Neuron

Disinhibition by VIP interneurons is orthogonal to cross-modal attentional modulation in primary visual cortex

Highlights

- Attention and VIP interneurons both strongly modulate sensory responses in V1
- Attentional and VIP modulations do not interact and are orthogonal to each other
- Imaging four cell classes shows the two modulations act through distinct mechanisms
- Circuit modeling confirms that V1 can multiplex strong yet non-interacting modulations

Authors

Dylan Myers-Joseph, Katharina A. Wilmes, Marian Fernandez-Otero, Claudia Clopath, Adil G. Khan

Correspondence

khan.adil@kcl.ac.uk

In brief

Attentional modulation of sensory processing is a critical aspect of cognition, but its neural circuit basis is poorly understood. Previous work implicates VIP interneurons as critical to this modulation. Myers-Joseph et al. demonstrate that modulations by VIP interneurons and attention are in fact orthogonal, allowing multiplexing diverse signals in V1.



Neuron

Article



Disinhibition by VIP interneurons is orthogonal to cross-modal attentional modulation in primary visual cortex

Dylan Myers-Joseph,¹ Katharina A. Wilmes,² Marian Fernandez-Otero,¹ Claudia Clopath,³ and Adil G. Khan^{1,4,*} ¹Centre for Developmental Neurobiology, King's College London, London SE1 1UL, UK

²Department of Physiology, University of Bern, Bern 3012, Switzerland

³Department of Bioengineering, Imperial College, London SW7 2AZ, UK

⁴Lead contact

*Correspondence: khan.adil@kcl.ac.uk https://doi.org/10.1016/j.neuron.2023.11.006

SUMMARY

Attentional modulation of sensory processing is a key feature of cognition; however, its neural circuit basis is poorly understood. A candidate mechanism is the disinhibition of pyramidal cells through vasoactive intestinal peptide (VIP) and somatostatin (SOM)-positive interneurons. However, the interaction of attentional modulation and VIP-SOM disinhibition has never been directly tested. We used all-optical methods to bi-directionally manipulate VIP interneuron activity as mice performed a cross-modal attention-switching task. We measured the activities of VIP, SOM, and parvalbumin (PV)-positive interneurons and pyramidal neurons identified in the same tissue and found that although activity in all cell classes was modulated by both attention and VIP manipulation, their effects were orthogonal. Attention and VIP-SOM disinhibition relied on distinct patterns of changes in activity and reorganization of interactions between inhibitory and excitatory cells. Circuit modeling revealed a precise network architecture consistent with multiplexing strong yet non-interacting modulations in the same neural population.

INTRODUCTION

Attention exerts a powerful influence on how cortical circuits process sensory information. There are multiple forms of visual attention, including spatial, feature-based, and cross-modality, and in all these cases, attention can modulate visual processing, leading to improved behavioral performance.^{1–5} Typically, stimulus selectivity in visual cortex increases with attention^{2,5–8} through increases^{7–9} or decreases^{10,11} in evoked firing rates or combinations of these.⁵ Changes in noise correlations, variability, and synchrony between neurons may also play a role in improving the discriminability of stimuli by downstream circuits.^{12–14} However, despite this thorough phenomenological description of the ways in which attention influences visual processing, very little is known about the circuit basis of attentional modulation in visual cortex.

Cortical activity is processed by interconnected networks of cells containing multiple classes of excitatory and GABAergic inhibitory interneurons with distinct molecular, cellular, and connectivity properties.^{15–19} Vasoactive intestinal peptide (VIP) interneurons are key regulators of cortical function, most prominently due to their strong inhibition of somatostatin (SOM) interneurons, which results in the disinhibition of excitatory pyramidal (PYR) cells.^{17,20,21} VIP-SOM-driven disinhibition provides useful computational opportunities.^{22,23} On a longer timescale, VIP-

SOM disinhibition gates the plasticity of inputs onto excitatory cells.^{24–27} Better studied, however, is the immediate effect of increasing activity in VIP interneurons, which leads to enhanced firing of local PYR cells.^{20,21,23,28–37}

Disinhibitory modulation of sensory cortex through VIP and SOM interneurons has in fact been proposed to be a core mechanism underlying attentional modulation in visual cortex.^{38,39} In this account, projections from frontal cortex activate VIP interneurons in V1, which leads to enhanced stimulus-evoked activation of PYR cells through VIP-SOM disinhibition and may lead to subsequent improvement in visual behavior.⁴⁰ Attentional modulation in V1 and visual discrimination behavior may also depend on cholinergic inputs to visual cortex, 41-43 and cortical VIP interneurons are activated by acetylcholine,^{44,45} raising the possibility of VIP interneurons being a key player in cholinergic visual gain enhancement during attention.⁴⁶ This account has parallels with the suggestion that locomotion elevates the gain of visual responses in V1 through VIP-SOM disinhibition.^{1,29,47} However, the interaction of attentional modulation and VIP modulation has never been directly tested.

If attention acts through VIP-SOM disinhibition, then three broad predictions exist. First, there should be an interaction between attentional modulation and VIP-SOM disinhibition. Second, silencing VIP interneurons should disrupt any attentional modulation. Third, there should be similarities in how attention



Figure 1. Modulation of stimulus selectivity in V1 neurons during a cross-modal attention-switching task (A) Schematic of the experimental apparatus.

(B) Schematic of the behavioral task. See also Figure S1.

(C) Behavioral discrimination performance (behavioral d') across attention (visual block, attend visual, 3.47 ± 0.77 median ± IQR here and below; olfactory block, ignore visual, 0.55 ± 0.87 , Wilcoxon signed-rank test p = 7.55×10^{-10} ; behavioral d' of olfactory discrimination, 3.90 ± 0.67 ; n = 50 sessions, 15 mice). Connected points indicate visual discrimination, individual points in the odor block represent olfactory discrimination. Gray lines and points are individual sessions, colored lines show the average of all sessions, and error bars indicate SD.

(D) Average responses from 2 example cells to the rewarded and unrewarded visual gratings in the odor and visual blocks, showing an increase in selectivity in the attend condition. Time zero marks the onset of the visual stimuli. Numbers indicate stimulus selectivity, shading indicates SEM.

(E) Stimulus selectivity of all non-VIP cells (calculated with stimulus-evoked responses averaged 0–1 s from stimulus onset) when attending or ignoring the same stimuli from an example session. Cells that significantly increase their selectivity with attention are highlighted in dark blue and purple for negatively and positively selective cells, respectively.

Neuron Article



and VIP-SOM disinhibition modify the activity of different cell classes and the interactions between them.

To test these predictions, we used an all-optical approach in V1, where we quasi-simultaneously imaged the activities of PYR, parvalbumin (PV), SOM, and VIP interneurons while optogenetically manipulating VIP interneurons as mice performed a cross-modal attention-switching task. We observed robust attentional modulation of stimulus selectivity in V1. We also observed a strong enhancement of activity when photoactivating VIP interneurons and a reduction in activity when photoinhibiting VIP interneurons. However, the two modulations did not interact and instead were orthogonal. The changes induced by these two modulations occurred through distinct mechanisms: although VIP photoactivation led to predominantly enhanced activity in VIP, PV, and PYR neurons and largely suppressed activity in SOM neurons, attention led to heterogeneous changes. Circuit modeling revealed that only a specific network architecture can account for the experimental findings. These results demonstrate that VIP-SOM disinhibition does not underlie the attentional modulation of selectivity in V1. At the same time, they reveal a remarkably versatile cortical circuit, which allows multiplexing of multiple non-interacting signals on the same neural populations.

RESULTS

Modulation of stimulus selectivity by attention in V1

To study the neural circuit basis of attentional modulation, we trained mice to perform a cross-modality attention-switching task (Figures 1A and 1B). Head-fixed mice switched between blocks of visual discrimination and olfactory discrimination, in which they licked a reward spout to obtain a reward in response to one of two visual grating stimuli or odors. During the olfactory discrimination blocks, the same grating stimuli used in the visual discrimination blocks were presented in 70% of the trials but were task irrelevant (Figure 1B; see Figure S1 for trial schematics). Mice attended to and accurately discriminated the grating stimuli in the visual block but ignored the same stimuli while successfully discriminating odors during the olfactory blocks (Figure 1C). Previous work using a similar paradigm^{5,48} established that this task requires V1.

We measured responses of layer 2/3 neurons expressing GCaMP7f⁴⁹ in V1 using two-photon calcium imaging during the task. We compared the responses to the same pair of visual stimuli in the attend and ignore conditions and observed a robust modulation of stimulus responses with attention (Figure 1D). These response changes modified stimulus selectivity (difference in the responses to the rewarded and unrewarded stimuli, normalized by the pooled standard deviation) such that cells preferring both the rewarded and unrewarded stimuli showed

increased stimulus selectivity (Figures 1D-1F). Across the population, the distribution of stimulus selectivity with attention broadened due to more positive and more negative values, indicating a higher preference for the rewarded and non-rewarded stimuli, respectively (Figure 1G). Consequently, the absolute stimulus selectivity of the neural population increased significantly with attention, as reported previously^{5,48} (Figure 1H). The increase in stimulus selectivity could not be accounted for by changes in running or licking behavior and remained after removing cells most influenced by running and licking (Figure S2). Furthermore, cells with preferred responses to the non-rewarded visual stimulus 2, which is not associated with reward in either block type, had a comparable increase in selectivity with cells preferring visual stimulus 1 (Figures 1D-1G). We restricted the analysis to cells that significantly increased their selectivity for either stimulus with attention and found that a combination of increased responses to the preferred and decreased responses to the non-preferred stimuli led to the increase in average selectivity (Figure 1I).

These results established that visual stimulus selectivity in V1 was strongly modulated in this attention-switching task. A candidate mechanism for this selectivity modulation is VIP interneuron-mediated disinhibition.^{38,39} To test this possibility, we first asked whether optogenetic VIP activation modified neural activity in the same mice passively viewing stimuli.

VIP activation strongly modulates cortical responses

To study the effect of activating VIP interneurons on visual stimulus responses, we expressed the red-shifted excitatory opsin Chrimson selectively in VIP interneurons and GCaMP7f nonselectively in V1 neurons using adeno-associated viruses (AAVs) (Figure 2A). We used an all-optical approach where we photoactivated VIP interneurons while measuring the activity of VIP and non-VIP cells in the same local circuit (Figure 2A, the population of all non-VIP cells is dominated by PYR neurons¹⁸; see also Figure 7 for the analysis of identified cell classes).

VIP interneurons were robustly activated by increasing light power *in vivo* (Figure 2C). For each imaged site, we conducted a calibration session (Figure 2D) and chose a "low" and "high" light power based on the shape of this curve (see STAR Methods). We established that our photoactivation was within a physiological range by comparing the response amplitude of VIP interneurons during photoactivation and spontaneous running bouts. Both light powers evoked activity in VIP interneurons that spanned the activity range naturally observed during locomotion (Figure 2E); mean VIP interneuron activities with low and high light power were 1.32 and 1.77 Δ F/F, respectively. Mean 75th and 95th percentiles of activity during locomotion without photostimulation for all VIP interneurons were 1.23 and 2.06 Δ F/F, respectively.

⁽F) Stimulus selectivity of the same cells in the attend and ignore conditions (columns). Cells were ordered by their mean selectivity across both contexts (n = 6,153 neurons, 15 mice).

⁽G) Histograms of stimulus selectivity when ignoring and attending the visual stimuli (n = 6,153 neurons, 15 mice).

⁽H) Boxplots of absolute stimulus selectivity during the ignore and attend conditions (ignore, 0.35 ± 0.10 , attend 0.61 ± 0.28 , significant fixed effect for attention, $p = 1.83 \times 10^{-14}$, linear mixed-effects model, n = 50 sessions, 15 mice).

⁽I) Average baseline subtracted responses to the preferred and non-preferred visual stimuli in the ignore and attend conditions for cells that significantly increased their stimulus selectivity with attention (n = 50 sessions, 15 mice, error bars indicate SEM).



Figure 2. VIP activation leads to multiplicative increase in activity of non-VIP cells

(A) Schematic of simultaneous imaging and optogenetic stimulation.

(B) Example region of an *in vivo* imaged plane showing all neurons expressing GCaMP7f and a VIP interneuron (arrowhead) additionally expressing Chrimson-tdTomato. (C) Mean responses of an example VIP interneuron to different light powers. Red bar indicates optogenetic stimulation duration (1.5 s), shading indicates SEM. (D) Boxplots of optogenetically evoked activity (mean 0-1 s, baseline subtracted) across different light powers for VIP interneurons (left, n = 91 cells, 6 mice) and non-VIP cells (right, n = 2,233 cells, 6 mice). Inset: schematic of VIP interneuron activation during passive gray screen viewing.

(E) Activity of an example VIP interneuron (black) and running speed (gray), dashed lines indicate mean optogenetically evoked activity in VIP interneurons (orange, low power, red, high power).

(F) Left, stimulus-evoked normalized activity in response to different oriented drifting gratings, averaged across all VIP interneurons (n = 141 cells) aligned to their preferred direction, shading indicates SEM here and below. ***p < 0.001 Wilcoxon signed-rank test for photoactivation compared with non-photoactivation conditions at each direction, corrected for multiple comparisons. Right, the same data shown as average activity with and without optogenetic light. Linear regression, low power, slope = 1.823, intercept = 0.327. High power, slope = 2.081, intercept = 0.52, 8 mice.

(G) Same as (F) for all orientation-selective non-VIP neurons, n = 1,044 cells. Low power, slope = 1.568, intercept = 0.009. High power, slope = 1.634, intercept = 0.032, 8 mice.

(H and I) Same as (F) and (G) for control mice expressing no opsin, n.s. indicates non-significant, n = 61 VIP interneurons, and 434 non-VIP cells, 3 mice.

Neuron Article

We photoactivated VIP interneurons while we passively presented a range of oriented, drifting grating stimuli to the mice. We confirmed that VIP interneurons increased their activity at each grating direction with increasing light power (Figure 2F). As a result, non-VIP cell stimulus responses were strongly enhanced in each direction (Figure 2G). This led to a largely multiplicative enhancement in orientation-tuned non-VIP cells (Figure 2G, right). Control experiments on mice without opsin expression, using identical light exposure, confirmed no activity changes, excluding any light-induced artifacts (Figures 2H and 21). Because a multiplicative increase in activity is unlikely to change stimulus selectivity,9 VIP interneurons may therefore not be involved in the stimulus selectivity modulations we observed with attention. However, directly testing this hypothesis required studying VIP activity during the attention-switching task.

Testing the interaction between VIP activation and cross-modal attention

To understand the role of VIP interneurons in attentional modulation, we first asked to what extent the activity of VIP interneurons themselves was modulated in our attention-switching task. We found a small but significant increase in pre-stimulus VIP interneuron activity with attention (Figure S3A). However, this could be accounted for by a difference in pre-stimulus running speed in the two blocks (Figure S3B), and correcting for pre-stimulus activity resulted in no significant difference in stimulus-evoked VIP activity (Figure S3C). This suggested that VIP interneurons may not be important in producing attentional modulation.

However, to conclusively determine whether VIP activity underlies attentional modulation, we photoactivated VIP interneurons in randomly interleaved trials (Figure 3A). This unilateral photoactivation had no effect on the behavior itself (Figures S4A and S4B), ruling out indirect effects through behavioral changes. VIP interneuron activity during the task was again enhanced with increasing light power, both when the visual stimuli were attended to and ignored (Figure 3B, top), leading to an increase in non-VIP cell stimulus-evoked activity (Figure 3B, bottom).

We next selected all non-VIP neurons that showed a significant increase in stimulus selectivity with attention. The stimulus responses of this population showed modulatory effects of both attention and VIP activation (Figure 3C, cells with preference for the unrewarded stimulus, similar effects seen for cells with a preference for the rewarded stimulus, data not shown, individual example cells shown in Figure S5). Increasing light power evenly increased the average activity of this population in both ignore and attend conditions for both stimuli (Figure 3D). We fit a linear mixed-effects model to predict the mean responses of negatively selective neurons and found a significant fixed effect for VIP activation (p = 1.75×10^{-10}), stimulus (p = 1.65×10^{-21}), and the interaction between stimulus and attention (p = 4.90 \times 10⁻⁶; attention results in an increase in preferred and decrease in non-preferred stimulus responses as in Figure 11; hence, this significant interaction term between stimulus and attention) but no significant interaction effect between attention and VIP activation (p = 0.92). Similar results of no interaction between attention and VIP activation were obtained for cells with a



preference for the rewarded stimulus (data not shown) and in control mice (Figure 3E; see also Tables S1–S6). The absence of an interaction effect between attention and VIP activation meant that attention led to similar changes in V1 responses, regardless of VIP activity levels.

Having found no interaction between attention and VIP activation on stimulus-evoked responses, we asked whether VIP activation led to any change in selectivity (Figure 3F). We fit a linear mixed-effects model on the data in Figure 3F (all non-VIP cells that significantly increased their selectivity with attention) to predict selectivity and found a significant fixed effect for attention $(p = 2.19 \times 10^{-17})$, but not for VIP activation (p = 0.095), and no interaction effect (p = 0.190, Table S2). The same result was obtained when taking all non-VIP cells (attention, p = 7.07 × 10^{-7} ; VIP activation, p = 0.0.845; and interaction, p = 0.847). Crucially, control mice displayed only a significant effect on selectivity from attention (p = 4.641×10^{-9}), but not from light delivery (p = 0.929), and no interaction effect (p = 0.753, Figure 3G). Thus, attention and VIP-driven disinhibition both induce robust modulation of the same neural population, but the two effects do not interact.

VIP inhibition leads to a modest suppression of cortical responses during passive viewing

Although the above results demonstrated that VIP activation did not interact with attentional modulation of stimulus selectivity, to better understand the role of VIP interneurons in shaping cortical stimulus-evoked responses, we next studied the effect of inhibiting VIP interneurons. We used a similar all-optical approach by expressing the inhibitory opsin ArchT in VIP interneurons (Figure 4A). VIP interneurons were progressively inhibited by increasing light power *in vivo* (Figure 4B). As with VIP activation, we conducted a calibration session for each site (Figure 4C) and chose a single light power, typically 1.5 mW (see STAR Methods).

We photoinhibited VIP interneurons while passively presenting drifting visual grating stimuli (Figure 4D, left), which led to decreased VIP activity at each grating direction (Figure 4D). As a result, stimulus responses of orientation-selective non-VIP cells were moderately but significantly inhibited at specific directions, including at the peaks of the tuning curve (Figure 4E). To confirm that we had inhibited VIP interneurons to minimal activity levels, we tested another group of mice in which we presented low-contrast gratings that activate VIP interneurons more strongly than high-contrast stimuli.⁵⁰ VIP interneuron activities were reduced to the same level whether they were highly active or less active (Figure 4F). VIP inhibition in both cases was accompanied by reductions in non-VIP cell activity (Figure 4G). Together, this confirmed that our photoinhibition was appropriate for reducing VIP activity to near the minimum physiological levels. Thus, VIP inhibition during passive viewing of stimuli led to a modest inhibition of non-VIP cell activity.

VIP inactivation during cross-modal attention-switching

To conclusively rule out the role of VIP interneurons in attentional modulation of stimulus selectivity, we photoinhibited VIP interneurons in randomly interleaved trials while mice performed the attention-switching task (Figure 5A). As with photoactivation,



Figure 3. No interaction between VIP modulation and cross-modal attentional modulation

(A) Schematic showing VIP photoactivation during the attention-switching task. Light onset (red bars) was from -0.1 to 1.5 s relative to visual stimulus onset. Light was ramped off over 0.2 s.

(B) Mean visual stimulus-evoked activity with increasing VIP photoactivation. Top, all VIP interneurons, bottom, all non-VIP cells. Left, responses when ignoring the visual stimuli, right, responses when attending the visual stimuli. Wilcoxon signed-rank test between photoactivation and non-photoactivation conditions, ***p < 0.001, n = 17 sessions, 7 mice. Gray lines indicate individual session averages and colored lines indicate overall average.

(C) Top, mean visual stimulus-evoked activity for all non-VIP cells with preference for the unrewarded stimulus that significantly increased their selectivity with attention (mean of n = 17 session averages, shading indicates SEM). Bottom, same sessions, responses with additional VIP photoactivation (red). Responses from top are superimposed for comparison (gray dashed lines). Light red shading indicates light onset.

(D) Black, mean visual stimulus-evoked activity (averaged 0–1 s) of all non-VIP cells with preference for the unrewarded stimulus that significantly increased their selectivity with attention, n = 17 sessions. Orange and red, same responses with additional VIP photoactivation. Error bars indicate SEM. (E) Same as (D) for control mice, n = 17 sessions, 3 mice.

(F) Boxplots of absolute stimulus selectivity with increasing VIP photoactivation for all non-VIP cells which significantly increased their selectivity with attention (n = 17 sessions). Stimulus selectivity measured when ignoring the visual stimuli (left) and attending the same stimuli (right).

(G) Same as (F) for control mice, n = 17 sessions, 3 mice. See also Tables S1, S2, S5, and S6.



Figure 4. VIP inhibition moderately suppresses responses to visual stimuli during passive presentation

(A) Example region of an *in vivo* imaged plane showing all neurons expressing GcaMP7f and a VIP interneuron (arrowhead) additionally expressing ArchT-tdTomato.

(B) Mean responses of an example VIP interneuron to different light powers. Responses are aligned to optogenetic light onset (dashed line). Red bar indicates optogenetic stimulation duration (1.5 s), shading indicates SEM.

(C) Boxplots of optogenetically inhibited activity (mean 0–1.5 s, baseline subtracted) across different light powers for VIP interneurons (left, n = 92 cells) and non-VIP cells (right, n = 1,648 cells, 5 mice).

(D) Left, stimulus-evoked normalized activity in response to different oriented gratings, averaged across all VIP interneurons (n = 122 cells, 6 mice) aligned to their preferred direction, shading indicates SEM here and below. *p < 0.05, **p < 0.01, ***p < 0.001 Wilcoxon signed-rank test for photoinhibition compared with non-photoinhibition conditions at each direction, corrected for multiple comparisons. Right, the same data shown as average activity with and without optogenetic light. Linear regression, slope = 0.748, intercept = -0.012.

(E) Same as (D) for all orientation-selective non-VIP neurons, n = 953 cells, slope = 0.965, intercept = -0.005. High power, slope = 1.634, intercept = 0.032, 8 mice. (F) Visual stimulus-evoked VIP interneuron activity (mean 0–1.5 s, baseline subtracted) in response to a drifting vertical grating at low and high contrast, with and without VIP photoinhibition, n = 37 cells, 5 mice, error bars indicate SEM.

(G) Same as (F) for non-VIP cells, n = 528 cells, 5 mice.

photoinhibition had no effect on mouse behavior (Figures S4C and S4D). VIP interneurons were inhibited during the behavior (Figure 5B, top). However, we observed no effect of VIP inhibition on stimulus responses of non-VIP neurons (Figure 5B, bottom). To confirm that there was no deficit in the degree of our photo-inhibition, we applied an average of 7.5 mW of light power (range

5.7–9 mW), higher than the 1.5 mW used so far. We found no further reduction in VIP interneuron activity at this higher light power (Figure S6A). Furthermore, we found that the higher light power led to a significant increase in measured neural activity in the control mice (Figure S6B), possibly induced through direct retinal activation and subtle changes in behavior linked to light onset







Locomotion, no light Locomotion, VIPs inhibited

(legend on next page)

Locomotion, no light - Locomotion, VIPs activated

Neuron Article



(Figures S6C and S6D). These results highlight the importance of light-only controls in animals performing the full behavioral task and confirm that our photoinhibition of VIP interneurons at 1.5 mW was close to complete.

To study the interaction of VIP inhibition with attention, we selected all non-VIP neurons that showed a significant increase in stimulus selectivity with attention. The stimulus responses of this population showed at most a subtle effect of VIP inhibition (Figure 5C, average of all cells with preference for unrewarded stimulus). We fit a linear mixed-effects model to predict the mean responses of neurons with a preference for the unrewarded stimulus and found a significant fixed effect for attention (p = 0.027), VIP inhibition (p = 0.024), stimulus ($p = 1.79 \times 10^{-28}$), and interaction between stimulus and attention ($p = 4.35 \times 10^{-6}$) but no significant interaction effect between VIP inhibition was also found in cells with a preference for the rewarded stimulus (Figure 5E; Table S3).

Having found no interaction between attention and VIP inhibition on non-VIP cell stimulus-evoked activity, we asked whether VIP inhibition led to any change in stimulus selectivity (Figure 5F). We fit a linear mixed-effects model on the data in Figure 5F (all non-VIP cells that significantly increased their selectivity with attention) and found a significant main effect on selectivity from attention ($p = 1.77 \times 10^{-11}$), but not from VIP inhibition (p = 0.56), and no interaction effect (p = 0.37, Table S4). The same result was obtained when taking all non-VIP cells (a significant main effect on selectivity from attention, $p = 3.66 \times 10^{-7}$), but not from VIP activation (p = 0.44), and no interaction effect (p = 0.47). Thus, VIP interneurons are not involved in producing the gain in stimulus selectivity during attention.

To confirm that our VIP manipulations were effective, we tested whether they did interact with another top-down modulation involving VIP interneurons. Locomotion increases the gain of visual stimulus responses in V1⁵¹ in a VIP interneuron-dependent manner.²⁹ We activated and inhibited VIP interneurons while presenting drifting oriented gratings and divided the data into locomotion and stationary epochs. We calculated the

locomotion modulation index (LMI, see STAR Methods) at the preferred stimulus for each neuron. Without VIP manipulations, there was a significant locomotion-induced gain increase in visual responses in non-VIP cells for mice expressing the excitatory (Figure 5G) and inhibitory opsin in VIP cells (Figure 5H). We photoactivated VIP interneurons on randomly interleaved trials and found that the LMI of non-VIP cells was significantly reduced on VIP photoactivation (Figure 5G). Similarly, we found that the LMI of non-VIP cells was significantly reduced on VIP photoinhibition (Figure 5H). These experiments act as a positive control, demonstrating the effectiveness of our VIP manipulations and strengthening our conclusion regarding the absence of a role of VIP interneurons in attentional modulation.

Overall, although VIP interneurons are capable of strongly modulating the activity of the non-VIP neural population, they are not the route through which cross-modal attention induces selectivity changes in V1. Any scenarios in which specific patterns of VIP cell activation may be required to induce the attentional modulation are ruled out by the results from the VIP inhibition experiments, because strong VIP inhibition left the attentional modulation of both stimulus responses and selectivity unperturbed.

Cross-modal attention and VIP modulations are orthogonal

Can attention-driven and VIP-driven modulations co-occur in the same neurons and yet have no adverse impact on each other? To address this question, we asked whether these two modulations were orthogonal to each other. We performed dimensionality reduction using linear discriminant analysis (LDA) and found two axes best separating the visual stimulus-evoked neural activity of the non-VIP population: first between attend and ignore conditions and second between photoactivation and no photoactivation conditions (VIP activation, VIP inhibition, and control) within the ignore condition (Figure 6A). We found the cosine similarity of these two axes as a measure of the alignment of the attentional and optogenetic manipulations.

Figure 5. No effect of VIP inhibition on cross-modal attentional modulation

(A) Schematic showing VIP photoinhibition during the attention-switching task. Light onset (red bars).

⁽B) Mean visual stimulus-evoked activity (baseline subtracted) with increasing VIP photoinhibition. Top, all VIP interneurons; bottom, all non-VIP cells. Left, responses when ignoring the visual stimuli; right, responses when attending the visual stimuli. Wilcoxon signed-rank test between photoactivation and non-photoactivation conditions, ***p < 0.001, n = 16 sessions, 5 mice. Gray lines indicate individual session averages, colored lines indicate overall average.

⁽C) Top, mean visual stimulus-evoked activity for all non-VIP cells with preference for the unrewarded stimulus that significantly increased their selectivity with attention (mean of n = 16 sessions, shading indicates SEM). Bottom, same sessions, responses with additional VIP photoinhibition (red). Responses from top are superimposed for comparison (gray dashed lines, light red shading indicates light onset).

⁽D) Black, mean visual stimulus-evoked activity (averaged 0–1 s, baseline subtracted) of all non-VIP cells with preference for the unrewarded stimulus that significantly increased their selectivity with attention, n = 16 sessions. Red, same responses with additional VIP photoinhibition. Error bars indicate SEM. (E) Same as (D) for cells with preference for the rewarded stimulus.

⁽F) Boxplots of absolute stimulus selectivity without and with VIP photoinhibition for all non-VIP cells that significantly increased their selectivity with attention (n = 16 sessions). Stimulus selectivity measured when ignoring the visual stimuli (left) and attending the same stimuli (right).

⁽G) Left, visual stimulus onset aligned mean responses during a passive viewing session from an example non-VIP cell in a mouse expressing the excitatory opsin Chrimson in VIP cells. Traces show average activity during visual stimulus presentation with and without photoexcitation of VIP cells, shading indicates SEM. Red shaded area indicates optogenetic light. Right, cumulative density curves of the locomotion modulation index (LMI) of all non-VIP cells with and without optogenetic light. LMI was significantly greater than 0 (Wilcoxon signed-rank test $p = 3.27 \times 10^{-247}$) and LMI with VIP photoexcitation was significantly smaller than without light (Wilcoxon signed-rank test $p = 9.69 \times 10^{-30}$, n = 1,513 cells, 4 mice).

⁽H) Same as (G) for mice expressing the inhibitory opsin ArchT in VIP cells. Left, example non-VIP cell with and without optogenetic light. Right, LMI was significantly greater than 0 ($p = 8.08 \times 10^{-287}$) and the LMI with VIP photoinhibition was significantly smaller than without light ($p = 3.65 \times 10^{-7}$, n = 1,760 cells, 4 mice). See also Tables S3 and S4.







Figure 6. Cross-modal attention and VIP modulations are orthogonal

(A) Mean stimulus-evoked activity from individual trials projected onto the first two axes obtained from dimensionality reduction using linear discriminant analysis (LDA). Example session from a mouse expressing Chrimson. Dashed lines indicate the directions along which attention and VIP photoactivation most strongly modulated activity.

(B) Mean cosine similarity for pairs of axes in neural activity space (Chrimson n = 17 sessions, ArchT n = 16 sessions, control n = 17 sessions, error bars indicate SEM). From left to right, absolute cosine similarity for: LDA axis separating VIP photoactivation vs. no photoactivation trials and LDA axis separating attend vs. ignore trials; pairs of random axes extracted based on the covariance of the original neural data (10,000 samples); two axes separating the rewarded vs. unrewarded visual stimuli where each axis was found using half of the data (mean of 50 shuffled repeats). Significance tests were against the corresponding random axes median using Wilcoxon signed-rank test, **p < 0.01, ***p < 0.001. From left to right p values for all tests: 0.009, 0.379, 0.653, 6.43 × 10⁻⁴, 2.93 × 10⁻⁴, 0.002.

We created a null distribution of cosine similarity between randomly chosen pairs of axes from the dataset based on the covariance of the original data (10,000 samples, same method as Elsayed et al.⁵²). We compared the cosine similarity of the LDA axes of modulation by attention and modulation by VIP manipulation for each session to its random mean. In the VIP activation experiment, the attention and VIP modulation axes were significantly less aligned than random axes drawn from the same space occupied by the data (Figure 6B). This demonstrates that attention and VIP disinhibition produce modulations that are significantly closer to orthogonality than expected by chance.⁵²

The VIP inhibition and control experiments showed no difference in cosine similarity from random vectors (Figure 6B), consistent with the absence of light-driven effects in these two groups. As a positive control, we confirmed that in all three groups, the vectors separating the two visual stimuli responses from two halves of the data had significantly higher cosine similarity than random vectors (Figure 6B). These results demonstrate that attention and VIP modulation not only interact but are also orthogonal to each other.

Distinct mechanisms of changes during cross-modal attention and VIP modulation

We next wished to understand the circuit mechanisms underlying these two modulations and establish how similar or distinct they were. To achieve this, we reidentified the same neurons in co-registered, immunohistochemically stained brain sections^{5,53} in 4 of the 8 same animals expressing Chrimson in VIP interneurons and detected simultaneously imaged PV-, SOM-, and VIP-positive interneurons (Figure 7A). The remaining cells were classified as putative PYR cells. VIP photoactivation during both passive viewing of a gray screen (Figures S7A–S7E) and oriented drifting grating stimuli (Figures S7F–S7J) led to an enhancement of VIP, PYR, and PV cell responses. Crucially, we saw concurrent suppression of average SOM cell activity (Figures 7B and 7G), consistent with VIP photoactivation leading to the disinhibition of PYR and PV cells via direct inhibition of SOM interneurons.

If attention and VIP modulations are indeed orthogonal, we would expect the underlying mechanisms of changes induced by either of them to be distinct. We tested this using different measures. First, as shown above, attention led to both increases and decreases in average stimulus-evoked activity (Figure 1I), but VIP activation only induced increases in stimulus-evoked activity (Figures 3B and 3D). Second, we measured noise correlations between the 4 simultaneously recorded cell classes. Noise correlations were measured as the stimulus-independent trialto-trial co-variability of responses and thus provided an estimate of mutual connectivity and shared inputs between and within cell classes. Attention significantly increased the noise correlation between VIP cell pairs and significantly decreased the correlation between PYR-SOM, SOM-VIP, and PV-SOM cell pairs (Figure 7E). In contrast, VIP activation significantly decreased noise correlation between PYR, VIP, and PV cell pairs and between PYR-VIP, PV-VIP, and SOM-VIP cell pairs (Figure 7F). p values



(legend on next page)





(uncorrected; unrewarded visual stimulus) for changes with attention are PYR-PYR: p = 0.320, VIP-VIP: p = 0.005, SOM-SOM: p = 0.966, PV-PV: p = 0.123, PV-VIP: p = 0.278, PYR-PV: p = 0.054, PYR-VIP: p = 0.123, PV-SOM: p = 0.014, PYR-SOM: p = 0.019, and SOM-VIP: p = 0.005 and for changes with VIP photoactivation are PYR-PYR: p = 0.010, VIP-VIP: p = 0.001, SOM-SOM: p = 0.240, PV-PV: p = 0.019, PV-VIP: p = 0.010, PYR-PV: p = 0.067, PYR-VIP: p = 0.032, PV-SOM: p = 0.083, PYR-SOM: p = 0.278, and SOM-VIP: p = 0.003.

Third, we studied how attention and VIP photoactivation modified stimulus-evoked activity across the populations of these 4 cell classes (Figure S8). By measuring the difference in stimulus-evoked responses between attend and ignore conditions for each neuron, we found that in the absence of photoactivation, attention led to heterogeneous changes (Figure 7G, unrewarded grating responses shown here to exclude reward-evoked responses). In contrast, VIP photoactivation led to predominantly enhanced activity in VIP, PV, and PYR neurons and largely suppressed activity in SOM neurons (Figure 7H).

Fourth, we compared the degree of selectivity modulation by attention and activity modulation by VIP activation on neurons from each cell class (Figure 7I). We found no strong correlation in any cell class, ruling out the possibility of a specific subset of cells driving both attentional and VIP-driven changes (Pearson's correlation coefficients -0.05, 0.03, 0.05, and -0.1 in VIP, SOM, PV, and PYR cells, respectively). In fact, we found a small but significant negative correlation in PYR (p = 3.4×10^{-6}) and VIP interneurons (p = 0.015, all other ps > 0.05), suggesting a moderate segregation of attention and VIP modulations across these cell populations.

Finally, we performed the analysis from Figure 6 for each cell class separately and found that orthogonality between attentional and VIP modulations of neural population responses was cell class specific. The orthogonality was significant for PYR and VIP populations but not for PV or SOM (Wilcoxon signedrank test, PYR p = 9.76×10^{-4} , PV p = 0.638, SOM p = 0.123, VIP p = 0.009, n = 11 sessions, 4 mice).

Overall, these results demonstrate that attention changes V1 stimulus processing by heterogeneous changes in activity and correlations across different cell classes, whereas VIP activation drives relatively homogeneous enhancement of PYR and PV cell activity along with SOM inhibition. Thus, attention and VIP disinhibition act on the same cortical circuit through distinct mechanisms.

A circuit model captures the effects of both attention and VIP perturbations

How can attention and VIP modulation each profoundly affect the same cell populations and yet not interact? To address this question, we developed a theoretical circuit model in which we represented the four cell types—PYR, PV, SOM, and VIP—by their population activity. The activity of each population was determined by baseline activity, bottom-up stimulus-related input, top-down attentional input, and connection strengths with other cell populations (Figure 8A). We used experimentally derived connectivity values from Pfeffer et al.¹⁷ and used simulation-based inference (SBI)⁵⁴ to determine connectivity values from PYR to all interneurons and from SOM to PV cells, and bottom-up, top-down, and baseline inputs that replicated a number of experimental findings, most importantly that of non-interacting attention and VIP modulations (see STAR Methods).

The model reproduced the suppression of SOM and enhancement of PYR and PV activities during VIP activation (Figure 8B). We next focused on PYR and PV cells because these cells showed a robust increase in stimulus selectivity (PYR average absolute selectivity, ignore 0.43 ± 0.77 [mean \pm SD], attend 0.49 ± 0.52 , Wilcoxon signed-rank test, p = 7.02×10^{-21} , n = 1,616; PV, ignore 0.29 ± 0.23 , attend 0.48 ± 0.34 , p = 8.53×10^{-5} , n = 67; SOM, ignore 0.35 ± 0.33 , attend 0.41 ± 0.34 , p = 0.08, n = 50; VIP, ignore 0.23 ± 0.17 , attend 0.41 ± 0.30 ,

Figure 7. Simultaneous VIP, SOM, PV, and pyramidal cell activity reveals distinct mechanisms of modulation with cross-modal attention and VIP photoactivation

(A) Example region of an *in vivo* image plane with GCaMP7f-expressing neurons (top) and the same region after post-hoc immunostaining for PV, SOM, and VIP (orange, blue, and magenta, respectively) following image registration (bottom). Identified interneurons are indicated by arrowheads.

(B) Difference (VIP photoactivation condition minus no photoactivation condition) of average visual stimulus-evoked response for each cell in all 4 cell classes (average of all orientations of visual stimuli during passive presentation) aligned to visual stimulus onset (dashed line). Optogenetic light onset here and below is -0.1 to 1.5 s from visual stimulus onset (red shading). Cells are sorted by their average response difference 0–1 s from stimulus onset. Left to right: VIP n = 85 cells, SOM n = 30 cells, PV n = 40 cells, PYR n = 1,567 cells.

(C) Mean of each column of (B) showing average change in activity with VIP activation; shading indicates SEM.

(D) Boxplots of visual stimulus-evoked activity with and without VIP photoactivation, averaged 0–1 s from visual stimulus onset, for each cell class (Wilcoxon signed-rank tests for differences in activity: PYR, n = 1,567 cells, p = 4.52×10^{-187} ; SOM, n = 30 cells, p = 5.71×10^{-4} ; VIP, n = 85 cells, p = 1.17×10^{-15} ; PV, n = 40 cells, p = 3.57×10^{-8}).

(E) Mean noise correlations between cell pairs belonging to the same or different cell classes, in the ignore and attend conditions, without photoactivation. Error bars represent SEM here and below (n = 11 sessions, 4 mice). Inset: changes in noise correlations due to attention as indicated by line thickness and color code. Shorter line segments indicate change in noise correlations between cells of the same type.

(F) Same as (E) for noise correlations with and without VIP photoactivation in the attend condition.

(G) Top, difference in mean visual stimulus-evoked response with attention (baseline subtracted, unrewarded grating), for each cell type, aligned to visual stimulus onset (dashed line). Cells are sorted by their average activity in the ignore condition (see also Figure S8). Bottom, average responses of cells from the top, middle and bottom 10th percentiles of the difference in responses (averaged 0–1 s) with attention shown above. Shaded area indicates SEM.

(H) Same as (G) but for differences in mean visual stimulus-evoked response with photoactivation of VIP interneurons (red bar and shading) compared with no photoactivation in the ignore condition. Top, cells are sorted the same as in (G), by their average activity in the ignore condition.

(I) Relationship between Δ Selectivity with attention (positive values indicate increased stimulus selectivity with attention) and change in stimulus-evoked activity with VIP photoactivation (mean 0–1 s, baseline subtracted), for VIP (n = 130 cells), SOM (n = 50 cells), PV (n = 67 cells), and PYR cells (n = 1,616 cells, 4 mice). Cells with values greater than the axes limits were pegged to the axes. Significant negative correlations were present only in VIP and PYR cells. See also Figures S7 and S8.



⁽legend on next page)



 $p = 7.89 \times 10^{-9}$, n = 130; Figure 8C; see also Poort et al.⁵). The model reproduced the increase in PYR and PV stimulus selectivity through a combination of increased and decreased responses to the preferred and non-preferred stimuli, respectively (Figure 8D). Crucially, VIP activation enhanced PYR activity almost equally in the ignore and attend conditions (Figure 8E), closely matching the experimental findings (Figure 3D). This led to small changes in selectivity with VIP activation but large selectivity changes with attention (Figure 8F; compare with Figure 3F). The model responses to VIP inhibition (Figure 8G) also closely followed experimental results (Figure 5D and 5E), even though this observation was not used to constrain the model. The model thus revealed a network architecture capable of sustaining strong yet non-interacting modulations by attention and VIP activation.

We found that PV cells received the strongest top-down inputs in our model (Figure 8H), which were tightly correlated with those arriving on PYR cells (Figure 8H), suggesting that the relative strength of top-down inputs on PYR and PV cells is highly regulated. We tested this prediction of the model, specifically that there is a higher degree of correlation between PYR and PV inputs compared with PYR and SOM in the attend condition. We found a significant trial-by-trial correlation of population activity between PYR and PV populations during the attend condition in 10/11 sessions, and furthermore, this correlation was higher than the correlation between PYR-SOM populations in 9/11 sessions. We performed bootstrapped significance testing between the correlation coefficients of PYR-PV and PYR-SOM populations and obtained a significantly higher correlation in the PYR-PV populations in 5/11 sessions (p < 0.05) compared with a significantly higher correlation in the PYR-SOM populations in 0/11 sessions. This result provided experimental validation of the circuit model.

To better understand how the model implemented the increase in selectivity with attention, we performed targeted manipulations in our model focused on two candidate circuit motifs, PYR-PV and VIP-SOM recurrently connected loops.¹⁷ On deleting the PV-to-PYR connection, we found a strong disruption of

PYR cell activity, resulting in the selectivity no longer being modulated by attention (Figure 8I, similar results were obtained on deleting PYR-to-PV connection). In contrast, deleting the VIP-to-SOM connection left the attentional modulation largely unperturbed (Figure 8J, similar results on deleting SOM-to-VIP connection). This suggests that top-down inputs to the PYR-PV recurrent circuit exploit distinct steady states of the network for preferred and non-preferred stimuli to implement the increase in selectivity with attention. Together, these results provide a mechanistic working model of attentional modulation for further exploration.

Neuron Article

Finally, we explored the origin of the experimental finding that VIP activation led to enhanced activity in PV cells (Figures 7B-7D). The enhancement of PV cell activity with VIP activation may be a result of disinhibition through SOM cells, because SOM cells inhibit PV cells,¹⁷ or a consequence of elevated PYR activity leading to increased PV cell activity. In our model, we first deleted the SOM-to-PV connection and compared the response of PV cells with and without VIP activation and found a minimal effect (Figure 8K). In contrast, when we deleted the connection from PYR to PV cells, we found strongly reduced PV responses that were not affected by VIP activation (Figure 8L). These results suggest that VIP activation suppresses SOM cells, enhancing PYR activity, which subsequently increases PV activity through PYR-to-PV connections. Overall, using this meanfield model provided the appropriate level of abstraction and experimental constraint, allowing us to comprehensively study the free parameter space and identify the circuit architecture consistent with our data.

DISCUSSION

In this study, we demonstrate that attention and VIP-SOM disinhibition both lead to strong modulations in V1. Critically, however, these two modulations do not interact and are in fact orthogonal to each other at the population level. Consistent with this conclusion, attention and VIP modulations are accompanied by distinct patterns of changes in activity and interactions

Figure 8. Circuit modeling reveals a network architecture for independent attentional and VIP modulations

(A) Schematic of the model architecture, indicating connectivity between different cell classes, bottom-up (stim) and top-down inputs (attention), and VIP photoactivation.

(B) Simulated responses of the 4 cell types to a visual stimulus with and without VIP photoactivation.

(C) Experimental results, average responses to the preferred and non-preferred visual stimuli in the ignore and attend conditions for positively selective cells that significantly increased their stimulus selectivity with attention, for PYR (left, n = 289 cells,) and PV cells (right, n = 15 cells), error bars indicate SEM. Inset: absolute selectivity of the full population of PYR (n = 1,616) and PV cells (n = 67), Wilcoxon signed-rank test, PYR $p = 7.02 \times 10^{-21}$, PV $p = 8.53 \times 10^{-5}$. (D) Model output, same as (C).

(E) Mean response of the model PYR population to visual stimuli with and without attention, and with no, low or high VIP photoactivation, for the preferred (left) and non-preferred stimuli (right). Compare with Figure 3D.

(F) Stimulus selectivity of the model PYR population with and without attention, and with no, low or high VIP photoactivation. Compare with Figure 3F. (G) Same as (E) for VIP photoinhibition. Compare with Figures 5D and 5E.

(H) Visualization of the posterior distribution over the three parameters for the top-down modulation of PYR, SOM, and PV that were consistent with the experimental observations and obtained by sampling 50,000 parameter sets from the estimated posterior distribution. Shown are the univariate marginals (histograms) and pairwise marginals (two-dimensional histograms). The y axis of each two-dimensional histogram corresponds to the values of the parameter on the left of it, and the x axis to the values of the parameter below.

(I) Left, model in which the PV to PYR connection was deleted (set to zero). Right, stimulus selectivity of the model PYR population in the ignore and attend condition, as in (F).

(J) Same as (I) with VIP to SOM connections deleted.

(K) Left, model in which the SOM to PV connection was deleted. Right, the model responses of the PV population, with and without VIP photoactivation.

(L) Same as (I) with PYR to PV connections deleted.



between excitatory and different inhibitory cell classes. Circuit modeling revealed a precise network architecture compatible with these experimental findings.

Circuit basis of attentional modulation

It is challenging to test hypotheses about the circuit basis of cognitive phenomena because this requires the measurement and manipulation of neural circuit components as animals perform a cognitive task.^{7,47,55–57} In this study, we advance this approach by simultaneously measuring the activity of four cell classes while manipulating the activity of one of these cell classes within a physiologically relevant range. As a result, we established that VIP cells are not the route through which attentional signals influence activity in V1.

An alternative explanation may be that a specific pattern of individual VIP cell activation needs to be engaged by attentional signals, and our non-specific full-field photoactivation approach is insufficient to rule out the role of VIP interneurons in attention modulation. However, this argument is refuted by our findings that VIP interneuron inhibition resulted in no changes to attentional modulations and that stimulus-evoked responses of VIP interneurons themselves were not modulated by attention.

Circuit modeling indicated that PV interneurons may instead be a critical target of attentional signals. In support of this idea, PV interneurons in V1 themselves show strong attentional modulations^{5,58} and may be capable of enhancing stimulus selectivity and discrimination behavior.⁵⁹ Manipulations of our model revealed that the increase in stimulus selectivity was achieved by top-down inputs acting through the recurrent PYR-PV circuit, utilizing distinct network steady states for preferred and non-preferred stimuli. Our results thus provide a mechanistic working model of top-down attentional modulation for further exploration.

This study used a cross-modal attention task, and our results do not rule out a role for VIP interneurons in other forms of attention, such as spatial or feature-based attention.^{38,39} Furthermore, several other circuit mechanisms for attention-related changes in sensory processing have been proposed, involving cortical and subcortical brain regions.^{60,61} Further investigations into the role of VIP interneurons in other forms of attention will be necessary to fully understand their contribution to attentional processing.

Role of VIP disinhibition

What is VIP-SOM disinhibition in V1 involved in, if not attention? We confirm that this circuit is involved in locomotion-induced gain modulation.²⁹ The effect of local visual context may also act through a VIP-SOM disinhibitory circuit.⁶² Crucially, VIP-SOM disinhibition may gate plasticity onto PYR cells.^{24,63,64}

The network effects of VIP activation were dominated by inhibited SOM activity and elevated activity in PYR and PV cells on average. However, we observed heterogeneity within cell populations, which is not surprising given that each molecularly defined class studied here comprises sub-types with distinct gene expression patterns, morphologies, network connectivity, and intrinsic properties.¹⁸



In general, caution is required in interpreting any results of circuit manipulations in densely interconnected and active networks.⁶⁵ In our study, however, several findings converge on the same conclusion: the absence of attentional modulation of VIP cell responses, the absence of interactions between attention and VIP activation effects, the absence of any impact of VIP inhibition on attention modulation, and the orthogonality of the two manipulations at the population level all support the conclusion that attention and VIP modulations are independent.

Orthogonal modulations

Orthogonal representations are important in segregating movement-related signals from stimulus-evoked activity in visual cortex.⁶⁶ In the primate motor cortex, during reaching movements, preparatory activity before the reach is actively maintained orthogonal to the activity during the reach, possibly supporting non-interfering computations.⁵² In the mouse somatosensory cortex, approximately orthogonal representations of whisker contacts may allow the flexible use of the same stimulus representations in distinct tasks.⁶⁷ Given our finding that attentional and VIP modulations are orthogonal, we speculate that various top-down signals may take advantage of this orthogonality to additionally engage the same V1 populations in other cognitive processes along with attention, such as working memory⁶⁸ or prediction error computation.⁶⁹

Overall, our results suggest the need to revise our understanding of the role of VIP-SOM-driven disinhibition in attentional modulation of stimulus selectivity. At the same time, we demonstrate the remarkable capacity of cortical circuits to combine multiple orthogonal and mechanistically distinct computations in the same neural populations.

STAR*METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
 - Lead contact
 - Materials availability
 - Data and code availability
- EXPERIMENTAL MODEL DETAILS
 - Mouse lines
- METHOD DETAILS
 - Animals and surgical procedures
 - Immunohistochemistry and ex vivo imaging
 - Two-photon calcium imaging
 - Behavioural training
 - Optogenetic manipulations
 - Direction tuning
 - Locomotion modulation
 - Optogenetic light during attention-switching
- QUANTIFICATION AND STATISTICAL ANALYSIS
 - Pre-processing
 - Behavioural controls
 - Selectivity





- Linear mixed-effects models
- Orthogonality
- Noise correlations
- Circuit model

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j. neuron.2023.11.006.

ACKNOWLEDGMENTS

We thank Florencia lacaruso, Takahiro Kanamori, Dimitar Kostadinov, and Petr Znamenskiy for discussions and comments on the manuscript and Matthew Grubb, Samuel Cooke, Juan Burrone, Attila Losonczy, and members of the Khan laboratory for discussions of the results in this manuscript. This work was supported by the Wellcome Trust (A.G.K., 206222/Z/17/Z), the BBSRC (A.G.K., BB/S015809/1), and start-up funds from the CDN, King's College London (A.G.K.). For the purposes of open access, the authors have applied a Creative Commons Attribution (CC-BY) license to any accepted author manuscript version arising from this submission.

AUTHOR CONTRIBUTIONS

D.M.-J. and A.G.K. designed the experiments. D.M.-J. performed the experiments and analyzed the data. M.F.-O. provided technical assistance and performed the immunolabeling and post-hoc cell matching. K.A.W. developed and analyzed the circuit model with inputs from C.C. A.G.K. and D.M.-J. wrote the paper with inputs from K.W. and C.C. All authors contributed to discussions and commented on the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: April 20, 2023 Revised: August 24, 2023 Accepted: November 8, 2023 Published: December 8, 2023

REFERENCES

- Speed, A., and Haider, B. (2021). Probing mechanisms of visual spatial attention in mice. Trends Neurosci. 44, 822–836.
- Maunsell, J.H.R. (2015). Neuronal mechanisms of visual attention. Annu. Rev. Vis. Sci. 1, 373–391.
- Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. Annu. Rev. Neurosci. 18, 193–222.
- Wang, L., and Krauzlis, R.J. (2018). Visual selective attention in mice. Curr. Biol. 28, 676–685.e4.
- Poort, J., Wilmes, K.A., Blot, A., Chadwick, A., Sahani, M., Clopath, C., Mrsic-Flogel, T.D., Hofer, S.B., and Khan, A.G. (2022). Learning and attention increase visual response selectivity through distinct mechanisms. Neuron 110, 686–697.e6.
- Reynolds, J.H., and Chelazzi, L. (2004). Attentional modulation of visual processing. Annu. Rev. Neurosci. 27, 611–647.
- Speed, A., Del Rosario, J., Mikail, N., and Haider, B. (2020). Spatial attention enhances network, cellular and subthreshold responses in mouse visual cortex. Nat. Commun. *11*, 505.
- Kanamori, T., and Mrsic-Flogel, T.D. (2022). Independent response modulation of visual cortical neurons by attentional and behavioral states. Neuron 110, 3907–3918.e6.
- McAdams, C.J., and Maunsell, J.H. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. J. Neurosci. 19, 431–441.

- Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. Science 229, 782–784.
- Luck, S.J., Chelazzi, L., Hillyard, S.A., and Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. J. Neurophysiol. 77, 24–42.
- Cohen, M.R., and Maunsell, J.H. (2009). Attention improves performance primarily by reducing interneuronal correlations. Nat. Neurosci. 12, 1594–1600.
- Fries, P., Reynolds, J.H., Rorie, A.E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. Science 291, 1560–1563.
- Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2009). Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. Neuron 63, 879–888.
- Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. Nat. Rev. Neurosci. 5, 793–807.
- 16. Kepecs, A., and Fishell, G. (2014). Interneuron cell types are fit to function. Nature 505, 318–326.
- Pfeffer, C.K., Xue, M., He, M., Huang, Z.J., and Scanziani, M. (2013). Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. Nat. Neurosci. 16, 1068–1076.
- Tremblay, R., Lee, S., and Rudy, B. (2016). GABAergic interneurons in the neocortex: from cellular properties to circuits. Neuron 91, 260–292.
- Campagnola, L., Seeman, S.C., Chartrand, T., Kim, L., Hoggarth, A., Gamlin, C., Ito, S., Trinh, J., Davoudian, P., Radaelli, C., et al. (2022). Local connectivity and synaptic dynamics in mouse and human neocortex. Science 375, eabj5861.
- Lee, S., Kruglikov, I., Huang, Z.J., Fishell, G., and Rudy, B. (2013). A disinhibitory circuit mediates motor integration in the somatosensory cortex. Nat. Neurosci. 16, 1662–1670.
- Pi, H.-J., Hangya, B., Kvitsiani, D., Sanders, J.I., Huang, Z.J., and Kepecs, A. (2013). Cortical interneurons that specialize in disinhibitory control. Nature 503, 521–524.
- Vogels, T.P., and Abbott, L.F. (2009). Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. Nat. Neurosci. 12, 483–491.
- Letzkus, J.J., Wolff, S.B.E., and Lüthi, A. (2015). Disinhibition, a circuit mechanism for associative learning and memory. Neuron 88, 264–276.
- Williams, L.E., and Holtmaat, A. (2019). Higher-order thalamocortical inputs gate synaptic long-term potentiation via disinhibition. Neuron 101, 91–102.e4.
- Letzkus, J.J., Wolff, S.B.E., Meyer, E.M.M., Tovote, P., Courtin, J., Herry, C., and Lüthi, A. (2011). A disinhibitory microcircuit for associative fear learning in the auditory cortex. Nature 480, 331–335.
- Krabbe, S., Paradiso, E., d'Aquin, S., Bitterman, Y., Courtin, J., Xu, C., Yonehara, K., Markovic, M., Müller, C., Eichlisberger, T., et al. (2019). Adaptive disinhibitory gating by VIP interneurons permits associative learning. Nat. Neurosci. 22, 1834–1843.
- Canto-Bustos, M., Friason, F.K., Bassi, C., and Oswald, A.-M.M. (2022). Disinhibitory circuitry gates associative synaptic plasticity in olfactory cortex. J. Neurosci. 42, 2942–2950.
- Guet-McCreight, A., Skinner, F.K., and Topolnik, L. (2020). Common principles in functional organization of VIP/calretinin cell-driven disinhibitory circuits across cortical areas. Front. Neural Circuits 14, 32.
- Fu, Y., Tucciarone, J.M., Espinosa, J.S., Sheng, N., Darcy, D.P., Nicoll, R.A., Huang, Z.J., and Stryker, M.P. (2014). A cortical circuit for gain control by behavioral state. Cell *156*, 1139–1152.
- Shapiro, J.T., Michaud, N.M., King, J.L., and Crowder, N.A. (2022). Optogenetic activation of interneuron subtypes modulates visual contrast responses of mouse V1 neurons. Cereb. Cortex 32, 1110–1124.

Neuron Article



- Jackson, J., Ayzenshtat, I., Karnani, M.M., and Yuste, R. (2016). VIP+ interneurons control neocortical activity across brain states. J. Neurophysiol. *115*, 3008–3017.
- Ayzenshtat, I., Karnani, M.M., Jackson, J., and Yuste, R. (2016). Cortical control of spatial resolution by VIP+ interneurons. J. Neurosci. 36, 11498–11509.
- Karnani, M.M., Jackson, J., Ayzenshtat, I., Hamzehei Sichani, A., Manoocheri, K., Kim, S., and Yuste, R. (2016). Opening holes in the blanket of inhibition: localized lateral disinhibition by VIP interneurons. J. Neurosci. 36, 3471–3480.
- Bigelow, J., Morrill, R.J., Dekloe, J., and Hasenstaub, A.R. (2019). Movement and VIP interneuron activation differentially modulate encoding in mouse auditory cortex. eNeuro 6, ENEURO.0164-19.2019.
- Garcia-Junco-Clemente, P., Ikrar, T., Tring, E., Xu, X., Ringach, D.L., and Trachtenberg, J.T. (2017). An inhibitory pull-push circuit in frontal cortex. Nat. Neurosci. 20, 389–392.
- Kamigaki, T., and Dan, Y. (2017). Delay activity of specific prefrontal interneuron subtypes modulates memory-guided behavior. Nat. Neurosci. 20, 854–863.
- 37. Lee, A.T., Cunniff, M.M., See, J.Z., Wilke, S.A., Luongo, F.J., Ellwood, I.T., Ponnavolu, S., and Sohal, V.S. (2019). VIP interneurons contribute to avoidance behavior by regulating information flow across hippocampalprefrontal networks. Neuron 102, 1223–1234.e4.
- Zhang, S., Xu, M., Kamigaki, T., Hoang Do, J.P.H., Chang, W.-C., Jenvay, S., Miyamichi, K., Luo, L., and Dan, Y. (2014). Selective attention. Longrange and local circuits for top-down modulation of visual cortex processing. Science 345, 660–665.
- Sridharan, D., and Knudsen, E.I. (2015). Selective disinhibition: A unified neural mechanism for predictive and post hoc attentional selection. Vision Res. 116, 194–209.
- Batista-Brito, R., Zagha, E., Ratliff, J.M., and Vinck, M. (2018). Modulation of cortical circuits by top-down processing and arousal state in health and disease. Curr. Opin. Neurobiol. 52, 172–181.
- Disney, A.A., Aoki, C., and Hawken, M.J. (2007). Gain modulation by nicotine in macaque V1. Neuron 56, 701–713.
- Herrero, J.L., Roberts, M.J., Delicato, L.S., Gieselmann, M.A., Dayan, P., and Thiele, A. (2008). Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. Nature 454, 1110–1114.
- 43. Pinto, L., Goard, M.J., Estandian, D., Xu, M., Kwan, A.C., Lee, S.H., Harrison, T.C., Feng, G., and Dan, Y. (2013). Fast modulation of visual perception by basal forebrain cholinergic neurons. Nat. Neurosci. 16, 1857–1863.
- Alitto, H.J., and Dan, Y. (2012). Cell-type-specific modulation of neocortical activity by basal forebrain input. Front. Syst. Neurosci. 6, 79.
- 45. Gasselin, C., Hohl, B., Vernet, A., Crochet, S., and Petersen, C.C.H. (2021). Cell-type-specific nicotinic input disinhibits mouse barrel cortex during active sensing. Neuron 109, 778–787.e3.
- Poorthuis, R.B., Enke, L., and Letzkus, J.J. (2014). Cholinergic circuit modulation through differential recruitment of neocortical interneuron types during behaviour. J. Physiol. 592, 4155–4164.
- McBride, E.G., Lee, S.-Y.J., and Callaway, E.M. (2019). Local and global influences of visual spatial selection and locomotion in mouse primary visual cortex. Curr. Biol. 29, 1592–1605.e5.
- 48. Poort, J., Khan, A.G., Pachitariu, M., Nemri, A., Orsolic, I., Krupic, J., Bauza, M., Sahani, M., Keller, G.B., Mrsic-Flogel, T.D., et al. (2015). Learning enhances sensory and multiple non-sensory representations in primary visual cortex. Neuron 86, 1478–1490.
- 49. Dana, H., Sun, Y., Mohar, B., Hulse, B.K., Kerlin, A.M., Hasseman, J.P., Tsegaye, G., Tsang, A., Wong, A., Patel, R., et al. (2019). High-performance calcium sensors for imaging activity in neuronal populations and microcompartments. Nat. Methods *16*, 649–657.
- Millman, D.J., Ocker, G.K., Caldejon, S., Kato, I., Larkin, J.D., Lee, E.K., Luviano, J., Nayan, C., Nguyen, T.V., North, K., et al. (2020). VIP interneu-

rons in mouse primary visual cortex selectively enhance responses to weak but specific stimuli. eLife 9, e55130.

- Niell, C.M., and Stryker, M.P. (2010). Modulation of visual responses by behavioral state in mouse visual cortex. Neuron 65, 472–479.
- Elsayed, G.F., Lara, A.H., Kaufman, M.T., Churchland, M.M., and Cunningham, J.P. (2016). Reorganization between preparatory and movement population responses in motor cortex. Nat. Commun. 7, 13239.
- 53. Khan, A.G., Poort, J., Chadwick, A., Blot, A., Sahani, M., Mrsic-Flogel, T.D., and Hofer, S.B. (2018). Distinct learning-induced changes in stimulus selectivity and interactions of GABAergic interneuron classes in visual cortex. Nat. Neurosci. 21, 851–859.
- Tejero-Cantero, A., Boelts, J., Deistler, M., Lueckmann, J.-M., Durkan, C., Gonçalves, P.J., Greenberg, D.S., and Macke, J.H. (2020). sbi: A toolkit for simulation-based inference. J. Open Source Software 5, 2505.
- Wang, L., Rangarajan, K.V., Gerfen, C.R., and Krauzlis, R.J. (2018). Activation of striatal neurons causes a perceptual decision bias during visual change detection in mice. Neuron 97, 1369–1381.e5.
- You, W.-K., and Mysore, S.P. (2020). Endogenous and exogenous control of visuospatial selective attention in freely behaving mice. Nat. Commun. *11*, 1986.
- Hu, F., and Dan, Y. (2022). An inferior-superior colliculus circuit controls auditory cue-directed visual spatial attention. Neuron *110*, 109–119.e3.
- Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2007). Differential attention-dependent response modulation across cell classes in macaque visual area V4. Neuron 55, 131–141.
- Lee, S.-H., Kwan, A.C., Zhang, S., Phoumthipphavong, V., Flannery, J.G., Masmanidis, S.C., Taniguchi, H., Huang, Z.J., Zhang, F., Boyden, E.S., et al. (2012). Activation of specific interneurons improves V1 feature selectivity and visual perception. Nature 488, 379–383.
- Mysore, S.P., and Knudsen, E.I. (2013). A shared inhibitory circuit for both exogenous and endogenous control of stimulus selection. Nat. Neurosci. 16, 473–478.
- Olsen, S.R., Bortone, D.S., Adesnik, H., and Scanziani, M. (2012). Gain control by layer six in cortical circuits of vision. Nature 483, 47–52.
- Keller, A.J., Dipoppa, M., Roth, M.M., Caudill, M.S., Ingrosso, A., Miller, K.D., and Scanziani, M. (2020). A disinhibitory circuit for contextual modulation in primary visual cortex. Neuron *108*, 1181–1193.e8.
- Larkum, M. (2013). A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. Trends Neurosci. 36, 141–151.
- Chen, S.X., Kim, A.N., Peters, A.J., and Komiyama, T. (2015). Subtypespecific plasticity of inhibitory circuits in motor cortex during motor learning. Nat. Neurosci. 18, 1109–1115.
- Seybold, B.A., Phillips, E.A.K., Schreiner, C.E., and Hasenstaub, A.R. (2015). Inhibitory actions unified by network integration. Neuron 87, 1181–1192.
- Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C.B., Carandini, M., and Harris, K.D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. Science 364, 255.
- Nogueira, R., Rodgers, C.C., Bruno, R.M., and Fusi, S. (2023). The geometry of cortical representations of touch in rodents. Nat. Neurosci. 26, 239–250.
- Voitov, I., and Mrsic-Flogel, T.D. (2022). Cortical feedback loops bind distributed representations of working memory. Nature 608, 381–389.
- Attinger, A., Wang, B., and Keller, G.B. (2017). Visuomotor coupling shapes the functional development of mouse visual cortex. Cell 169, 1291–1302.e14.
- Pachitariu, M., Stringer, C., Dipoppa, M., Schröder, S., Rossi, L.F., Dalgleish, H., Carandini, M., and Harris, K.D. (2017). Suite2p: beyond 10,000 neurons with standard two-photon microscopy. Preprint at bioRxiv. https://doi.org/10.1101/06150710.1101/061507.





- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Müller, A., Nothman, J., Louppe, G., et al. (2018). Scikitlearn: Machine Learning in Python. Preprint at arXiv, 10.48550/ arXiv.1201.0490 10.48550/arXiv.1201.0490.
- Schindelin, J., Arganda-Carreras, I., Frise, E., Kaynig, V., Longair, M., Pietzsch, T., Preibisch, S., Rueden, C., Saalfeld, S., Schmid, B., et al.

(2012). Fiji - an Open Source platform for biological image analysis. Nat Methods 9. https://doi.org/10.1038/nmeth.2019.

- 73. Brainard, D.H. (1997). The Psychophysics Toolbox. Spat Vis 10, 433–436.
- Mazurek, M., Kager, M., and Van Hooser, S.D. (2014). Robust quantification of orientation selectivity and direction selectivity. Frontiers in Neural Circuits 8.

Neuron Article



STAR***METHODS**

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER			
Antibodies					
Rat anti-Somatostatin	Millipore	Cat# MAB354, RRID:AB_2255365			
Mouse anti-Parvalbumin	Swant	Cat# 235, RRID:AB_10000343			
Rabbit anti-Vasoactive Intestinal Peptide (VIP)	Immunostar	Cat# 8726017, RRID:AB_2922959			
Goat anti-Rat Alexa 647	Thermo Fisher	Cat# A-21247, RRID:AB_141778			
Donkey anti-Mouse Dylight 405	Jackson Immunoresearch	Cat# 715-475-150, RRID:AB_2340839			
Goat anti-Rabbit Alexa 594	Thermo Fisher	Cat# A-11012, RRID:AB_2534079			
Bacterial and virus strains					
AAV1-hSyn-GCaMP7f	Zurich Vector Core	Cat# v292-1			
AAV5-hSyn-FLEX-ChrimsonR-tdTomato	Addgene	Cat# 62723-AAV5			
AAV5-hSyn-FLEX-ArchT-tdTomato	Addgene	Cat# 28305-AAV5			
AAV1-FLEX-tdTomato	Addgene	Cat# 28306-AAV1			
AAV8-hEF1a-jGCaMP7f	Zurich Vector Core	Cat# v401-8			
Experimental models: Organisms/strains					
Mouse: VIP-Cre	Jackson Laboratory	RRID:IMSR_JAX:031628			
Software and algorithms					
Suite2p	Pachitariu et al. ⁷⁰	https://github.com/MouseLand/suite2p			
Scikit-learn	Pedregosa et al. ⁷¹	https://scikit-learn.org/stable/			
Fiji	Schindelin et al. ⁷²	https://imagej.net/software/fiji/			
The Neural Decoding Toolbox	Meyers	readout.info			
sbi	Tejero-Cantero et al.54	https://github.com/mackelab/sbi			
Psychtoolbox-3	Brainard ⁷³	http://psychtoolbox.org/			
Python version 3.8	Python Software Foundation	https://www.python.org/			
MATLAB	MathWorks	https://www.mathworks.com/			
Attentional modulation model	This paper; zenodo.org	10.5281/zenodo.10041197			

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Adil G. Khan (khan.adil@ kcl.ac.uk).

Materials availability

This study did not generate new unique reagents.

Data and code availability

All data reported in this paper will be shared by the lead contact upon request. All original code has been deposited at https://zenodo. org/records/10041198 and is publicly available as of the date of publication. DOIs are listed in the key resources table. Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

EXPERIMENTAL MODEL DETAILS

Mouse lines

A total of 20 VIP-Cre mice (MGI:4431361) were used in this study (5 male, 15 female). Mice were aged between postnatal days 78 and 84 at the time of surgery.





All experimental procedures were carried out in accordance with institutional animal welfare guidelines and licensed by the UK Home Office. All mice were healthy and were not involved in previous experiments.

METHOD DETAILS

Animals and surgical procedures

Mice were anaesthetised using isoflurane, at 4% concentration for induction and at 1-1.5% for maintenance. At the start of the surgery additional drugs were given to provide analgesia (Metacam 5mg/kg), anti-inflammatory effects (dexamethasone 3.8mg/kg), and to reduce mucus secretions (Atropine 0.08mg/kg). Eye-cream (Maxitrol) was applied to the eyes to prevent drying and body temperature was maintained at 37°C using a heating mat and rectal temperature probe (Harvard Apparatus).

Mice were implanted with a chronic imaging window (3-5 mm diameter) above the right V1 (2.7mm medial, 0.6mm anterior to lambda) following viral injections using a pressure micro-injection system (Picospritzer III, Parker) of AAV1-hSyn-GCaMP7f mixed with either AAV5-hSyn-FLEX-ChrimsonR-tdTomato (8 mice), AAV5-hSyn-FLEX-ArchT-tdTomato (5 mice) or AAV1-FLEX-tdTomato (3 mice). Mice used in the low contrast stimulus experiments were injected with a mixture of AAV8-hEF1a-jGCaMP7f and AAV5-hSyn-FLEX-ArchT-tdTomato (4 mice). The craniotomy was then sealed with a glass coverslip and cyano-acrylic glue (Loc-tite). A custom machined aluminium head-plate was cemented onto the skull using dental cement (C&B Superbond).

Before the removal of anaesthesia mice were injected with antibiotic (Betamox 120mg/kg) and analgesia (methadone hydrochloride 10mg/kg). Mice were closely monitored for 4 days after surgery and further analgesia was given daily for 1-2 days during recovery of the animal. Imaging and behavioural training were not started until at least one week after surgery.

Immunohistochemistry and ex vivo imaging

Brains were fixed by transcardial perfusion with 4% paraformaldehyde in phosphate buffer 0.1 M, followed by 24h of postfixation in the same solution at 4°C. The whole brains were incubated successively in 15% and 30% sucrose in phosphate-buffered saline (PBS) at 4°C for 2 and 12h respectively. Brains were sectioned tangentially to the surface of visual cortex at 80µm thickness on a microtome (Leica). Slides were washed and permeabilized with 0.4% Triton X-100 in PBS for 4 × 15 minutes and then incubated with blocking buffer (0.3% Triton X-100 + 5% BSA + 10% Normal Donkey Serum and 10% Normal Goat Serum in PBS) for 3h at room temperature. Primary antibodies were incubated with blocking buffer (0.3% Triton X-100 + 1% BSA + 5% Normal Donkey Serum and 5% Normal Goat Serum in PBS) overnight at 4°C. The next day, slides were washed and incubated for 2h with secondary antibodies, then mounted in DABCO-PVA (2.5% DABCO, 10% polyvinyl alcohol (Sigma; Type II), 5% glycerol and 25 mM Tris buffer at pH 8.7). The slides were imaged with a confocal microscope (Zeiss LSM 800), and confocal z-stacks were compared with the previously acquired in vivo imaging planes and z-stacks of the recording sites. We determined the approximate location of the injection site using GCaMP7f fluorescence and then used blood vessel patterns and cellular morphology to identify the imaging site. We matched at least three points in the confocal z-stack to points in the in vivo imaging plane to obtain a three-dimensional transformation matrix that was applied to the entire confocal z-stack. Cells were then manually identified and assigned to cell classes based on immunostaining. Primary antibodies and dilutions used: Rat anti-Somatostatin, 1:200 (Millipore MAB354); Mouse anti-Parvalbumin, 1:5000, (Swant PV235); Rabbit anti-Vasoactive Intestinal Peptide (VIP), 1:500 (Immunostar #20077). Secondary antibodies and dilutions used: Goat anti-Rat Alexa 647, 1:500 (Thermo Fisher #A21247); Donkey anti-Mouse Dylight 405, 1:500 (Jackson Immunoresearch #715-475-150); Goat anti-Rabbit Alexa 594, 1:500 (Thermo Fisher #A11012).

Two-photon calcium imaging

Two-photon imaging was performed using a custom-built resonant scanning two-photon microscope (Cosys) and a Chameleon Vision S laser (Coherent) at 930nm using a 16X,.8NA objective (Nikon). Images were acquired using a 12 KHz resonant scanner (Cambridge Technology) and an FPGA module (PXIe-7965R FlexRIO, National Instruments). Multi-plane imaging was performed using a piezoelectric objective scanner (Physik Instrumente). All recordings were made of neurons in L2/3 (generally 150-250um below the surface). Each imaging volume consisted of 6 planes, 20µm apart, approximately 450x450um, 512x512 pixels in size. Images were captured at an effective framerate of 6.3 Hz per volume. At the beginning of each session anatomical landmarks were used to find and record from the same imaging site as on previous days. Mice which were found to have bone regrowth under the window, poor viral expression or many brightly labelled cells with nuclear GCaMP7f expression were excluded from the study.

The coarse receptive field position of each imaging site was determined on the first imaging day to ensure the visual stimuli were approximately centred in the receptive field. The monitor in front of the contralateral eye (covering $\sim 100 \times 60$ degrees of visual space) was divided into a 4 \times 3 grid and stimuli alternating between black and white at 2Hz were presented at each grid position on a grey background in randomized order (10 repetitions). Stimuli were generated using Psychtoolbox-3 in MATLAB. At the end of all *in vivo* imaging data collection, a high-quality image stack of all recording sites was acquired under anaesthesia to aid subsequent registration with immunohistochemically labelled brain slices.

Behavioural training

The equipment and method used for behavioural training was similar to previous studies.^{5,48} Mice were trained first on a visual discrimination go-no go task followed by the full odour-visual attention switching task described below. Mice were food restricted

Neuron Article



to maintain at least 80% of their free-feeding body weight (typically 85-90%, 2-3g of standard food pellets per animal per day) but had free access to water. A 10% solution of soy milk powder (SMA Wysoy) was used as reward during the task and delivered through a spout positioned near the snout of the mouse. Licks to this spout were detected through a piezo disc sensor and reward was released by opening a pinch valve (NResearch), both controlled by custom electronics. The visual stimuli were presented on two luminance-corrected monitors (luminance meter LS-100, Konica Minolta) positioned at 45° angles and 25cm distance relative to the mouse. Visual stimuli were generated using psychoolbox-3 and all behavioural tasks were controlled using custom scripts written in MATLAB and with a Teensy microcontroller board.

Mice were first habituated to handling and gentle restraint over two to three days and were then head-fixed and trained to run on a polystyrene cylinder (20cm diameter) for a further one to four days. Mice were free to run on the polystyrene cylinder during all awake recordings and their running speed on this cylinder was measured using an incremental rotary encoder (Kübler).

Once mice were running reliably on the wheel, they performed one closed-loop visual discrimination behavioural session during which the movement of the mouse on the wheel controlled the movement of visual gratings on the screen. After this closed-loop session all subsequent sessions were with fixed spatial and temporal frequency of the drifting gratings, and mice were trained to run for sustained periods of time to initiate trials - at least 2.8s with an added random duration drawn from an exponential distribution (mean 0.4s). The stimuli used for visual discrimination were two sinusoidal gratings drifting in the opposite direction to the direction of running, with a fixed spatial and temporal frequency of 0.1 cycles per degree and 2Hz respectively. Unless otherwise specified the rewarded and unrewarded gratings were oriented +/-15° relative to vertical, symmetrically on both screens. The stimulus presented on a given trial was selected at random.

The mouse could trigger the release of a drop of soya milk when the rewarded grating was displayed by licking the reward spout during the 'reward period'. The reward period started 1.5s (with added random duration, mean 0.2s) after the rewarded visual stimulus onset and lasted until the offset of the stimulus 1s later. If the mouse licked during the 'reward period' the trial was recorded as a 'hit', if the mouse did not lick it was recorded as a 'miss' and a drop of soya milk was dispensed automatically shortly before the disappearance of the visual stimulus. Miss trials were rare after the first few days of training. A lick at any time when the unrewarded grating was displayed was recorded as a 'false alarm' and the mouse was punished with a 4s time-out period, during which the unrewarded grating persisted on the screen and any more licks reset the time-out duration. Ignoring the unrewarded visual stimulus by not licking was recorded as a 'correct rejection'. In the initial stages of training, to discourage incorrect licking the probability of unrewarded trials was sometimes temporarily increased from 0.5 to 0.7. Mice typically learned the visual discrimination task in 5–10 days, with the threshold for learning defined as three consecutive days of discrimination with a behavioural d-prime score of 2.0 or above. Behavioural d-prime was calculated as: $d' = \Phi^{-1} (H) - \Phi^{-1} (F)$, where Φ^{-1} is the normal inverse cumulative distribution function, H is the rate of hit trials, and F is the rate of false alarm trials. H and F were defined as licking in response to visual stimulus 1 and visual stimulus 2 respectively in both block types.

After learning the visual discrimination task, mice were trained to perform odour discrimination. The odour discrimination task was identical in structure to the visual discrimination task except that instead of visual stimuli one of two odour stimuli were presented to the mouse via polyethylene tubing positioned above the snout of the mouse. A custom-built flow-dilution olfactometer calibrated with a mini PID (Aurora) delivered 10%–20% saturated vapour concentration of two solutions, 10% soy milk (rewarded odour) and 10% soy milk with 0.1% p-Cymene mixture (unrewarded odour). Mice typically started accurately discriminating between the odours after 30-40 trials, after which they were trained to switch between blocks of the olfactory and visual discrimination task.

Mice typically learned to perform the attention switching task in a further 1-3 days. In the olfactory blocks, 70% of odour stimuli were preceded by one of the two visual gratings presented in the visual discrimination task. These irrelevant visual stimuli were displayed for a fixed duration of 1.8s, with an onset delay distribution identical to the visual block, and neither grating was rewarded or punished. Mice learned to accurately discriminate between the odours while ignoring, i.e., not licking in response to either irrelevant visual grating. One of the two odours followed the irrelevant visual grating offset with a delay of 1.5s, plus an added random duration drawn from an exponential distribution with mean 0.2s.

Mice performed the switching task with two-photon imaging for sessions of 60-90 minutes, with a median of 4 total blocks of trials (odour and visual combined, range 3 to 6 blocks) and a median of 100 trials per block (range 50 to 175 trials). Since visual stimuli were presented in only 70% of odour trials, odour blocks were proportionately longer.

Optogenetic manipulations

Expression of the tdTomato conjugated opsin (Chrimson or ArchT) was first verified in each imaging site through two-photon imaging at 1030nm excitation wavelength. Optogenetic light was delivered using a digitally triggered 637nm laser (OBIS 637nm LX, Coherent), through a 200µm diameter 0.39 NA optic fibre (Thorlabs) positioned above the cranial window. To allow for quasi-simultaneous two-photon imaging and optogenetic activation, the laser and stimulus monitors were blanked during the linear phase of the resonant scanner.

For each two-photon imaging site an optogenetic calibration session was performed. During these calibration sessions the screens were grey, and 8 light powers (including 0%) were applied in pseudo-random sequence to the imaging window for 1.5s with 5s intervals. The effective maximum output used in mice expressing ArchT was 9mW, and in mice expressing Chrimson was 3mW. The average activity of each ROI in the 1s before the optogenetic laser onset was subtracted from the light period and the resulting base-line corrected activity was used for the calibration plots. Based on the shape of these calibration curves, 2 powers were chosen for





the behavioural task for each imaging site: the lowest power producing a saturated or just below saturated response, and a second power with approximately half of this effect. For Chrimson these light powers were typically 0.6 mW (low, range 0.5 to 0.6 mW) and 1.5 mW (high, range 1.5 to 2.25 mW) although tailored values were selected for each site. For ArchT, only one laser power was selected, which was typically 1.5mW (range 1.15 to 1.8mW).

The power calibration was performed in the first session only. The maximum difference in time between the calibration session for each recording site and the last optogenetic session performed was 3 to 22 days, median 11 days).

There was no significant change in the expression level of Chrimson-tdTomato between the power calibration and final optogenetic sessions, as measured by comparing the fluorescence in the red channel for all recorded VIP cells (Wilcoxon signed rank test, p = 0.813, n = 8 mice). There was also no significant change in ArchT-tdTomato expression (Wilcoxon signed rank test, p = 0.438, n = 6 mice) or no-opsin tdTomato expression in control mice (Wilcoxon signed rank test, p = 0.625, n = 3 mice). The same was true when pooling all mice regardless of injected virus (Wilcoxon signed rank test, p = 0.836, n = 17 mice).

Direction tuning

To examine the effect of the optogenetic light on visual processing in mice passively viewing stimuli, two-photon imaging sessions were conducted while head-fixed mice (free to run on a polystyrene wheel) were shown sinusoidal visual gratings drifting in one of 8 directions separated by 45°, with a spatial frequency of 0.1 cycles per degree and temporal frequency of 2Hz. These visual stimuli were randomly interleaved and one of three laser powers (including 0%, the same as those used in the attention switching task) was selected with equal probability for each visual stimulus presentation. The visual gratings were presented for 2s and the optogenetic laser lasted from 100ms before the stimulus onset to 1.5s after the start of stimulus presentation. There was a 5s interval before the start of the next visual stimulus presentation.

Direction tuning curves were constructed for each cell using their mean responses to each direction after baseline correction. The activity of each neuron was 'soft' normalized⁵² so that neurons with strong responses had approximately unity firing rate range (normalization factor = firing rate range + 0.2).

When aligning direction tuning curves across neurons to each neuron's preferred direction, the preferred direction was determined by taking the greatest response after pooling all light power and no light conditions. This prevented artefactual results which would occur from selecting the maximum firing rate in one condition (e.g., no light) and comparing this to other conditions (e.g., optogenetic stimulation). An orientation selectivity index was calculated as:

$$OSI = \frac{R_{pref} - R_{orth}}{R_{pref} + R_{orth}}$$

Where R_{pref} and R_{orth} are the average responses to the preferred and orthogonal directions respectively.⁷⁴ Neurons were considered orientation selective if their OSI was greater than 0.33, such that the response at the preferred direction was twice as large as it's response to the orthogonal direction.

Locomotion modulation

To assess the effect of locomotor activity on the visual responses of neurons a locomotion modulation index (LMI) was calculated. Two-photon imaging sessions were performed with visual stimuli presented as described above (Direction tuning). Visual stimulus presentations were separated according to whether the mouse was stationary (average speed < 1 cm/s during visual stimulus presentation) or locomoting during the visual stimulus (average speed > 1 cm/s during visual stimulus presentation).

Direction tuning curves were then constructed as described above, pooling across stationary, running, optogenetic light and no-light conditions to calculate each neuron's preferred direction. Each neuron's LMI was then calculated as:

$$LMI = \frac{L - S}{L + S}$$

Where L and S are the average responses to the preferred orientation when the mouse was locomoting and stationary respectively. The LMI was calculated for optogenetic light and no light trials separately using the same preferred orientation for each neuron.

Optogenetic light during attention-switching

Application of the optogenetic light during attention switching sessions was similar to the orientation mapping sessions. Optogenetic laser powers were interleaved and randomly selected on each trial from 1 of 3 powers (No laser, low power, high power) with equal probability. The optogenetic laser was on from 100ms before stimulus presentation to 1.5s after visual stimulus onset in Chrimson mice, or to the offset of the visual stimulus for ArchT mice. The light was ramped off over 0.2s. The optogenetic laser was delivered only during presentation of the visual stimuli in both the visual and odour blocks.

QUANTIFICATION AND STATISTICAL ANALYSIS

Statistical parameters for each analysis can be found in the appropriate figure legends.

Neuron Article



Pre-processing

Pre-processing of two-photon calcium imaging data was performed using the software Suite2p (https://github.com/MouseLand/ suite2p) to correct for motion, detect regions of interest (ROIs) and extract the raw fluorescence time series of those ROIs, F(t). Each site yielded between 164 and 688 cells, median = 432 cells. We corrected the calcium traces for out of focus neuropil fluorescence using the neuropil masks identified by suite2p. For each frame we subtracted 0.7 * (neuropil - median neuropil fluorescence). Subsequent analysis unless otherwise specified was done with custom code in MATLAB and Python. Baseline fluorescence $F_0(t)$ was computed by smoothing F(t) (causal moving average of 0.75s) and taking the 2.5th percentile of the smoothed data. The change in fluorescence relative to baseline, $\Delta F/F$, was computed by taking the difference between F and F_0 , and dividing by F_0 . Video recording (The Imaging Source) of the eye contralateral to the imaging site was performed during all sessions, and the time-points of saccades and blinks identified. Frames in which the mouse made a saccade or blinked were removed from further analysis. To identify VIP interneurons labelled with tdTomato *in vivo*, a brief dual channel recording of the imaging planes in red and green was taken before each imaging session at an excitation wavelength of 1020nm.

Behavioural controls

To assess the proportion of neural activity which was attributable to overt behaviour recorded during our task, a linear model was fit using ridge regression to predict neural activity. The model was constructed by combining multiple sets of variables into a design matrix, to capture signal modulation by the following different task or behavioural events: 2 visual stimuli, 2 odour stimuli, reward delivery, licks, running speed, block type, and an interaction term for visual stimuli and block type. Each stimulus/event variable was structured to capture a time-varying event kernel. Variables therefore consisted of a vector of the relevant stimulus/event, and copies of this vector, each shifted in time by one frame for specific durations. For sensory stimuli, the time-shifted copies ranged up to 2s after the original. For motor events (running and licking) the time-shifted copies spanned the frames from 0.5s before until 2s after the original. The model was fit with 5-fold cross validation and the coefficient of determination (R^2) was calculated based on the predictions of the model on held out data not used during training. We then assessed the predictive power of the behavioural model variables by comparing the R^2 value for the full model to a model without the running and licking predictors, taking the proportion 1-(no behaviour model R^2).

To decode block type based on neural activity or running speed the neural decoding toolbox readout.info was used with the max_ correlation_coefficient_CL classifier. To separate sessions according to their time-point of divergence in running speed the pre-stimulus baseline speed was subtracted and the first time-point found at which there was a significant difference in running speed between the attended and ignored rewarded visual stimulus trials using Wilcoxon rank sum tests at each time-point.

Selectivity

We computed a selectivity index for individual ROIs as the difference between the mean response to each of the two gratings divided by the pooled standard deviation of that ROIs responses. Unless otherwise specified, all selectivity values presented here are from an analysis of the activity in the first 1s of visual stimulus presentation. To calculate Δ Selectivity we took the difference selectivity(attend) - selectivity(ignore). For cells that were negatively selective in the attend condition we multiplied the resultant values by -1, to ensure that cells that became more selective with attention had positive values.

To test if an ROI was significantly selective within a certain time window, a two-sided Wilcoxon rank sum test was performed comparing the activity on trials for the two visual stimuli. ROIs were excluded from further analysis if they displayed selectivity in the period before visual stimuli were presented. To find ROIs that significantly changed their selectivity with attention, an ROI's selectivity in the attend condition was compared to a distribution produced through bootstrapping using the data in the ignore condition (1000 repeats). If the attend selectivity was below or above the 2.5th or 97.5th percentiles respectively of the bootstrapped distribution, then the ROI was considered to have significantly changed its selectivity with attention. To avoid artefactual effects from selecting cells in one condition and testing in another, the test for significance was performed with no light and all light powers pooled.

After each behavioural block transition, transient periods of less accurate behaviour were discarded by identifying the trial in each block beyond which behaviour was stably accurate, that is, where mice displayed greater than 75% accuracy on both the go and nogo stimuli for the remainder of the block, and for the odour block, where mice licked in response to fewer than 25% of either of the irrelevant visual gratings for the remainder of the block. Light and no light trials were pooled to identify this point to avoid artefactual results.

Linear mixed-effects models

For the analysis of the effects of optogenetic manipulation of VIP interneurons on stimulus evoked responses and stimulus selectivity during the attention switching task, multiple sessions were captured per mouse. Linear mixed- effects models were used to analyse this nested data. Models were fit using the matlab function *fitIme* with a random intercept for each mouse and maximum likelihood estimation to find the parameters.

Orthogonality

Visualisation of the LDA transformation of neural activity was done using the Python library scikit-learn. All other analysis for testing orthogonality was done on axes found in MATLAB using the *fitcdiscr* function. The alignment of the axes best separating the





optogenetic modulation and attentional modulation were found through the cosine similarity of the coefficients of two linear discriminant analysis models. One model separated visual stimulus trials in the odour block from visual stimulus trials in the visual block, both sets of trials without the optogenetic perturbation. The second model separated visual stimulus trials in the odour block without optogenetic perturbation from visual stimulus trials in the odour block with optogenetic perturbation.

The cosine similarity for the split data control was found using a similar approach. In the visual block, the rewarded and unrewarded visual stimulus trials were separated into two halves and the coefficients used to calculate cosine similarity were for two models separating the rewarded and unrewarded trials using non-overlapping halves of the data. This process was repeated 50 times for each session and the resulting values were averaged to produce one value for cosine similarity for split data for each session.

To ensure that we do not obtain orthogonal axes simply because they lie within a high-dimensional neural subspace, random axes were found using the method from Elsayed et al.⁵² The neural covariance was estimated from each session and a Monte Carlo analysis used to sample pairs of random axes (10,000 samples) that were then used to calculate an expectation of cosine similarity based only on the dimensionality of the data.

Noise correlations

To calculate noise correlation, the average stimulus evoked response across all trials of a particular type was taken for each cell and subtracted from each trial of the corresponding type. There were 14 trial types in total, rewarded and unrewarded odour stimuli, and rewarded and unrewarded visual stimuli in the visual and odour blocks, all at multiple laser powers. The Pearson correlation coefficient was then used to quantify the correlation between responses of pairs of cells for each trial type. Changes in noise correlations between different cell types with attention or optogenetic modulation were tested using a Wilcoxon signed-rank test on all sessions for which the post-hoc immunomatching had been successful (n = 4 mice).

Circuit model

We modelled a circuit consisting of an excitatory population of layer 2/3 pyramidal cells (PYR), and three inhibitory populations, corresponding to PV, SOM, and VIP interneurons. The activity of the population *i* is described by its response r_i , which evolves over time according to the following equation:

$$\tau_i \frac{dr_i}{dt} = -r_i + \varphi \left(I_i^{TD} \left(I_i^{b} + I_i^{s} \right) + \sum_j \gamma W_{ij} r_j \right)$$

where $i, j \in \{PYR, PV, SOM, VIP\}$ and τ_i is the time constant of population i, I_i^b is the baseline input to population i,

 I_i^s is the stimulus-dependent feedforward input to population *i*,

 I_i^{TD} is the modulatory top-down input - the attentional modulation of population *i*, and

 $\sum_{i} \gamma W_{ii} r_i$ is the recurrent input from the local circuit and W_{ii} is the effective synaptic weight from cell population *j* to *i*.

 γ is a factor by which all weights are multiplied.

 $\varphi(x)$ is the activation function:

$$\varphi(x) = \begin{cases} 0 & \text{if } x \le 0\\ (r_{max} - r_0) tanh(x/(r_{max} - r_0)) & \text{if } x > 0 \end{cases}$$

where $r_0 = 1.0$ and $r_{max} = 20.0$ denote the baseline and maximum activity, respectively. PYR and PV populations received a stimulusselective input current l_i^s upon presentation of their preferred stimulus representing thalamic inputs. They received a fraction of this input current $(0.5 \cdot l_s)$ upon presentation of their non-preferred stimulus. Thus, the PYR and PV cell population activity corresponded to the average response of all cells with a given stimulus preference in the population. The SOM population also received a stimulus input current, which was the same for both presented stimuli. All populations received a constant baseline current input l_i^b . Each modulated population *i* (PYR, SOM, PV) received a multiplicative top-down modulation l_i^{TD} during attention (see Table S8). The VIP population did not receive top-down inputs in this model reflecting the lack of attentional modulation observed during the task (Figure S3).

To find parameter values consistent with the experimental observations, we used the python module sbi,⁵⁴ which implements neural inference algorithms for simulation-based inference. We used uniform prior distributions within the ranges shown in Table S7 from which parameter values were sampled. We ran circuit simulations to obtain training data consisting of the sampled parameter values and the output of the simulation for these values. This training data was then used to train a deep neural density estimator, provided by the sbi (v0.21.0) package (sequential neural posterior estimation – SNPE). The trained neural density estimator returns a posterior distribution $P(\theta|x)$ of the parameters θ given the desired model output *x*. To obtain the distribution of parameter values that are consistent with the desired model output (the experimental observations), we sampled parameter sets from the posterior. The observations we used to match the model were the PYR, PV, SST and VIP peak stimulus-evoked activity without the optogenetic light, the PYR, PV and SST peak stimulus-evoked activity with the optogenetic VIP activation, the increased PYR and PV activity to the preferred stimulus with attention, the decreased PYR and PV activity to the non-preferred stimulus with attention, and the increased

Neuron Article



PYR activity to both preferred and non-preferred stimuli with the VIP activation. For the final model, we took parameter values that had the maximum posterior probability. The inferred values of parameters for the baseline inputs (l_i^b) , the feedforward inputs (l_i^s) , and the modulatory top-down inputs (l_i^{TD}) can be seen in Table S8.

For most values of connection strengths, we used the relative weights reported by Pfeffer et al.¹⁷ We multiplied all weights with a factor γ =1.485, which we inferred with sbi. Hence, the relative weights match the ones reported by Pfeffer et al.,¹⁷ which were normalised such that W_{EP} was 1.00. We used simulation-based inference as described above to infer the parameter values of W_{EE}, W_{PE}, W_{SE} and W_{VE} which were norterported by Pfeffer et al.¹⁷ Since the SOM to PV connection also plays an important disinhibitory role, we adjusted the value for W_{PS} to match the observations in the data. In our model the value of W_{PS} was larger compared to the value reported by Pfeffer et al.¹⁷ (0.5 vs. 0.3). Critically, with the smaller value, it was not possible to match the observations from the data.

The final connections between cell types were as follows:

W =	W_{EE}	W_{EP}	W _{ES}	W_{EV}	=	(0.73	1.00	0.54	0.02
	W _{PE}	W _{PP}	W_{PS}	W_{PV}		0.57	1.01	0.50	0.02
	W _{SE}	W _{SP}	W _{SS}	W _{SV}		0.37	0.03	0.02	0.34
	W _{VE}	W_{VP}	W _{VS}	W _{vv})		0.26	0.22	0.77	0.02

We simulated the network without stimulus input for 4s until the neural activity for each cell class reached steady state. Then we presented the stimulus (either preferred or non-preferred) for 1.5s. In the attend condition, we introduced multiplicative top-down modulation 2s after the start of the stimulation (2s before the stimulus). Optogenetic light stimulation happened at the same time as stimulus presentation. During light stimulation, we set the activity of the VIP population to a fixed value such that the normalized response was 2.0, in the case of VIP activation or 0 in the case of VIP inactivation. The simulation time step was 0.2ms. τ_i with $i \in \{PYR, SOM, VIP, PV\}$ was 100ms. All responses were baseline-subtracted (baseline at 50ms before stimulus onset) and to ensure that the responses were comparable to fluorescence intensity measurements, the model responses were divided by a factor of 5.0.

To calculate the mean response for each cell type and each condition, we took the average of the response over the stimulus window. To calculate the selectivity of cell populations in the model, we subtracted the mean activity to the non-preferred stimulus \overline{x}_{N} from the mean activity to the preferred stimulus \overline{x}_{P} during the stimulus presentation:

 $SI = \overline{X}_P - \overline{X}_N$

To study the impact of W_{PE} and W_{PS} on how much the VIP activation increased PV activity, we set either connection to 0.0 and compared the mean activity of the PV population in response to their preferred stimulus in the three conditions (control, without W_{PE} , without W_{PS}).

To visualize the parameter distributions of the top-down modulation used in the final model that is consistent with the desired model output we fixed the baseline and feedforward inputs, varied the parameters W_{EE} , W_{PE} , W_{VE} , W_{PS} , γ and I_i^{TD} , ran 10,000 simulations on which we trained the neural density estimator (SNPE), and then sampled 50,000 parameter sets from the estimated posterior distribution.