 2B). A similar delay in responses to contextual information is observed in mouse V1 neurons responding to occluded images (Figure 2C). Such late enhancement of V1 activity is necessary for figure-ground segregation, both in monkeys (e.g. Poort et al., 2012) and mice (Kirchenberger et al., 2021), and to enhance object regions in natural images (Papale et al.,

2021). These representations of non-occluded scene information in the cortical region corresponding to the visually occluded region were correlated between humans and monkeys, suggesting cross-species similarity in the content of perceptual internal models supplied by higher visual areas to V1 (Figure 2B).

Mice

In mice, two-photon imaging of neuronal calcium signals in L2/3 neurons in V1 shows that distinct populations of L2/3 neurons code for either sensory or feedback information, and contextual responses to occluded scenes are stronger in trained animals, suggesting that these responses can be explained in a predictive processing framework (Seignette et al., 2023a, Figure 2C, *cf. Pennartz et al. 2019*). Subsets of inhibitory interneurons might help to regulate the weighting given to sensory and feedback streams. Somatostatin-expressing (SST) deep-layer Martinotti neurons very effectively control apical dendritic activity and are activated by deep layer pyramidal neurons but are disinhibited in awake versus anaesthetised animals (Murayama et al., 2009a; Murayama et al., 2009b). Similarly, SST interneurons in L2/3 that are active during normal sensory input are disinhibited by feedback input in occluded image paradigms (Keller et al., 2020). This disinhibition is mediated by Vasoactive Intestinal Peptide (VIP) expressing interneurons that are particularly active when the cortex processes inputs that strongly activate top-down inputs. This is revealed convincingly by making use of stimuli lacking bottom up inputs, in situations of occlusion or uncertainty, for example (Kirchberger et al., 2023). The strong activation of VIP cells in otherwise inactive regions of cortex silences nearby SST cells, releasing the apical dendrites from inhibition, thus rendering pyramidal cells more sensitive to top-down inputs. Optogenetically silencing VIP cells reduces activity in pyramidal cells, particularly in situations where the pyramidal neurons are only being activated by contextual inputs (Kirchberger et al., 2023). In addition to being strongly activated when feedforward-input is missing, VIP-cells are also activated by contextual feedback information during normal visual processing (Zhang et al., 2014, Kirchberger et al., 2021). VIP activity is enhanced in regions of V1 receiving feedback (e.g. on a figure compared to background regions) again causing a release of the apical dendrites in V1 from inhibition by inhibiting somatostatin (SST) cells. These findings obtained using occluded stimuli match well with results obtained using a visuomotor mismatch paradigm to study predictive processing, which also revealed involvement of this canonical disinhibitory circuit (Attinger et al., 2017). Intriguingly, visuomotor mismatch also strongly activates another subset of interneurons, the elusive axon-targeting chandelier cells (Seignette et al 2023b). How their activity modulates the activity of L2/3 pyramidal neurons and how this ties into predictive processing remains unknown however.

Models

To further test the intuition that internal models are reconstructing the image portion under the occluder, Svanera et al., (2021) presented two neural network models with the same partially occluded stimuli shown to humans, monkeys and mice. After being trained to fill in the missing scene quadrant, a self-supervised deep convolutional neural network (CNN) with an encoder/decoder architecture outperformed a classical object-recognition supervised network (VGG16) in terms of similarity to brain data. The branch of the network more similar to the early visual cortex was not the portion compressing retinotopic features into higher level

representations (i.e., the forward encoder pathway), but instead, the portion reconstructing retinotopic features from higher level representations (Svanera et al., 2021, Figure 2D). The network supports the interpretation that in the human, monkey and mouse data there is information in V1 not related to sensory features but to contextual information or generative internal models of the expected external information. This modelling work is complemented by multi-layer predictive coding models styled after the feed-forward/feedback architecture of the cortex and trained by unsupervised, Hebbian learning (Dora et al. 2021; Brucklacher et al. 2023).

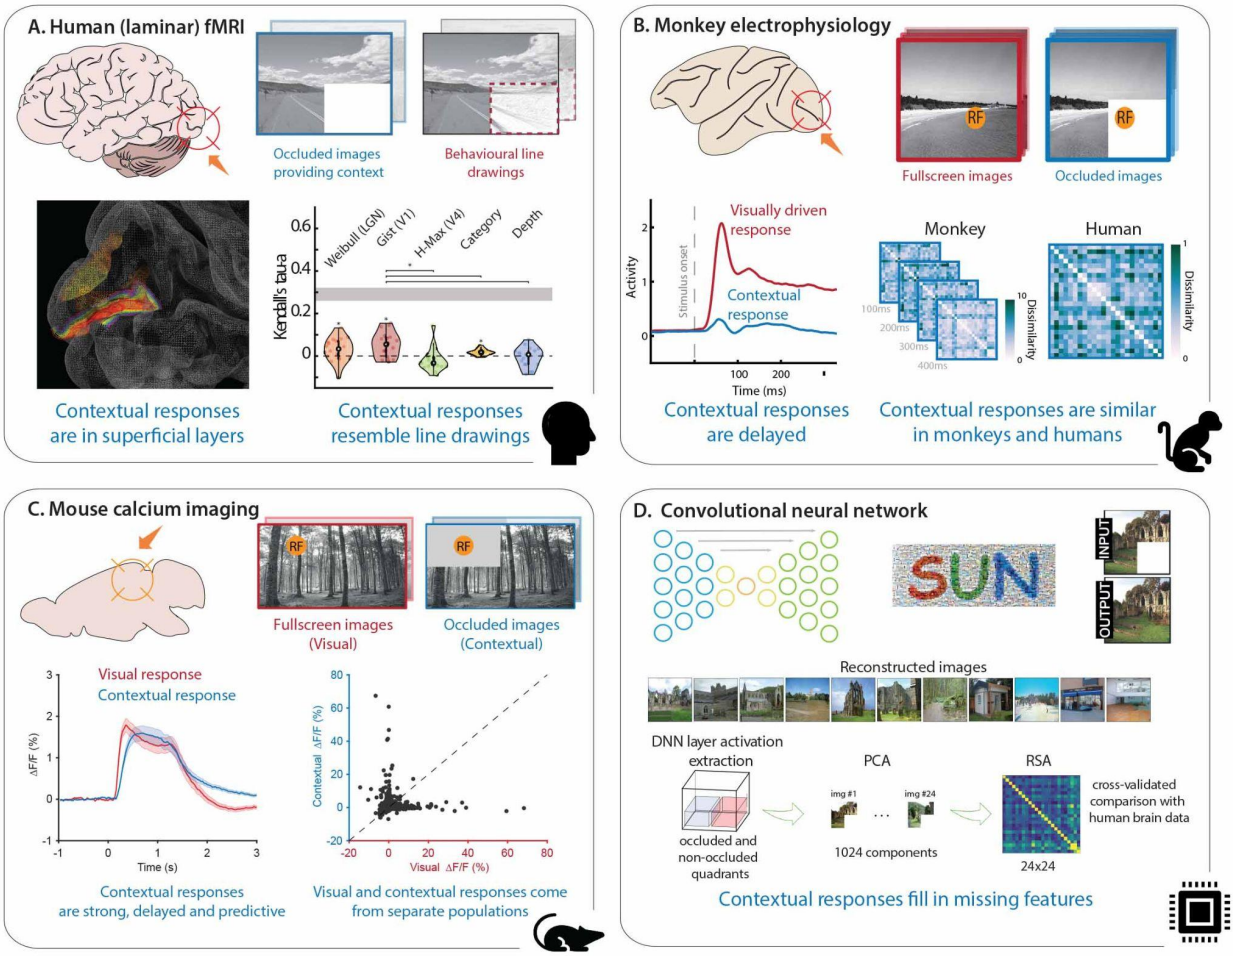

Figure 2: Multispecies partial visual occlusion tasks to parameterise contextual feedback responses in non-stimulated primary visual cortex. In human, monkey and mouse V1, scene identity can be decoded using multivariate pattern analyses of data taken from a region processing only the occluded image patch. Contextual responses are in response to this occluded image patch. A. Laminar fMRI data localises these contextual responses to occluded stimuli in the superficial depths of V1 (Muckli et al., 2015). Contextual responses in the non-stimulated retinotopic region of V1 correlate with orientation information found in internal model predictions of scenes, sampled in humans as line drawings of the masked scene content (Morgan et al., 2019. See also Revina et., 2018, Ortiz-Tudela et al 2023; Petro et al., 2023; Lazarova et al., 2023). B. Spiking activity of V1 neurons in a monkey viewing either natural scenes or scenes where the information in the RF was occluded reveals that neurons responded rapidly and selectively to occluded scenes, but 20ms later than to non-occluded scenes. Monkey V1 spiking responses to occluded stimuli correlate strongly with the representations of the same scenes in humans measured with fMRI (Papale et al., 2023). C. Two photon calcium imaging in V1 of awake mice. Animals trained to

detect fullscreen natural scenes were subsequently presented with fullscreen or occluded natural scenes. L2/3 neurons in V1 that had their RF on the occluder responded to either fullscreen or occluded stimuli. This suggests the existence of separate neuronal populations responding to feedforward or contextual information (Seignette et al., 2023a). D. Exploiting a generative encoder-decoder architecture to model human fMRI responses measuring cortical feedback processing (Svanera et al., 2021). The autoencoder model, and mouse, monkey and human occlusion data are available on EBRAINS to download, or as Jupyter notebooks that can be used to reproduce the experimental data (link).

Multiscale and multimethods

The human, monkey, mouse and model findings presented above illustrate how we came together as a collaborative group with a common interest in the functions and microcircuits of cortical feedback processes. Such a systematic accumulation of evidence towards multilevel spatiotemporal characteristics of top-down circuits in sensory processing sits within broader schemes for the future of integrated neuroscience (Amunts et al., 2022, Paquola et al., 2022). There, ambitious interdisciplinary and large-scale collaborations are innovating towards the computational principles and circuits of human intelligence, brain models for artificial intelligence, transformational neuromorphic systems, and strategic applications in brain medicine (e.g. Human Brain Project, Center for Brains, Minds and Machines). In these endeavours, understanding corticocortical and thalamocortical loops is necessary for understanding principles of intelligence, depending crucially on the mechanisms of cortical feedback processing (Larkum, 2013). But functional descriptions of cognition need to be anchored to the level of neuroanatomical features that capture sufficient explanatory detail upon which to ground our theories. At first glance, these remarkably similar cross-species stimulus-specific responses in neurons with no informative sensory stimulation are not fully explained by current functional descriptions of laminar circuits, and may even involve streaming of subcortical information to thalamus and cortex (Suzuki et al., 2023). There is substantial evidence that cortical feedback inputs target the apical dendrites of pyramidal neurons (Pardi et al., 2022; Schuman et al., 2021), however there is also a significant projection to basal dendrites in deep layers (Manita et al., 2015). At present, it is not understood what function this second projection serves, although it has been suggested that it might allow certain top-down projections to activate the coincidence detection mechanism in layer 5 pyramidal neurons by activating both the apical and basal compartments simultaneously (Manita et al., 2017). In any case, it is clear that the apical dendrites have specialised computational abilities afforded by a range of linear and nonlinear mechanisms (Larkum et al., 2022). It is now essential to ask how long-range feedback signalling in the cortex dovetails with processing at the cellular level and under what circumstances. Further, what are useful biophysical abstractions of this microscale anatomy and function for constraining macroscale measures, computational models and artificial networks?

Subcellular basis of cognitive function

Despite the abundance of evidence (reviewed above) that pyramidal neurons can perform complex operations on feedforward and feedback inputs embedded in a rich local network of interacting interneurons, pyramidal neurons are still most frequently modelled or conceptualised as point neurons, that is, without dendrites (Larkum, 2022). In vivo and in vitro studies suggest that, far from operating as point neurons, pyramidal neuron activity is strongly

determined by dendritic activation via L1 feedback inputs using their repertoire of Ca^{2+} (calcium), Na^{2+} (sodium) and NMDA (N-methyl-D-aspartate) spikes (London and Häusser, 2005; Larkum et al., 2009; Larkum, 2022). Therefore, at least a two-compartment abstraction might be necessary to capture input/output properties of neocortical pyramidal neurons with sufficient biological realism (Major et al., 2013). Recent evidence that the apical dendritic compartment can become isolated from their cell bodies in L5 pyramidal neurons suggests that a 3rd compartment (a "coupling" compartment) is necessary to capture the influence of higher-order thalamus, possibly related to attention, that is absent during anaesthesia (Larkum et al., 2001; Suzuki and Larkum, 2020; Aru et al., 2020a). At minimum, one can say that the way neurons operate is computationally richer than is assumed in both empirical and computational neuroscience and especially in the field of machine learning. Even more, increasing evidence reveals that mechanisms in apical dendrites play a key role in cognitive cortical functions including sensory perception, prediction, learning, memory, and consciousness. For example, functional imaging of cortical L1 in behaving mice during sensory perception, where top-down inputs provide information related to behavioural saliency, shows that calcium spikes in the dendrites of L5b neurons control the perception of an external stimulus (Takahashi et al., 2016, 2020, see also Xu et al., 2012).

Fundamentally, a multi-compartmental model of neuronal function allows for the segregation of categorically separable information streams. This may be particularly important for distinct learning rules for top-down versus bottom-up pathways, both for their biological and computational implementation. For internal models to be functional, they should adapt with learning such that stimulus-evoked activity is modulated according to prior expectations. Indeed, a recent study showed that hippocampus-dependent learning involved a gating signal to L1 of mouse somatosensory cortex suggesting that consolidation of semantic memory involves a specific process in the apical compartment of pyramidal neurons (Doron et al., 2020; see also Aru et al., 2020a). Cortical feedback signals to V1 in humans encode more detailed stimulus-specific information after learning (Lazarova et al., 2023; Ortiz-Tudela et al., 2023), which presumably also relies on top-down activity in L1. In sum, cortical function relies on the integration of internally-generated top-down and bottom-up sensory information, that might ultimately be best explained at the subcellular level, considering the anatomical and computational distinction between apical and basal dendritic trees. Now the challenge is to attempt to elucidate the computations at the cellular level, for instance determining what functions are carried out in specific dendritic compartments (see Herz et al., 2006; Jordan et al., 2020; Kreutzer et al., 2020; Leugering et al., 2023; Mikulasch et al., 2023; Wybo et al., 2023), the consequence for neuronal output, the subsequent interactions between L2/3 and L5 pyramidal neurons, and the roles of interneurons in the local circuit. Only then will it be possible to determine the implications for cortical functions, the consolidation of semantic and abstract information in networks and its extraction for solving tasks and generating relevant behaviour. In short, the time is now right for determining the conceptual usefulness of dendritic computations in biological and artificial networks (Larkum, 2022).

Embedding multicompartment neurons in functional laminar microcircuits

Two-compartment neurons as described above are not universally or even widely recognised in understanding neocortical function, but their inputs and computations are essential to explain

architectures enabling the top-down processes of attention, arousal, learning and prediction, (the latter two being closely intertwined). These processes require different classes of L2/3 and L5 pyramidal neurons, allowing for different types of dendritic computation, that do not currently exist in artificial neurons. The predictive processing function of the cortex might be the core computational motif upon which to situate our theories of learning, attention and conscious awareness (Mumford, 1991; Rao and Ballard, 1999; Hohwy, 2013; Clark, 2015; Parr and Friston, 2018; Pennartz, 2022). We remain open to different ways to implement predictive processing but align with the notion that higher areas generate predictions about sensory inputs, that are compared with actual inputs, to learn and be aware of the external world (Clark, 2015; Phillips, 2023; Hawkins and Blakeslee, 2004). Awareness has been described as a 'superinference' that encompasses many low-level predictions pertaining to individual modalities and submodalities, but also spatial and situational aspects of conscious experience (e.g. Pennartz, 2015, 2022). The granularity of this predictive process might need to be revised to not only include classes of neurons (e.g. prediction error neurons, representation neurons, inhibitory interneurons) but also the apical and basal neuronal compartments via which pyramidal neurons communicate, along with the neuronal dynamics to support spatio-temporal predictions (Haider et al., 2021; Senn et al., 2023). Influential theories about the functional classes of pyramidal neurons required for predictive processing include prediction error neurons in L2/3 and internal representation neurons in L5 (Keller and Mrsic-Flogel, 2018, figure 3). Representation neurons associate top-down predictions with the bottom-up signals, and prediction error neurons compute the difference between the top-down prediction and actual input and send this error signal up to higher levels to optimise internal models if necessary (Rao and Ballard 1999; Batsos et al., 2012; see also Richter at al., 2023). This classical setup has been shown to result in object-selective and view-invariant firing patterns in higher layers of predictive coding networks (Dora et al. 2021; Brucklacher 2023).

Multiple streams of integration

The reality may be more complicated in both regimes. For example, not all L2/3 neurons show signatures of prediction error, and L2/3 contains at least two classes of pyramidal neurons (Figure 3). One class processes bottom-up sensory input and another processes top-down expectations. This functional separation is supported by evidence in mouse V1 cortex during visual occlusion, with different subsets of L2/3 neurons responding to either the full image (bottom-up information), or the missing scene content in occluded scene regions (encoding top-down predictions or a negative prediction error, a prediction that is not matched even though expected, Seignette et al., 2023a). Also using a visuomotor mismatch paradigm, two populations of L2/3 pyramidal cells have been observed, possibly encoding positive and negative mismatch signals (Jordan and Keller, 2020; Seignette et al., 2023b), and may even represent genetically distinct cell types (O'Toole et al., 2023). These two classes have been modelled as 'discriminator' neurons processing low-level activity and 'generator' neurons that project from higher to lower areas (Deperrois et al., 2022). In this scheme, classes of L2/3 pyramidal neurons are involved in calculating error signals or different types of attention signals, either bottom-up salience signals, or top-down prospective attention signals. Both types of attention signals require positive and negative prediction errors that are represented in a different subset of L2/3 pyramidal neurons, that themselves only represent a fraction of L2/3

pyramidal cells. We further propose distinct classes of L5 pyramidal neurons, one that integrates sensory input in the basal dendrite and top-down attention signals in the apical dendrites, and a second whereby the basal dendrites represent bottom-up salience signal while the apical dendrite of L5 pyramidal cells integrates top-down expectations. These L5 'awareness' neurons could integrate the two streams, via a coupling mechanism (Aru et al., 2020a). Here, L2/3 neurons would represent unconscious content while in L5, coincident apical and basal input triggering a dendritic calcium spike (whether apical input are errors or content) would contribute to awareness. This proposal is supported by evidence that L5a neurons project to L1 interneurons and L5b neurons (Ledderose et al., 2023).

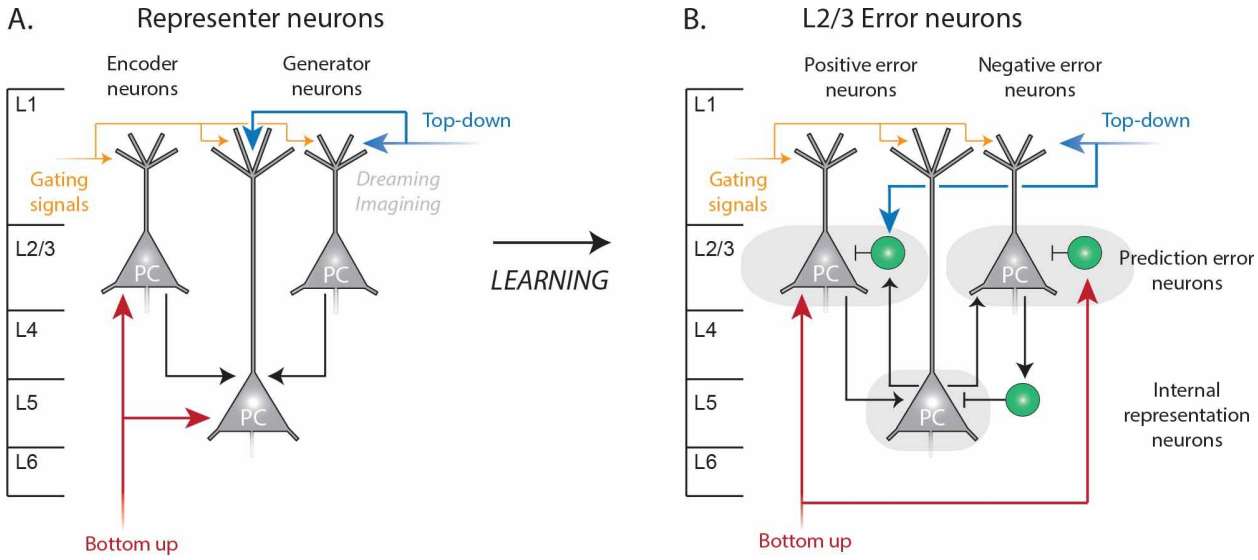

Figure 3: Separate bottom up classification networks and top-down generative networks realised in classes of L2/3 and L5 pyramidal neurons. A. One class of L2/3 pyramidal neurons, "representer" neurons, might differentially encode bottom-up information at basal sites (encoder neurons), and top-down information as inputs to apical dendrites (generator neurons). B. These representer neurons might be transformed by learning into a second class of L2/3 error neurons. During learning, encoder neurons that are driven by bottom-up inputs develop top-down inhibition. This inhibition explains away the bottom-up drive, and the neurons become positive error neurons (bottom-up minus top-down). Analogously, negative prediction error neurons develop from generator neurons by learning to subtract away the top-down drive through bottom-up inhibition. The separation of the encoding (bottom-up and generative (top-down) pathways permits the hypothesis that L2/3 neurons (encoding and generative) represent unconscious content. L5 pyramidal neurons integrate the two pathways, and matching bottom-up and top-down inputs in these L5 pyramids is a necessary ingredient for conscious awareness.

Apical input can create conscious experience

What is the nature of top-down information streams when internally-generated information is not used for contextualising bottom up streams, but for imagination or dreaming? It has been hypothesised that during dreaming, a mode of apical function, so-called "apical drive" originates from internally generated information, thus becoming the main cause of the cell's action potential generation and output. This mechanism has been proposed as a theory of

dreaming during REM sleep (Aru et al., 2020b). Although it remains to be explored empirically it offers a testable framework. The apical drive hypothesis is compatible with "the overfitted brain" hypothesis for dreaming (Hoel, 2021), which suggests that dreams may serve to avoid overfitting of the brain's internal models and to promote generalisations.

Deep convolutional neuronal network models also offer novel ideas into the biological function of dreaming, where it is proposed that rapid-eye-movement (REM) dreaming is essential for efficient memory semantisation by randomly combining episodic memories to create new, virtual sensory experiences (Deperrois et al., 2022). The states of wakefulness, NREM and REM sleep differ in their function, and in whether they have external inputs. During wakefulness, sensory inputs drive early cortex and feedforward encoding pathways, and store latent representations in higher areas, stored in the hippocampus as episodic memories. According to the synaptic tagging hypothesis, this pathway should 'tag' that there was external input (Frey and Morris, 1997), representing a delayed-in-time solution to the credit assignment problem (Gütig, 2016). Generative feedback pathways can reproduce low level activities from these abstract representations, and should minimise the mismatch between these streams. This replay scenario is compatible with electrophysiological recordings of the "Up" and "Down" states during NREM sleep (Destexhe et al., 2007).

During sleep, the network is largely disconnected from the sensory world, and it is proposed that during NREM sleep, information we stored during wakefulness can be replayed from the hippocampus (even if it was occluded), generating low-level activities, with the encoder pathway ensuring that latent representations are similar to the episodic memories (Deperrois et al., 2022). However during REM sleep, random hippocampal memories combine with spontaneous activity to create new visual content. From this new high-level representation, activity in the lower sensory cortex is generated and is again fed through the encoder or feedforward pathway. Synaptic plasticity adjusts feedforward connections to silence the activity of the discriminator output as it should learn to distinguish it from externally evoked sensory activity. Simultaneously, feedback connections are adjusted adversarially to generate activity patterns which appear externally-driven and thereby trick the discriminator into believing that the lower-level activity was externally-driven. This network architecture allows the neuronal network to learn to generalise to many hypothetically possible situations, including those that have never been experienced (visual experience of a particular car in a street) but make the neuronal architecture robust to include these simulated data sets in its repertoire of recognisable situations.

The cellular hypothesis of dreaming (Aru etal., 2020b) and the functional model (Deperrois et al., 2022) can be reconciled by postulating that the forward encoding pathway and the feedback generative pathway are represented by separate streams of L2/3 pyramidal neurons. The content of these L2/3 streams is not perceived, but they are integrated in L5 pyramidal neurons which do contribute to perceived content. L5 pyramidal neurons encode the match between the encoding and generative pathway through dendritic calcium spikes (Aru et al., 2020a, Takahashi et al., 2020). These L5 pyramidal neurons may also be active when a subject dreams, as opposed to cortical activity during sleep that is not experienced as dream. This hypothesis that dreaming is a mode of apical function in the superficial layers (Aru et al., 2020b) could help to reconcile puzzling findings in individuals with aphantasia, an inability to voluntarily imagine visual content. Aphantasics are unable to perform voluntary visual imagery

during awake states but around 63% report to dream visually (compared to around 89% of individuals with average imagery vividness; Zeman et al., 2020). Aphantasia might therefore involve functional differences in deep layer feedback processing, whereas superficial layer feedback processes may be less affected. This hypothesis is compatible with laminar brain imaging in humans showing that imagery is preferentially detectable in deeper layers of primary visual cortex (Bergmann et al., 2019). Here, top-down generated input might arrive to apical dendrites in L1, that drive L5b neurons, because we are aware of the imagined content, and this can be detected with laminar fMRI. This is quite different from cases of visual occlusion, where we do not consciously perceive scene-specific content in the occluded image region (but that leads to prediction error signals related to the absence of expected information in superficial layers). Under circumstances where internal activity drives conscious experience, as in dreaming or imagining, the higher-order thalamus, which is in a recurrent loop with L1 (also L4-5) of perceptual regions and with prefrontal cortex, might be involved in gating ascending sensory information with descending information from the prefrontal cortex, downscaling the weight given to the ascending information in the cases where we generate imagined content at the same time as perceiving our environment.

Context-dependent neural computation in next-generation AI and neuromorphic computing

Does the microcircuit structure laid out above provide a blueprint for a computational motif that is computationally advantageous for building artificial intelligent systems? While artificial intelligence offers novel approaches for neuroscientific investigations (Richards et al., 2022), experimental observations and computational insights into behaving cortical circuits may also advance AI. Predictions and contextual computations, as discussed above, are also key computational motifs of the two most important modern AI architectures: convolutional neural networks and transformers. At the most general level, computations that integrate spatial and temporal context, including predictions, require a recurrent computational architecture (van Bergen et al., 2020). Although feedforward deep neural networks still dominate the engineering of perceptual systems, adding recurrence to convolutional neural networks can improve their performance at perceptual inference tasks (Liang and Hu, 2015; Spoerer et al. 2017; Linsley et al. 2018). Recurrent convolutional models have been shown to flexibly trade off speed and accuracy depending on current task demands and predict how long it will take humans to recognize images (Spoerer et al. 2020). While these models implement dynamic recurrent inference processes that take context into account, they do not yet incorporate the detailed circuit motifs discussed above and it is unclear to what extent their computations employ predictive mechanisms.

One of the most impactful classes of machine learning architectures in recent years is transformers models enabling breakthroughs in generative vision and language modelling. The key mechanism of transformers is “self-attention” (Vaswani et al., 2017), which enables context-dependent computations. Like convolution, self-attention in vision transformers allows a functional column (e.g., a location in a spatial representation in the network) to integrate information from other, neighbouring columns. Convolution uses fixed integration weights over the neighbouring columns. Self-attention, in contrast, computes the integration weights dynamically and as a function of the affinity with the neighbouring columns. Affinity-dependent dynamic computations enable complex information integration schemes and the success of

transformer architectures emphasises the importance of contextual computations. Unlike the rodent and primate visual system, transformer models are typically feedforward and the contextual signal is only computed within the same representational layer (i.e., "self"-attention). Cortical microcircuits also implement top-down attention signals, implying recurrent processing and integration of more global information. It remains to be seen to what extent integrating these inspirations from biology can benefit modern artificial neural network architectures like transformers.

Contextual information integration with dynamic weights between representational units implies a departure from point-units with fixed nonlinearities (Pagalos et al., 2023b; Poirazi et al., 2003; Beniaguev et al., 2021; Poirazi and Papoutsi, 2020; Gidon et al., 2020; Guerguiev et al., 2017; Grewal et al., 2021; Wu et al., 2018; Tzilivaki et al., 2019; Jones and Kording, 2021). A number of AI experts have been inspired by the explanatory power of few-compartment neuron models for cortical computation (Lillicrap, 2020; see also Guerguiev et al., 2017; Haider et al., 2021; Payeur et al., 2021; Max et al., 2022 for the computational role of compartmental neurons), and one next step will be showing the information processing capabilities of these neurons to process large-scale data. Some earlier machine learning approaches inspired by these neurophysiological discoveries focused predominantly on learning, but using context to guide both ongoing processing and learning might be central to unlocking the potential of these networks (Adeel et al., 2022). It has recently been shown that deep networks of CS-TPNs (context-sensitive two point neurons) can process large-scale real-world multisensory data 1250X more efficiently than networks composed of point neurons (PNs, Adeel et al., 2023a, and Figure 1C). Furthermore, going beyond a point-neuron-inspired Transformer model (the backbone of ChatGPT), a CS-TPN-inspired Cooperator model with the same number of parameters, learns quicker than a Transformer (Adeel et al., 2023b, 2023c). Context-sensitive two-point neurons amplify and suppress the transmission of information when the context shows it to be relevant and irrelevant, respectively. A deep neural network composed of such local processors seeks to maximise agreement between the active neurons, extracting relevant features at very early stages, thus restricting the transmission of conflicting information to higher levels and reducing the neural activity required to process large amounts of heterogeneous real-world data. Such findings contribute to the shift in modern neuromorphic computing to implement multicompartment neurons, offering new possibilities for AI systems (Schemmel et al., 2017; Amir et al., 2018; Yang et al., 2019; Davies et al., 2021; Gao et al., 2022; Khacef et al., 2022; Ward and Rhodes, 2022; Pagkalos et al., 2023a; see also Urbanczik and Senn, 2014).

Ongoing directions in experimental, modelling and machine learning domains

Experimental challenges include bringing the (local) neuronal compartmentalization and (local) laminar fMRI findings into a wider brain context, i.e. simultaneously assessing the source and target of feedback signals at multi-scales is crucial. This information would provide the means to test models casually at a brain-wide scale. The obvious next step would then be to characterise the underlying cellular mechanisms although this may be only possible in the long term. There is growing evidence that AI algorithms and neuromorphic systems benefit from incorporating dendritic properties into their architectures. In terms of machine learning, forms of deep learning that integrate two-point neurons are a step change in transforming the cellular

foundations of deep network architectures. Future research should advance cross-links of these artificial networks to behavioural phenomena and cortical microcircuit architectures. For example, how do neural networks (such as Figure 1C) with built-in algorithms that infer the relevance of their inputs using 'apical' processes relate to the brain measures we presented showing the 'filling-in' of the occluded image regions based on internal context? And how do DNNs that incorporate two-point neurons map onto the microcircuit functions we expressed above using encoding streams and generative streams in L2/3, and a L5 stream that integrates the bottom-up and top-down streams?

Concluding section

We recognise that a number of important studies were beyond our already far-reaching review of mechanistic neuronal models for how the cortex uses feedback processing in behaviour. In reviewing, we made a number of points. Assessing cognitive functions and neuronal substrates of top-down microcircuits can be achieved with the occlusion paradigm. We showed evidence in humans that top-down signals include predictive world models that are updated by information from our senses. The occlusion paradigm can also be implemented in a multispecies approach, with surprisingly consistent results in human and nonhuman primates and mice, opening the door for studying the involvement of finer-grained (cellular) neural mechanisms in feedback processing. Top-down feedback inputs arrive mainly to a sub-compartment of the neuron in the apical dendrites, and apical dendrites have specialised computational abilities afforded by a range of linear and nonlinear mechanisms. This spatial separation of inputs within a pyramidal neuron has causal implications for perception, learning, memory, and consciousness. Algorithms that describe dendritic computation and that are consequential for behaviour should inspire more powerful machine learning approaches. While making these points, at times we bridged to topics and bodies of literature in a more cursory manner, but we intended to give the breadth of what we consider the fundamental implications of context-sensitive neurons, and their role in cognition.

We propose that the next radical shift in our understanding of the neocortex will come from algorithmic and conceptual points of view that incorporate the morphological and biophysical features of cortical pyramidal neurons and their distinct computational sub-compartments. These features are often overlooked in neurobiological investigations and when modelling artificial neurons. This intraneuronal architecture offers the evolutionary gain of increased computational power and the parsing of information streams in single cells, necessary for the internal modelling of the outside world with input from the senses. The emerging field of 'cellular psychology', the cellular foundations of cognition, suggests that the role of context-dependent coding in pyramidal neurons ranges from perception to self-development, from levels of consciousness to theory of mind (Philips, 2023). This paradigm shift should provide explanation at a level incorporating the biophysical specialisations of individual pyramidal neurons into a mechanistic understanding of principles underlying basic cognitive functions and conscious experience.

Box 1 Future research themes

Pathological states and dendritic integration dysfunction

Several pathologies of conscious experience and intellectual function are associated with apical dendritic malfunction or its regulation, including epilepsy, psychoses in the schizophrenia spectrum, anti-NMDA encephalitis, a heterogeneous class of pervasive neurodevelopmental disorders in the autism spectrum including Fragile-X, other neurodevelopmental disorders of intellectual function such as Down's and Timothy syndromes, and foetal alcohol spectrum disorders (see Palmer, 2014; Johnston, Frick, and Poolos, 2016; Sanders et al., 2018; Granato and Merighe, 2021; Nelson and Bender, 2021). Though prima facie different, there is overlap between these disorders, involving various combinations and extents of malfunction in basic mechanisms on which the functions of apical input to neocortical pyramidal cells depends. These mechanisms include the over- or under-activation of mechanisms including BAC-firing, the regulation of apical function via current flow through HCN channels, the integration of apical branch input and communication of that to the soma via NMDA spikes, enhancing coincidence detection via metabotropic glutamate receptors, and the prenatal development of the apical branches. The precise aetiology presumably involves various combinations of gene polymorphisms, mutations, and interactions with unique prenatal and postnatal environmental factors. But the appeal of linking these pathologies also to apical dendritic malfunction is that it provides a coherent functional explanation whilst simultaneously suggesting ways to investigate and even design therapeutic approaches.

Species-specific morphological differences

One fundamental question that has to be rigorously tested is whether the cellular properties of cortical neurons and their connectivity in rodents and humans are largely the same. Does the size of neurons, or their membrane properties matter? Is laminar connectivity in mice and humans the same in temporal, sensory or prefrontal cortices? Are feedback circuits organised similarly? There is increasing evidence that dendritic and other postsynaptic mechanisms are a key site of differences between species. A major issue for future investigation is the extent to which distinctively human cognitive capabilities arise from distinctively human morphological neuronal features and subcellular processes, particularly those involving enhanced forms of cellular context-sensitivity. Comparisons between human and rodent L5 pyramidal cells *in vitro* have revealed that the greater apical dendritic length in human L5 neurons severely weakens their somato-dendritic electrical coupling (Beaulieu-Laroche et al., 2018). Thus, distal dendritic synaptic input, even when boosted by dendritic calcium spikes, caused weaker somatic spiking in humans than in rat L5 neurons. However, *in vivo* recordings in mice suggest that the degree of electrical compartmentalization is far less in the intact brain than in brain slices (Beaulieu-Laroche et al. 2019), suggesting that functional implications of this species difference may be other than assumed from *in vitro* studies. Thus, a challenge for future research will be to identify functionally important species differences in somato-dendritic communication *in vivo*, and the actual roles of dendritic processing in the intact, functioning cortex, particularly in L5 and other cortical principal neurons. To this end, it

will be essential to develop better tools for specifically targeting dendritic mechanisms, in order to determine their functions, as distinct from confounding off-target circuit effects (Francioni and Harnett 2022).

Acknowledgements
This work was supported by the European Union's Horizon 2020 Framework Program for Research and Innovation under the Specific Grant Agreements 720270, 785907 and 945539 (Human Brain Project SGA1, SGA2, SGA3). We thank Arnau Manasanch for technical support with data and live figures, and the Human Brain Project for assistance with data and model curation. We thank Feng Wang, Xing Chen and Amparo Gilhuis for contributions to the monkey experiment.

Interactive Live Paper
An interactive version of this manuscript is available at: https://live-papers.brainsimulation.eu/#2023-muckli-et-al
The Cooperator vs Transformer demo is available here: http://cmilab.org/research/#robotics

Data availability - partial visual occlusion
The human, monkey and mouse occlusion data are available on EBRAINS to download, or as Jupyter notebooks that can be used to reproduce the experimental data. Access to this resource requires free user registration at https://www.ebrains.eu/page/sign-up
Morgan, A. T., Petro, L., & Muckli, L. (2021). fMRI data of early visual cortex while viewing 24 occluded scenes [Data set]. EBRAINS. https://doi.org/10.25493/Z60A-BGY
Papale, P., Wang, F., Morgan, A. T., Chen, X., Gilhuis, A., Petro, L., Muckli, L., Roelfsema, P. R., & Self, M. W. (2022). Electrophysiological recordings in macaque V1 during passive viewing of full and occluded natural scenes (v1) [Data set]. EBRAINS. https://doi.org/10.25493/KABE-GS0
Seignette, K., Van Der Togt, C., & Levelt, C. N. (2022). Two-photon calcium imaging of layer 2/3 somas and layer 5 dendrites in mouse visual cortex during visual occlusion (v1) [Data set]. EBRAINS. https://doi.org/10.25493/NXRY-0W6
(See Code availability below for the autoencoder model used to reconstruct the human occlusion data).

Model availability
Deperrois, N., Petrovici, M. A., Senn, W., & Jordan, J. (2023). Learning cortical representations through perturbed and adversarial dreaming (v1). EBRAINS.
https://doi.org/10.25493/A1WA-5RG
Svanera, M., Morgan, A. T., Petro, L., & Muckli, L. (2022). Self-Supervised Deep Neural Network for Image Completion (v1.0.0). EBRAINS. https://doi.org/10.25493/5Z86-K1Z
Spoerer, C.J., Kriegeskorte, N. Recurrent deep networks with feedback for modern AI (1.0).
https://search.kg.ebrains.eu/instances/eae6a13b-4006-49df-bdcc-ec27c2409e72

References

Aamir, S. A., Müller, P., Kiene, G., Kriener, L., Stradmann, Y., Grübl, A., ... & Meier, K. (2018). A mixed-signal structured AdEx neuron for accelerated neuromorphic cores. *IEEE transactions on biomedical circuits and systems*, *12*(5), 1027-1037.
Acharya, J., Basu, A., Legenstein, R., Limbacher, T., Poirazi, P., & Wu, X. (2022). Dendritic Computing: Branching Deeper into Machine Learning. *Neuroscience*, 489, 275–289.
Adeel, A. (2020). Conscious multisensory integration: introducing a universal contextual field in biological and deep artificial neural networks. *Frontiers in Computational Neuroscience*, *14*, 15.
Adeel, A., Franco, M., Raza, M., & Ahmed, K. (2022). Context-sensitive neocortical neurons transform the effectiveness and efficiency of neural information processing. *arXiv preprint arXiv:2207.07338*.
Adeel, A., Adetomi, A., Ahmed, K., Hussain, A., Arslan, T., & Phillips, W. A. (2023a). Unlocking the potential of two-point cells for energy-efficient and resilient training of deep nets. *IEEE Transactions on Emerging Topics in Computational Intelligence*.
Adeel, A., Muzaffar, J., Ahmed, K., & Raza, M. (2023b). Cooperation Is All You Need. *arXiv preprint arXiv:2305.10449*.
Adeel, A., Adetomi, A., Ahmed, K., Hussain, A., Arslan, T., & Phillips, W. A. (2023c). Unlocking the potential of two-point cells for energy-efficient and resilient training of deep nets. IEEE Transactions on Emerging Topics in Computational Intelligence, Erratum (in preparation).
Adeel, A., Muzaffar, J., Ahmed, K., & Raza, M. (2023d). Cooperation Is All You Need. arXiv preprint arXiv:2305.10449, Erratum (in preparation).
Amunts, K., DeFelipe, J., Pennartz, C., Destexhe, A., Migliore, M., Ryvlin, P., ... & Jirsa, V. (2022). Linking brain structure, activity, and cognitive function through computation. *eneuro*, *9*(2).
Aru, J., Suzuki, M., & Larkum, M. E. (2020a). Cellular mechanisms of conscious processing. *Trends in Cognitive Sciences*, *24*(10), 814-825.
Aru, J., Siclari, F., Phillips, W. A., & Storm, J. F. (2020b). Apical drive—A cellular mechanism of dreaming?. *Neuroscience & Biobehavioral Reviews*, *119*, 440-455.
Attinger, A., Wang, B., & Keller, G. B. (2017). Visuomotor coupling shapes the functional development of mouse visual cortex. *Cell*, *169*(7), 1291-1302.
Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, *76*(4), 695-711.
Beaulieu-Laroche, L., Toloza, E. H., Van der Goes, M. S., Lafourcade, M., Barnagian, D., Williams, Z. M., ... & Harnett, M. T. (2018). Enhanced dendritic compartmentalization in human cortical neurons. *Cell*, *175*(3), 643-651.
Beaulieu-Laroche, L., Toloza, E. H., Brown, N. J., & Harnett, M. T. (2019). Widespread and highly correlated somato-dendritic activity in cortical layer 5 neurons. *Neuron*, *103*(2), 235-241.
Beniaguev, D., Segev, I., & London, M. (2021). Single cortical neurons as deep artificial neural networks. *Neuron*, *109*(17), 2727-2739.
Bergmann, J., Morgan, A. T., & Muckli, L. (2019). Two distinct feedback codes in V1 for 'real' and 'imaginary' internal experiences. *BioRxiv*, 664870.
Brucklacher, M., Bohte, S. M., Mejias, J. F., & Pennartz, C. M. (2022). Local minimization of prediction errors drives learning of invariant object representations in a generative network model of visual perception. *Frontiers in Computational Neuroscience.*
Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.

Davies, M., Wild, A., Orchard, G., Sandamirskaya, Y., Guerra, G. A. F., Joshi, P., ... & Risbud, S. R. (2021). Advancing neuromorphic computing with loihi: A survey of results and outlook. *Proceedings of the IEEE*, *109*(5), 911-934.
Dayan, P., Hinton, G. E., Neal, R. M., & Zemel, R. S. (1995). The helmholtz machine. *Neural computation*, *7*(5), 889-904.
Deperrois, N., Petrovici, M. A., Senn, W., & Jordan, J. (2022). Learning cortical representations through perturbed and adversarial dreaming. *Elife*, *11*, e76384.
Destexhe, A., Hughes, S. W., Rudolph, M., & Crunelli, V. (2007). Are corticothalamic 'up' states fragments of wakefulness?. *Trends in neurosciences*, *30*(7), 334-342.
Dora, S., Bohte, S. M., & Pennartz, C. M. (2021). Deep gated Hebbian predictive coding accounts for emergence of complex neural response properties along the visual cortical hierarchy. *Frontiers in Computational Neuroscience*, *15*, 666131.
Doron, G., Shin, J. N., Takahashi, N., Drüke, M., Bocklisch, C., Skenderi, S., ... & Larkum, M. E. (2020). Perirhinal input to neocortical layer 1 controls learning. *Science*, *370*(6523), eaaz3136.
Francioni, V., & Harnett, M. T. (2022). Rethinking single neuron electrical compartmentalization: dendritic contributions to network computation in vivo. *Neuroscience*, *489*, 185-199.
Frey, U., & Morris, R. G. (1997). Synaptic tagging and long-term potentiation. *Nature*, *385*(6616), 533-536.
Friston, K. (2005). A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological sciences*, *360*(1456), 815-836.
Gao, T., Deng, B., Wang, J., & Yi, G. (2022). Highly efficient neuromorphic learning system of spiking neural network with multi-compartment leaky integrate-and-fire neurons. *Frontiers in Neuroscience*, *16*, 929644.
Garner, A. R., & Keller, G. B. (2022). A cortical circuit for audio-visual predictions. Nature neuroscience, 25(1), 98-105.
Gau, R., Bazin, P. L., Trampel, R., Turner, R., & Noppeney, U. (2020). Resolving multisensory and attentional influences across cortical depth in sensory cortices. *Elife*, *9*, e46856.
Gidon, A., Zolnik, T. A., Fidzinski, P., Bolduan, F., Papoutsi, A., Poirazi, P., ... & Larkum, M. E. (2020). Dendritic action potentials and computation in human layer 2/3 cortical neurons. *Science*, *367*(6473), 83-87.
Granato, A., & Merighi, A. (2022). Dendrites of neocortical pyramidal neurons: the key to understand intellectual disability. *Cellular and Molecular Neurobiology*, *42*(1), 147-153.
Grewal, K., Forest, J., Cohen, B. P., & Ahmad, S. (2021). Going Beyond the Point Neuron: Active Dendrites and Sparse Representations for Continual Learning (p. 2021.10.25.465651). bioRxiv.
Guerguiev, J., Lillicrap, T. P., & Richards, B. A. (2017). Towards deep learning with segregated dendrites. eLife, 6, e22901.
Gütig, R. (2016). Spiking neurons can discover predictive features by aggregate-label learning. *Science*, *351*(6277), aab4113.
Haider, P., Ellenberger, B., Kriener, L., Jordan, J., Senn, W., & Petrovici, M. A. (2021). Latent equilibrium: A unified learning theory for arbitrarily fast computation with arbitrarily slow neurons. *Advances in Neural Information Processing Systems*, *34*, 17839-17851.
Hawkins, J., & Blakeslee, S. (2004). *On intelligence*. Macmillan.
Herz, A. V., Gollisch, T., Machens, C. K., & Jaeger, D. (2006). Modeling single-neuron dynamics and computations: a balance of detail and abstraction. *science*, *314*(5796), 80-85.

Hoel, E. (2021). The overfitted brain: Dreams evolved to assist generalization. *Patterns*, *2*(5).
Hohwy, J. (2013). *The predictive mind*. OUP Oxford.
Iyer, A., Grewal, K., Velu, A., Souza, L. O., Forest, J., & Ahmad, S. (2022). Avoiding catastrophe: Active dendrites enable multi-task learning in dynamic environments. *Frontiers in neurorobotics*, *16*, 846219.
Johnston, D., Frick, A., & Poolos, N. (2016). Dendrites and disease. *Dendrites*.
Jones, I. S., & Kording, K. P. (2021). Might a single neuron solve interesting machine learning problems through successive computations on its dendritic tree?. *Neural Computation*, *33*(6), 1554-1571.
Jordan, R., & Keller, G. B. (2020). Opposing influence of top-down and bottom-up input on excitatory layer 2/3 neurons in mouse primary visual cortex. *Neuron*, *108*(6), 1194-1206.
Jordan, J., Petrovici, M. A., Senn, W., & Sacramento, J. (2020, March). Conductance-based dendrites perform reliability-weighted opinion pooling. In *Proceedings of the 2020 Annual Neuro-Inspired Computational Elements Workshop* (pp. 1-3).
Keller, G. B., & Mrsic-Flogel, T. D. (2018). Predictive processing: a canonical cortical computation. *Neuron*, *100*(2), 424-435.
Keller, A. J., Roth, M. M., & Scanziani, M. (2020). Feedback generates a second receptive field in neurons of the visual cortex. *Nature*, *582*(7813), 545-549.
Khacef, L., Klein, P., Cartiglia, M., Rubino, A., Indiveri, G., & Chicca, E. (2022). Spike-based local synaptic plasticity: A survey of computational models and neuromorphic circuits. *arXiv preprint arXiv:2209.15536*.
Kirchberger, L., Mukherjee, S., Schnabel, U. H., van Beest, E. H., Barsegyan, A., Levelt, C. N., ... & Roelfsema, P. R. (2021). The essential role of recurrent processing for figure-ground perception in mice. *Science advances*, *7*(27), eabe1833.
Kirchberger, L., Mukherjee, S., Self, M. W., & Roelfsema, P. R. (2023). Contextual drive of neuronal responses in mouse V1 in the absence of feedforward input. *Science advances*, *9*(3), eadd2498.
Kok, P., Bains, L. J., van Mourik, T., Norris, D. G., & de Lange, F. P. (2016). Selective activation of the deep layers of the human primary visual cortex by top-down feedback. *Current Biology*, *26*(3), 371-376.
Kreutzer, E., Petrovici, M. A., & Senn, W. (2020). Natural gradient learning for spiking neurons. *eLife*.
Larkum, M. E., Zhu, J. J., & Sakmann, B. (2001). Dendritic mechanisms underlying the coupling of the dendritic with the axonal action potential initiation zone of adult rat layer 5 pyramidal neurons. *The Journal of physiology*, *533*(2), 447-466.
Larkum, M. E., Nevian, T., Sandler, M., Polsky, A., & Schiller, J. (2009). Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle. *Science*, *325*(5941), 756-760.
Larkum, M. (2013). A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. Trends in neurosciences, 36(3), 141-151.
Larkum, M. E., Petro, L. S., Sachdev, R. N., & Muckli, L. (2018). A perspective on cortical layering and layer-spanning neuronal elements. *Frontiers in neuroanatomy*, 56.
Larkum, M. E. (2022). Are dendrites conceptually useful?. *Neuroscience*, *489*, 4-14.
Larkum, M. E., Wu, J., Duverdin, S. A., & Gidon, A. (2022). The guide to dendritic spikes of the mammalian cortex in vitro and in vivo. *Neuroscience*, *489*, 15-33.

Lawrence, S. J., Norris, D. G., & De Lange, F. P. (2019). Dissociable laminar profiles of concurrent bottom-up and top-down modulation in the human visual cortex. *Elife*, *8*, e44422.
Lazarova, Y., Huang, Y., Muckli, L, Petro, L.S (2023). Perceptual priors add sensory detail to contextual feedback processing in V1. BioRxiv.
Ledderose, J., Zolnik, T. A., Toumazou, M., Trimbuch, T., Rosenmund, C., Eickholt, B. J., Jaeger, D., Larkum, M.E. & Sachdev, R. N. (2023). Input to layer 1 of somatosensory cortex: Local input outweighs long-range and can modulate the activity of layer 1 interneurons. *Cerebral Cortex, in press, 1-19.*
Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *JOSA A*, *20*(7), 1434-1448.
Leugering, J., Nieters, P., & Pipa, G. (2023). Dendritic plateau potentials can process spike sequences across multiple time-scales. *Frontiers in Cognition*, *2*, 1044216.
Liang, M. and Hu, X., (2015). Recurrent convolutional neural network for object recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3367-3375.
Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J., & Hinton, G. (2020). Backpropagation and the brain. *Nature Reviews Neuroscience*, *21*(6), 335-346.
Linsley, D., Kim, J., Veerabadran, V., Windolf, C. and Serre, T. (2018). Learning long-range spatial dependencies with horizontal gated recurrent units. *Advances in neural information processing systems*, *31*.
Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, *453*(7197), 869-878.
London, M., & Häusser, M. (2005). Dendritic computation. *Annu. Rev. Neurosci.*, *28*, 503-532.
Major, G., Larkum, M. E., & Schiller, J. (2013). Active properties of neocortical pyramidal neuron dendrites. *Annual review of neuroscience*, *36*, 1-24.
Manita, S., Suzuki, T., Homma, C., Matsumoto, T., Odagawa, M., Yamada, K., ... & Murayama, M. (2015). A top-down cortical circuit for accurate sensory perception. *Neuron*, *86*(5), 1304-1316.
Markov, N. T., & Kennedy, H. (2013). The importance of being hierarchical. *Current opinion in neurobiology*, *23*(2), 187-194.
Marquardt, I., De Weerd, P., Schneider, M., Gulban, O. F., Ivanov, D., Wang, Y., & Uludağ, K. (2020). Feedback contribution to surface motion perception in the human early visual cortex. *Elife*, *9*, e50933.
Max, K., Kriener, L., Pineda García, G., Nowotny, T., Senn, W., & Petrovici, M. A. (2023, September). Learning efficient backprojections across cortical hierarchies in real time. In *International Conference on Artificial Neural Networks* (pp. 556-559). Cham: Springer Nature Switzerland.
Mikulasch, F. A., Rudelt, L., Wibral, M., & Priesemann, V. (2023). Where is the error? Hierarchical predictive coding through dendritic error computation. *Trends in Neurosciences*, *46*(1), 45-59.
Morgan, A. T., Petro, L. S., & Muckli, L. (2019). Scene representations conveyed by cortical feedback to early visual cortex can be described by line drawings. *Journal of Neuroscience*, *39*(47), 9410-9423.
Muckli, L., De Martino, F., Vizioli, L., Petro, L. S., Smith, F. W., Ugurbil, K., ... & Yacoub, E. (2015). Contextual feedback to superficial layers of V1. *Current Biology*, *25*(20), 2690-2695.
Mumford, D. (1991). On the computational architecture of the neocortex: I. The role of the thalamo-cortical loop. *Biological cybernetics*, *65*(2), 135-145.

Murayama, M., & Larkum, M. E. (2009). Enhanced dendritic activity in awake rats. Proceedings of the National Academy of Sciences, 106(48), 20482-20486.
Murayama, M., Pérez-Garci, E., Nevian, T., Bock, T., Senn, W., & Larkum, M. E. (2009). Dendritic encoding of sensory stimuli controlled by deep cortical interneurons. Nature, 457(7233), 1137-1141.
Nelson, A. D., & Bender, K. J. (2021). Dendritic integration dysfunction in neurodevelopmental disorders. *Developmental neuroscience*, *43*(3-4), 201-221.
Ortiz-Tudela, J., Bergmann, J., Bennett, M., Ehrlich, I., Muckli, L., & Shing, Y. L. (2023). Concurrent contextual and time-distant mnemonic information co-exist as feedback in the human visual cortex. *Neuroimage*, *265*, 119778.
O'Toole, S. M., Oyibo, H. K., & Keller, G. B. (2023). Molecularly targetable cell types in mouse visual cortex have distinguishable prediction error responses. *Neuron*.
Pagkalos, M., Makarov, R., & Poirazi, P. (2023a). Leveraging dendritic properties to advance machine learning and neuro-inspired computing. *arXiv preprint arXiv:2306.08007*.
Pagkalos, M., Chavlis, S., & Poirazi, P. (2023b). Introducing the Dendrify framework for incorporating dendrites to spiking neural networks. *Nature Communications*, *14*(1), 131.
Pak, A., Ryu, E., Li, C., & Chubykin, A. A. (2020). Top-down feedback controls the cortical representation of illusory contours in mouse primary visual cortex. *Journal of Neuroscience*, *40*(3), 648-660.
Palmer, L. M. (2014). Dendritic integration in pyramidal neurons during network activity and disease. *Brain research bulletin*, *103*, 2-10.
Papale, P., Zuiderbaan, W., Teeuwen, R. R., Gilhuis, A., Self, M. W., Roelfsema, P. R., & Dumoulin, S. O. (2021). The influence of objecthood on the representation of natural images in the visual cortex. *bioRxiv*, 2021-09.
Papale, P., Wang, F., Morgan, A. T., Chen, X., Gilhuis, A., Petro, L. S., Muckli, L., Roelfsema, P.R., Self, M. W. (2023). The representation of occluded image regions in area V1 of monkeys and humans. *Current Biology*.
Paquola, C., Amunts, K., Evans, A., Smallwood, J., & Bernhardt, B. (2022). Closing the mechanistic gap: the value of microarchitecture in understanding cognitive networks. *Trends in Cognitive Sciences.*
Pardi, M. B., Schroeder, A., & Letzkus, J. J. (2022). Probing top-down information in neocortical layer 1. *Trends in Neurosciences*.
Parr, T., & Friston, K. J. (2018). The anatomy of inference: generative models and brain structure. *Frontiers in computational neuroscience*, *12*, 90.
Payeur, A., Guerguiev, J., Zenke, F., Richards, B. A., & Naud, R. (2021). Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nature neuroscience*, *24*(7), 1010-1019.
Pearson, M. J., Dora, S., Struckmeier, O., Knowles, T. C., Mitchinson, B., Tiwari, K., ... & Pennartz, C. (2021). Multimodal representation learning for place recognition using deep Hebbian predictive coding. Frontiers in Robotics and AI, 8, 732023.
Pennartz, C. M. (2015). *The brain's representational power: on consciousness and the integration of modalities*. MIT Press.
Pennartz, C. M., Dora, S., Muckli, L., & Lorteije, J. A. (2019). Towards a unified view on pathways and functions of neural recurrent processing. *Trends in neurosciences*, *42*(9), 589-603.

Pennartz, C. M. (2022). What is neurorepresentationalism? From neural activity and predictive processing to multi-level representations and consciousness. Behavioural Brain Research, 432, 113969.
Pennartz, C. M., Oude Lohuis, M. N., & Olcese, U. (2023). How 'visual' is the visual cortex? The interactions between the visual cortex and other sensory, motivational and motor systems as enabling factors for visual perception. Philosophical Transactions of the Royal Society B, 378(1886), 20220336.
Peters, B., & Kriegeskorte, N. (2021). Capturing the objects of vision with neural networks. *Nature human behaviour*, *5*(9), 1127-1144.
Petro, L. S., Smith, F. W., Abbatecola, C., & Muckli, L. (2023). The spatial precision of contextual feedback signals in human V1. *Biology*, *12*(7), 1022.
Philips, W.A. (2023). The Cooperative Neuron: Cellular Foundations of Mental Life. OUP Oxford.
Poirazi, P., Brannon, T., & Mel, B. W. (2003). Pyramidal Neuron as Two-Layer Neural Network. Neuron, 37(6), 989–999.
Poirazi, P., & Papoutsi, A. (2020). Illuminating dendritic function with computational models. *Nature Reviews Neuroscience,* 21(6), 303–321.
Poort, J., Raudies, F., Wannig, A., Lamme, V. A., Neumann, H., & Roelfsema, P. R. (2012). The role of attention in figure-ground segregation in areas V1 and V4 of the visual cortex. *Neuron*, *75*(1), 143-156.
Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, *2*(1), 79-87.
Revina, Y., Petro, L. S., & Muckli, L. (2018). Cortical feedback signals generalise across different spatial frequencies of feedforward inputs. *Neuroimage*, *180*, 280-290.
Richards, B., Tsao, D., & Zador, A. (2022). The application of artificial intelligence to biology and neuroscience. *Cell*, *185*(15), 2640-2643.
Richter, D., Kietzmann, T. C., & de Lange, F. P. (2023). High-level prediction errors in low-level visual cortex. *bioRxiv*, 2023-08.
Sanders, S. J., Campbell, A. J., Cottrell, J. R., Moller, R. S., Wagner, F. F., Auldridge, A. L., ... & Bender, K. J. (2018). Progress in understanding and treating SCN2A-mediated disorders. *Trends in neurosciences*, *41*(7), 442-456.
Schemmel, J., Kriener, L., Müller, P., & Meier, K. (2017). An accelerated analog neuromorphic hardware system emulating NMDA-and calcium-based non-linear dendrites. In *2017 International Joint Conference on Neural Networks (IJCNN)* (pp. 2217-2226). IEEE.
Schuman, B., Dellal, S., Prönneke, A., Machold, R., & Rudy, B. (2021). Neocortical layer 1: an elegant solution to top-down and bottom-up integration. *Annual review of neuroscience*, *44*, 221-252.
Seignette, K., Papale, P., De Kraker, L., Neering, P., van Der Aa, M., Hobo, B., van der Togt, C., M. W. Self, M.W., Levelt., C.N. (2023a). Layer-specific plasticity of feedforward and contextual neuronal populations in mouse visual cortex. Society for Neuroscience abstract.
Seignette, K., Jamann, N., Papale, P., Terra, H., Porneso, R. P., de Kraker, L., ... & Levelt, C. N. (2023b). Visuomotor experience induces functional and structural plasticity of chandelier cells. *bioRxiv*, 2023-04.
Senn, W., Dold, D., Kungl, A. F., Ellenberger, B., Jordan, J., Bengio, Y., ... & Petrovici, M. A. (2023). A neuronal least-action principle for real-time learning in cortical circuits. *eLife.*

Smith, F. W., & Muckli, L. (2010). Nonstimulated early visual areas carry information about surrounding context. *Proceedings of the National Academy of Sciences*, *107*(46), 20099-20103.
Spoerer, C. J., McClure, P., & Kriegeskorte, N. (2017). Recurrent convolutional neural networks: a better model of biological object recognition. *Frontiers in psychology*, *8*, 1551.
Spoerer, C. J., Kietzmann, T. C., Mehrer, J., Charest, I., & Kriegeskorte, N. (2020). Recurrent neural networks can explain flexible trading of speed and accuracy in biological vision. *PLoS computational biology*, *16*(10), e1008215.
Stöckel, A., & Eliasmith, C. (2022). Computational properties of multi-compartment LIF neurons with passive dendrites. *Neuromorphic Computing and Engineering*, *2*(2), 024011.
Suzuki, M., & Larkum, M. E. (2020). General anesthesia decouples cortical pyramidal neurons. *Cell*, *180*(4), 666-676.
Suzuki, M., Pennartz, C.M.A., Aru, J. (2023) How deep is the brain? The shallow brain hypothesis. *Nature Reviews Neuroscience* (in press).
Svanera, M., Morgan, A. T., Petro, L. S., & Muckli, L. (2021). A self-supervised deep neural network for image completion resembles early visual cortex fMRI activity patterns for occluded scenes. *Journal of Vision*, *21*(7), 5-5.
Takahashi, N., Oertner, T. G., Hegemann, P., & Larkum, M. E. (2016). Active cortical dendrites modulate perception. *Science*, *354*(6319), 1587-1590.
Takahashi, N., Ebner, C., Sigl-Glöckner, J., Moberg, S., Nierwetberg, S., & Larkum, M. E. (2020). Active dendritic currents gate descending cortical outputs in perception. *Nature Neuroscience*, *23*(10), 1277-1285.
Tzilivaki, A., Kastellakis, G., & Poirazi, P. (2019). Challenging the point neuron dogma: FS basket cells as 2-stage nonlinear integrators. *Nature Communications,* 10(1).
Uludağ, K., & Blinder, P. (2018). Linking brain vascular physiology to hemodynamic response in ultra-high field MRI. *Neuroimage*, *168*, 279-295.
Urbanczik, R., & Senn, W. (2014). Learning by the dendritic prediction of somatic spiking. Neuron, 81(3), 521-528.
van Bergen, R.S. and Kriegeskorte, N., 2020. Going in circles is the way forward: the role of recurrence in visual inference. *Current Opinion in Neurobiology*, *65*, pp.176-193.
Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, *30*.
Vezoli, J., Magrou, L., Goebel, R., Wang, X. J., Knoblauch, K., Vinck, M., & Kennedy, H. (2021). Cortical hierarchy, dual counterstream architecture and the importance of top-down generative networks. *Neuroimage*, *225*, 117479.
Ward, M., & Rhodes, O. (2022). Beyond LIF Neurons on Neuromorphic Hardware. *Frontiers in Neuroscience*, *16*, 881598.
Wu, X., Liu, X., Li, W., & Wu, Q. (2018). Improved expressivity through dendritic neural networks. *Advances in neural information processing systems*, *31*.
Wybo, W. A., Tsai, M. C., Tran, V. A. K., Illing, B., Jordan, J., Morrison, A., & Senn, W. (2023). NMDA-driven dendritic modulation enables multitask representation learning in hierarchical sensory processing pathways. *Proceedings of the National Academy of Sciences*, *120*(32), e2300558120.

Xu, N. L., Harnett, M. T., Williams, S. R., Huber, D., O'Connor, D. H., Svoboda, K., & Magee, J. C. (2012). Nonlinear dendritic integration of sensory and motor input during an active sensing task. *Nature*, *492*(7428), 247-251.

Yang, S., Deng, B., Wang, J., Li, H., Lu, M., Che, Y., ... & Loparo, K. A. (2019). Scalable digital neuromorphic architecture for large-scale biophysically meaningful neural network with multi-compartment neurons. *IEEE transactions on neural networks and learning systems*, *31*(1), 148-162.

Zeman, A., Milton, F., Della Sala, S., Dewar, M., Frayling, T., Gaddum, J., ... & Winlove, C. (2020). Phantasia–the psychological significance of lifelong visual imagery vividness extremes. *Cortex*, *130*, 426-440.

Zhang, S., Xu, M., Kamigaki, T., Hoang Do, J. P., Chang, W. C., Jenvay, S., ... & Dan, Y. (2014). Long-range and local circuits for top-down modulation of visual cortex processing. *Science*, *345*(6197), 660-665.