Active filtering of sequences of neural activity by recurrent circuits of sensory cortex

Ciana E. Deveau*[1,2,3], Zhishang Zhou*[1], Paul K. LaFosse[1,2,4], Yanting Deng[1], Saghar Mirbagheri[5], Nicholas Steinmetz[5], Mark H. Histed[1]
* equal contributions

1 Intramural Program, National Institute of Mental Health, National Institutes of Health, Bethesda, MD USA, 2 NIH Graduate Partnership Program, Bethesda, MD USA, 3 Department of Neuroscience, Brown University, Providence RI USA, 4 Neuroscience and Cognitive Science Program, University of Maryland, College Park MD USA 5 Department of Biological Structure, University of Washington, Seattle, WA USA

Email: mark.histed@nih.gov, ciana.deveau@nih.gov

Abstract

In daily life, organisms interact with a sensory world that dynamically changes from moment to moment. Recurrent neural networks can generate dynamics, but in sensory cortex any dynamic role for the dense recurrent excitatory-excitatory network has been unclear. Here we show a new role for recurrent connections in mouse visual cortex: they support powerful dynamical computations, but via filtering sequences of input instead of generating sequences. Using two-photon optogenetics, we measure responses to natural images and play them back, showing amplification when played back during the correct movie dynamic context and suppression in the incorrect context. The sequence selectivity depends on a network mechanism: inputs to groups of cells produce responses in different local neurons, which interact with later inputs to change responses. We confirm this mechanism by designing sequences of inputs that are amplified or suppressed by the network. Together, these data suggest a novel function, sequence filtering, for recurrent connections in cerebral cortex.

Main text

A defining feature of all cerebral cortical areas is their extensive local recurrent connectivity. Excitatory cortical cells make up approximately 80% of the neurons in all cortical areas, and each excitatory cell receives hundreds or thousands of inputs (*1*). A majority of those input synapses come from other local neurons, within a few hundred microns distance (*2*, *3*). The resulting massive local excitatory-excitatory connectivity is a common feature of all cortical areas (Fig. 1A-D). In several cortical regions this dense recurrent connectivity is thought to generate ongoing or slow dynamics: for example, in motor cortex, dynamics associated with muscle movement, and in prefrontal cortex, delay dynamics associated with short-term memory (*4*, *5*).

In sensory areas, however, the role of recurrent connections has been less clear (Fig. 1E). In visual cortex, for example, responses to visual input end within a few tens of milliseconds after visual stimuli are extinguished (*6*), suggesting that sustained or generative responses are not produced by sensory recurrent connections (*7*).

Here we study the roles of these extensive recurrent connections in sensory cortical regions. Sensory stimuli activate thousands or more neurons in the brain, and thus many kinds of input — feedforward, feedback, and recurrent — arrive together to cortical neurons during visual stimulation. Thus, it has historically been difficult to isolate and study the influences of recurrent connections. To make progress, here we use two-photon holographic stimulation (Fig. 1F) (*8*, *9*), delivering patterns of input to cortical neural populations (Fig. 1G-J) and measuring how neurons in the recurrent cortical network respond.

We used a viral preparation we developed (*10*) to express an opsin and a calcium indicator in excitatory neurons in layer 2/3 of mouse V1, and used a stimulation laser and holographic device to deliver input to groups of chosen single neurons (Fig. 1F). With this approach, we deliver inputs that are precise spatially (Fig. 1G), and also precise temporally, with one input pattern changing to the next within tens of milliseconds (Fig. 1J; 30 ms pattern duration).

When animals are shown natural visual movies, many neurons in the visual cortex respond (Fig. 1H), and using two-photon stimulation we can simulate similar activity patterns by giving input to 20 or more neurons at once (Fig. 1I). This stimulation approach allows us to measure visual responses, select neurons that respond, and replay activity patterns, driving population activity in the network (Fig 1H-I).

Sequential order affects V1 neurons' responses through off-target suppression

If recurrent effects influence V1 dynamics, then changing the order of patterns of input might be expected to produce differences in neurons' responses. To test this, we chose patterns of neurons at random (three patterns; A, B, and C, each a fixed duration, 30 or 60 ms in different experiments) and presented these patterns in two different orders (ABC or CBA, Fig. 2A).

We find the sequential order of patterns does affect population responses (Fig. 2B-D). The pattern which is presented first produces a greater response than the same pattern when it is delivered later in the sequence (Fig. 2B-D).

One way that this dependence on input order might occur is if population responses to a given pattern could in some way affect later patterns. To study this, we separated the three stimulation patterns in time, presenting some patterns in sequence and inserting a delay of several seconds (7.9 sec) between others (Fig. 2E). When patterns A and B are presented in sequence (AB, Fig. 2E), we found striking off-target responses — some nonstimulated cells are activated, and some are suppressed. These off-target neural responses, in neurons that are not receiving stimulation input, are likely to be primarily due to local recurrent interactions, not due to axonal or dendrite activation, as here we used a somatically-targeted opsin (stChRimsonR, (*10*)).

These off-target responses can occur in neurons that will later be stimulated in the C pattern (Fig. 2G). Thus, we asked whether overlap of off-target responses of one input with later inputs from other patterns could change responses.

Off-target responses might interact in an additive, or linear, way with later inputs, or there could be nonlinear interactions that could produce complex sequential responses. In prior work we have found attenuation-by-suppression (*11*), where V1 neurons, when suppressed, produce smaller (sublinear) responses to input. To understand if sequence selectivity could depend on such nonlinearities, we examined off-target responses in C cells generated by the AB sequence. We compared the ABC response to a prediction computed by summing the C response and the AB response. We found, as we saw previously, that excited off-target cells produced linear responses when they were stimulated next in sequence, well-predicted by the sum of off-target responses to AB and C (Fig. 2H, Supp. Fig. 2E). However, suppressed off-target cells showed nonlinear responses (Fig. 2I, Supp. Fig. 2E) — the response of the C pattern embedded in the full sequence was attenuated compared to the response to the C pattern alone.

Thus, these data suggest a mechanism for generating sequence selectivity: Early patterns cause off-target effects, likely due to the local recurrent network, that interact with later inputs. This interaction creates sequence responses that are not merely sums of sequential input (Fig. 2J, Supp. Fig. 2E) but are nonlinearly transformed.

Sequence components corresponding to natural visual inputs are amplified

With confirmation that V1 neurons' responses are sensitive to input pattern sequence, we next asked how this might affect vision. One of the most common kinds of sequential input that the visual cortex receives comes in response to natural motion. For example, when we see a person or animal running, the changing visual input creates changing patterns of input to the cortex. In principle, the cortical network may have, in its recurrent connections, information about the temporal structure of the natural visual world, which allows the network to preferentially respond to sequences corresponding to natural vision.

In that case, sequences of input arising from natural vision would be expected to produce different responses in the network than sequences that do not correspond to natural vision. Then earlier parts of a dynamic natural stimulus should influence responses to input arriving at later times. We tested this hypothesis using two-photon stimulation to mimic responses from one part of a natural movie while we changed the prior visual context.

We first measured responses to a single frame taken from a movie (Fig. 3A). We constructed a 2-photon input pattern from the cells activated by that single frame. We then showed animals the movie and replaced that single frame with the 2-photon input pattern, either at the correct time in the movie (Fig. 3A: matched context), or at a different time, when the preceding movie frames were not matched to the stimulation pattern (Fig. 3A: unmatched context).

If V1 responses were dependent on the dynamic context of natural vision, we would expect that neurons would produce different responses to the same, fixed input pattern when presented in different contexts. That is what we found (Fig. 3B-E). The input pattern produced a larger

response in the sequential context matched to the visual movie than when presented in the unmatched context. Responses were not just amplified in the matched context, but attenuated in the unmatched context (Fig. 3CD), relative to the response to the frame when it was not preceded by a dynamic visual stimulus.

Above, we showed that one stimulation pattern can interact with a next pattern to produce nonlinear responses (Fig. 2H; using stimulation patterns that overlap in time; Methods). Therefore, to determine if this type of mechanism could explain the effects of visual context, we asked how neurons' firing was affected by the visual movie segment comprising the matched and unmatched contexts. Consistent with a nonlinear response generated by off-target suppression, we found that the matched context indeed excited more neurons that would be stimulated than the unmatched context did (Supp. Fig. 3A) — and conversely, the unmatched context generated more suppressed cells.

For a recurrent network interaction to modify responses to dynamic visual input, it would also seem to be necessary for the response to one visual stimulus (movie frame) to be sustained long enough for responses to later frames to arrive. Thus, to determine if single-frame responses were sustained for enough time to allow overlap between responses, we turned to Neuropixels electrophysiological recordings to measure response timecourses. We recorded V1 neuron responses to flashed single image frames (as in Fig. 3A), and sorted neurons into groups based on whether they showed activated or sustained responses. We found that suppressed responses were sustained for several hundred milliseconds (225ms, Fig. 3F) after a single image frame, long enough to interact with inputs from later frames. For comparison, a movie at 25 frames per second, a rate comfortably within the range where humans see smooth motion, has a 40 ms frame time. The mouse visual system likely supports even faster processing than the human, and yet here we found that natural images flashed for short times produced responses 75 ms long for activated neurons (Fig. 3F) and 225 ms for suppressed cells (3F, inset). It is possible that the one frame's response could be shorter during dynamic natural vision due to network state. Still, these data suggest that the dynamics of input during natural vision is sufficient to allow overlap between patterns. This overlap would allow attenuating or filtering responses to later inputs based on responses induced by earlier inputs (Fig. 2H-J).

Together, these observations show that patterns of input arising from frames of a natural movie produce amplified responses in the correct natural visual context. This suggests that sequences of input corresponding to dynamic natural vision are amplified — boosted relative to other sequences which are not associated with visual inputs found in the world.

We confirm this mechanism can explain sequence filtering by constructing amplified and suppressed sequences

We have observed that the visual cortex is sensitive to the temporal order of patterns of input (Figs. 2, 3). We also find that single input patterns generate off-target responses, which then interact with later patterns to change responses to those later patterns (Fig. 2E-J, Supp. Fig. 2).

To confirm that this network mechanism can explain sequence filtering, we constructed sequences that should, in this way, be either amplified or suppressed.

To construct such sequences (Fig. 4A), we measured responses to single input patterns. We then created later stimulation patterns based on the off-target responses to those single patterns. We created 'amplified' sequences by adding a pattern of stimulated cells that were activated by the first pattern, and 'attenuated' sequences by adding a pattern of stimulated cells suppressed by the first pattern (Supp. Fig. 4D-G). We repeated this step a second time to create amplified and attenuated sequences that were three patterns long.

We then asked whether the putative amplified and attenuated sequences created different responses. Confirming our hypothesis, we found they generated markedly different responses. The sequence designed to be amplified produced a much larger response than the sequence designed to be attenuated (Fig. 4B-D).

Interestingly, the amplified patterns produced responses that were sparser than suppressed sequences — a small number of outlier cells produced large responses that elevated the mean (Fig. 4CE). This sparsity is a hallmark of responses to natural movies (*12*–*14*), and supports the idea that the cortical network can amplify sequences of responses associated with natural vision.

In sum, the responses to our designed sequences support the idea that recurrent interactions across time amplify or attenuate sequences via off-target interactions. Responses to an earlier pattern in a sequence influence responses to later patterns.

Discussion

Here, we find a role for the dense recurrent network in layer 2/3 of sensory (visual) cortex: to amplify some sequences of input and attenuate others. V1 does not generate complex sustained or ongoing dynamics (e.g. after the end of a visual or optogenetic input, Fig. 3; (*15*)). But we find the V1 network does in fact support complex temporal processing, via filtering input sequences.

During spontaneous activity, neurons in the cortex show activity that fluctuates from moment to moment (*16*–*18*) with measurable, but weak, connection to natural visual responses (*19*). These past studies, combined with our data, suggest that while sequences of neural responses during spontaneous activity states are largely suppressed by the cortical network, some spontaneous patterns, which partially match natural patterns, are weakly amplified by the recurrent network into an echo of natural visual responses.

The sequence filtering effect we observe would have been hard to identify by studying responses to visual inputs alone, without the external input provided by optogenetics. Because visual stimuli change neural activity in thousands or more neurons across the brain, it is hard to know from visual responses whether neural activity is driven by recurrent inputs or is principally shaped by feedforward or feedback inputs. Here, by controlling inputs to neurons in vivo with 2p stimulation, we are able to identify recurrent influences and reveal these sequence amplification effects.

For the network mechanism we identify, it is the details of the recurrent off-target responses, presumably created by recurrent connectivity, that act to choose which sequences to amplify. That is, one pattern of input falls on a set of neurons, and those neurons' local connections lead to activation of some other local neurons, and suppression of others. This interaction between one input pattern to the next is in principle enough to create selectivity for longer sequences. Interaction between neighboring states leading to complex sequences is analogous to what is seen in Markov, first-order autoregressive, or renewal processes — which can produce complex with transitions or events depending only on the previous event, without dependence on the full history of events.

Because many thousands of cortical neurons can fire differently in each sequence, the patterns of population responses to visual input are high-dimensional. That is, inputs can fall on the set of neurons in L2/3 in many different ways, and the potential off-target patterns can also occur in many different ways. To support such high-dimensional responses would seem to require a large number of synapses. The dense excitatory-excitatory connectivity in the cortex seems well-placed to support the transformations required for sequence filtering. Our results suggest that the numerous excitatory recurrent connections in the cortex gives the cortex great capacity to selectively process the high-dimensional space of natural visual inputs (*20*).

Our data are consistent with the idea that the recurrent network in V1 learns the temporal statistics of the visual sensory world by boosting patterns of input that correspond to natural vision. While we have made these measurements in visual cortex, we might speculate that recurrent connections in other sensory regions — auditory, somatosensory, etc — also are used to process sensory inputs. Those sensory modalities also process natural inputs that vary in time, though often with faster changes than in vision. Thus, we might predict that tuning for faster sequences will be seen in other regions. And beyond biology, recurrent networks in artificial systems also often are used to create temporally-structured computations. Recurrent artificial networks can learn temporal statistical structure (*21*), and transformers like ChatGPT (*22*) have been used to generate highly complex natural language sequences. Thus, densely-connected recurrent networks seem to be useful for sequence processing both in artificial systems and in biological brains.

In summary, these data show a new and powerful purpose for recurrent connectivity in the sensory cerebral cortex: to confer sensitivity to sequential input. The visual cortical L2/3 network is sensitive to dynamic visual context, boosting responses to sequences of input corresponding to dynamic natural sensation, and suppressing others.

Figures

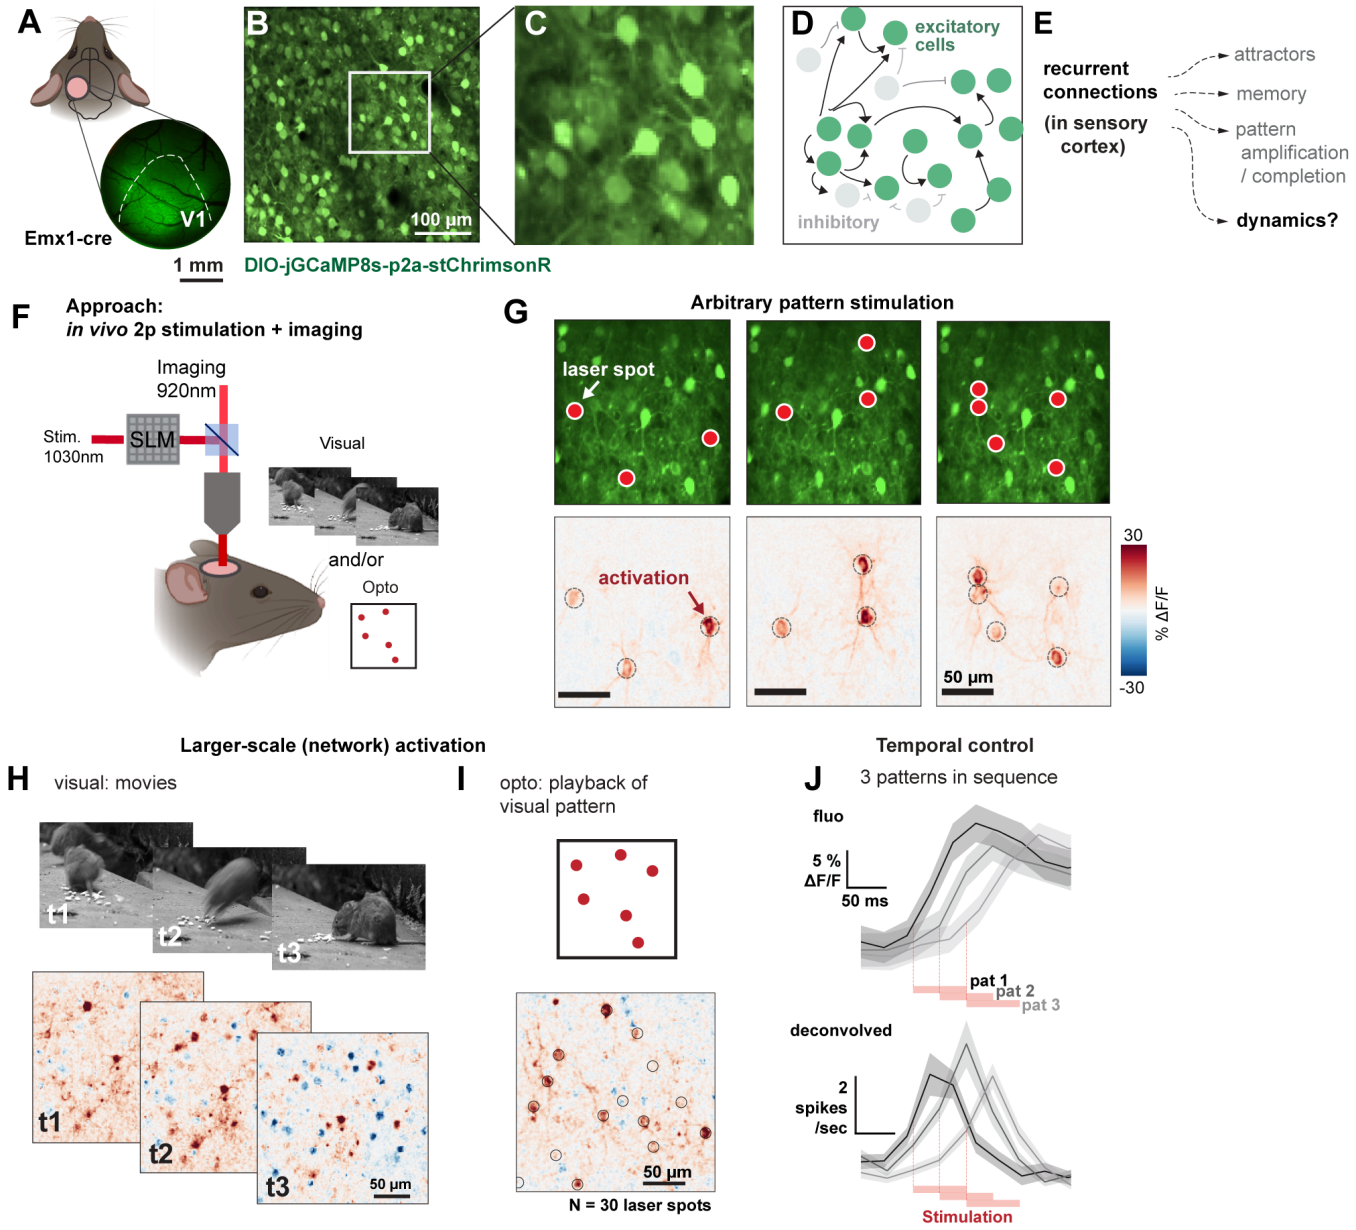
A
V1
Emx1-cre
1 mm
B
100 µm
C
DIO-jGCaMP8s-p2a-stChrimsonR
D
excitatory cells
inhibitory
E
recurrent connections (in sensory cortex)
attractors
memory
pattern amplification / completion
dynamics?
F
Approach:
in vivo 2p stimulation + imaging
Imaging 920nm
Stim. 1030nm
SLM
Visual
and/or
Opto
G
Arbitrary pattern stimulation
laser spot
activation
50 µm
30
-30
% ΔF/F
Larger-scale (network) activation
H
visual: movies
t1
t2
t3
50 µm
I
opto: playback of visual pattern
50 µm
N = 30 laser spots
Temporal control
J
3 patterns in sequence
fluo
5 % ΔF/F
50 ms
pat 1
pat 2
pat 3
deconvolved
2 spikes /sec
Stimulation

Figure 1: Using two-photon holographic stimulation and imaging to probe the role of excitatory recurrent connections. (A) 3mm window surgical implant over V1. (**B)** FOV image of excitatory cells expressing both GCaMP8s and opsin stChrimsonR (via AAV-DIO-jGCaMP8s-p2a-stChrimsonR in Emx1-Cre line). **(C)** Enlarged region of (B). **(D)** Schematic of recurrent connections between cortical neurons. **(E)** Possible functions for cortical recurrent connections in V1. **(F)** Schematic of two-photon imaging and stimulation setup, allowing pairing of patterned optogenetic and visual input. **(G)** Example of spatial precision of 2p stimulation; left to right: three different patterns, with 3, 3, and 5 laser spots each targeted to one neuron. **(H)** Network activation from natural movie input at three distinct timepoints. **(I)** Achieving similar network activation with optogenetic patterns of 30 laser spots. **(J)** Example of temporal precision of sequential 2p stimulation: 3 patterns, 15 cells each, 60 ms stim per pattern with 30 ms between pattern onset. Above: fluorescence traces due to calcium change, below: corresponding deconvolved traces using OASIS (Methods). These traces show that even with some overlap, imaging with stimulation every 30 ms can resolve distinct temporal peaks.

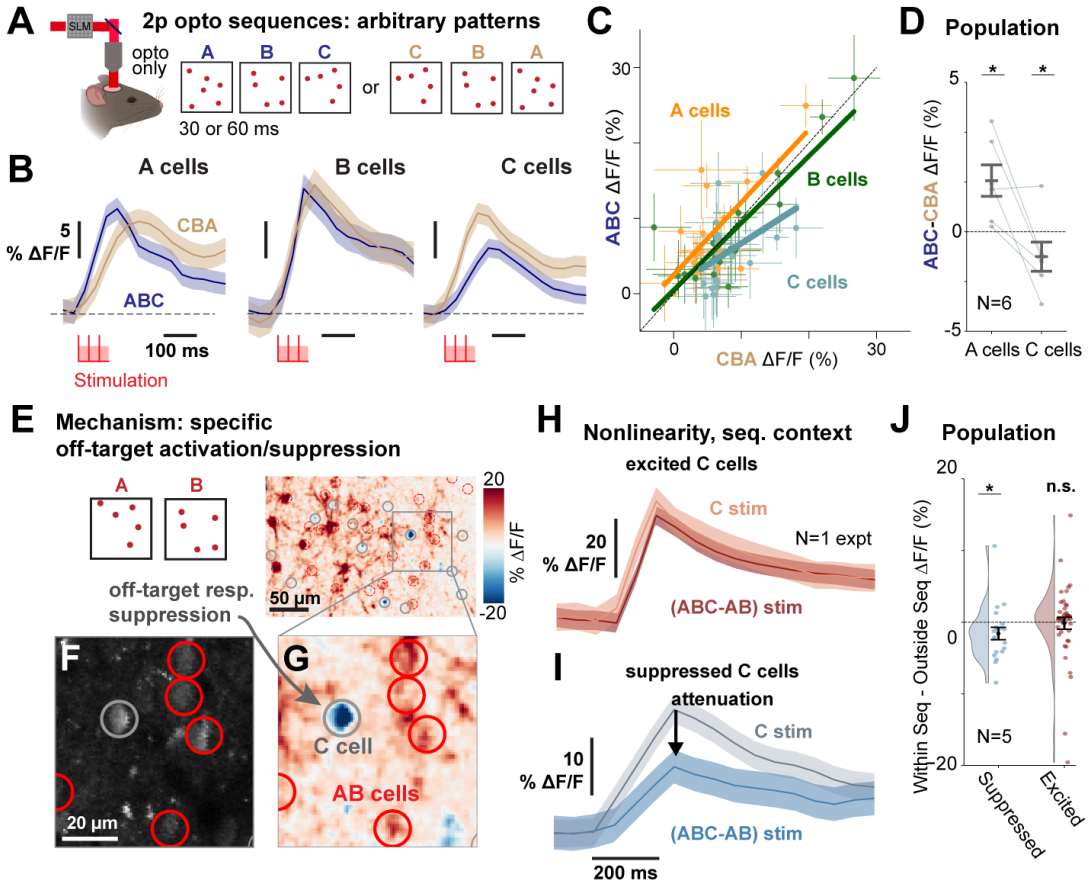

Figure 2: Patterns of input produce different responses in the V1 network depending on sequential context (**A**) Experiment schematic: stimulation of three distinct patterns (e.g. N=30 cells per pattern, cells chosen at random for each pattern). Two sequences: ABC and reverse CBA order. (**B**) Responses depend on sequential order. Mean responses averaged across cells of A (left), B (middle), and C (right) cells in the ABC (blue) or CBA (beige) trials (N=1 expt, pattern duration: 30 ms). (**C**) Points: average response (across trials) of individual neurons, same data as in B; change in slope for A, B, C indicate differences in response based on sequential order. (**D**) Population data (N=6 expts). gray lines: ABC - CBA responses of A (left) and C (right) cells; averaged over A, C cells per experiment. Error bars: SEM. (t-test, A cells: $p<0.0001$, C cells: $p=0.001$). (**E-G**) Separating A, B, and C patterns in time reveals off-target interactions. (**E**) Responses to AB sequence (schematic, left, GCaMP responses, right. Red circles: stimulated (AB) cells. Gray circle: C cells, here unstimulated. (**F**) Gray image: field of view. (**G**) Zoomed view shows suppression in unstimulated neuron (gray circle). (**H**) Linear summation in excited off-target cells (N=5 cells). Light red: response to C stimulus alone. Dark red: prediction, response to ABC sequence minus response to AB sequence; response and prediction overlap (**I**) Sublinear summation in suppressed off-target cells (N=5 cells): Same conventions as (H) but for cells suppressed by prior patterns (AB). Response to C within the ABC sequence (ABC-AB, blue) is smaller than resp. to C alone (gray; t-test, 300 ms period after stim onset, $p=0.013$). (**J**) Nonlinear responses to suppressed cells across experiments (N=5; N=62 C cells). Average single cell response outside of the sequence (C) subtracted from response within the sequence (ABC-AB). Blue: suppressed cells (response to AB stim<0; t-test, $p = 0.04$). Red: excited cells (response to AB stim>0; t-test, $p=0.86$). Errorbars: SEM.

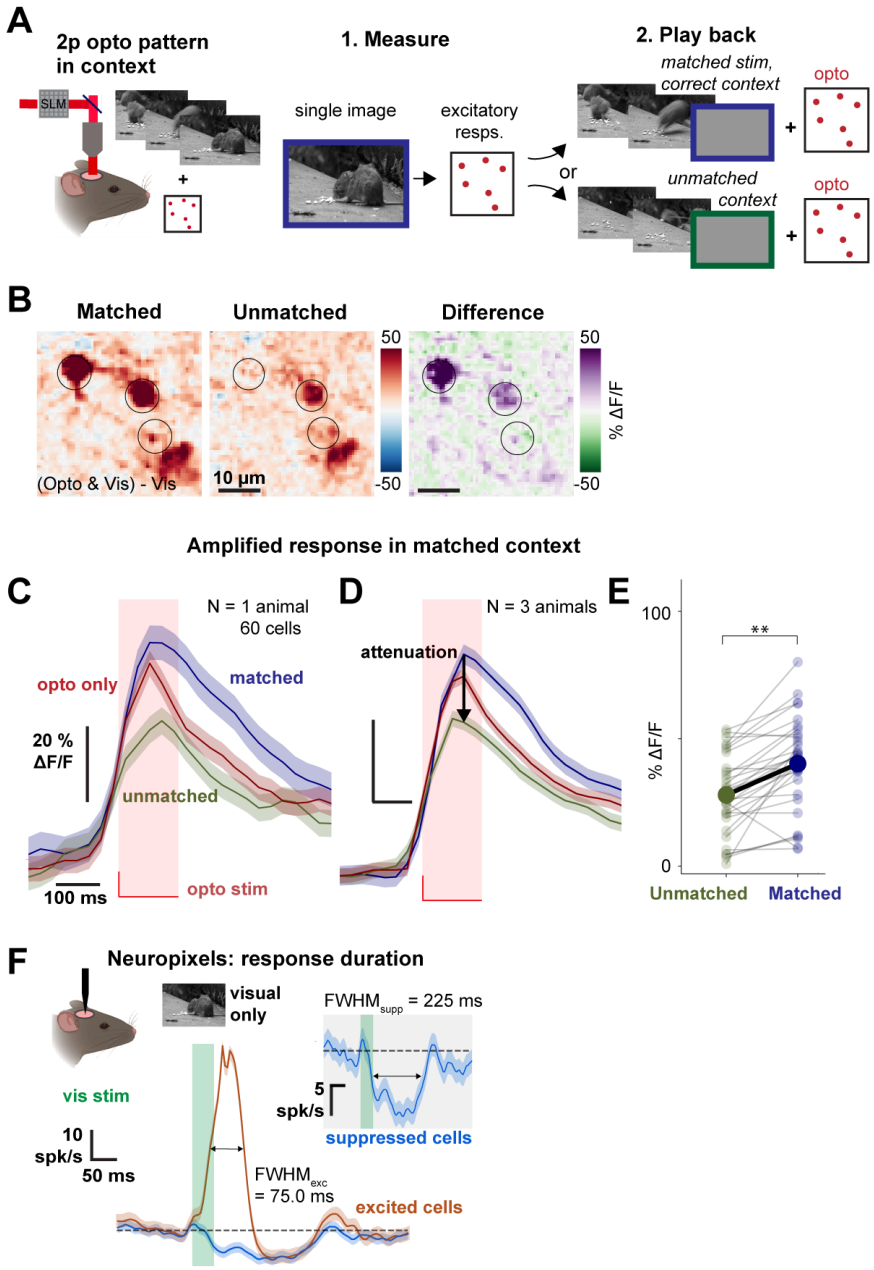

Figure 3: Strength of response to an input pattern is dependent on the dynamic context of visual input. (**A**) Experimental schematic: Responses to a single frame from a natural movie are recorded; a stimulation pattern is formed by selecting excited cells. That pattern is stimulated during visual movie playback at the correct time (matched context) or an incorrect time (unmatched). (**B**) Example responses in matched and unmatched contexts. Black circles: stimulated cells. Responses in matched context are larger. (**C**) Mean responses averaged across cells in matched (blue), unmatched (green), and opto only (red) trials (N=1 expt, pattern duration: 120ms). (**D**) Population time courses, N=3 animals, same conventions as (C). Responses are larger in matched context. (**E**) Points are cell responses in matched (blue) and unmatched (green) contexts from data in D. Mean optogenetic responses are significantly larger in the matched context (t-test, p=0.004). (**F**) Electrophysiological responses to a flashed natural image show responses sustained for up to hundreds of milliseconds. Excited cells (orange; errorbar, SEM, largely hidden by mean). Suppressed cells (blue). FWHM: Full width at half max. Error bars: SEM.

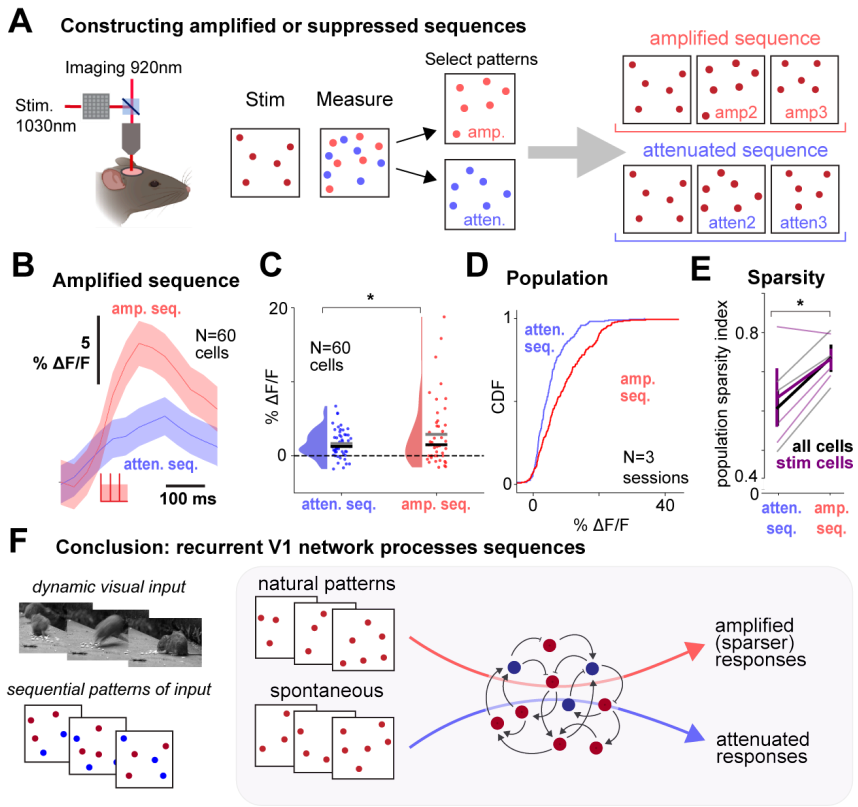

Figure 4: Sequences designed to be amplified based on network interactions produce stronger and sparser responses than suppressed sequences. (**A**) Experimental schematic: sequences that should be amplified or attenuated are constructed by imaging responses to stimulation and selecting cells based on response to form the next pattern in the sequence. For example, amplified sequence (top): we stimulate with a random pattern A (left), select excited neurons to create pattern B, then stimulate AB and select excited neurons to create pattern C. (**B**) Mean responses of amplified (pink) and suppressed (blue) sequences (N=20 cells per pattern, 3 patterns in the sequence; 30 ms pattern duration). (**C**) Distribution of responses of stimulated cells. Mean (gray horizontal line), median (black), means are sig. different (t-test, p=0.033; N=60 cells from N=3 patterns). (**D**) Distribution of amplified responses. (**E**) Population sparsity is also greater for amplified patterns. Each point is an average across neurons, all cells (gray) significantly sparser in amplified sequences (t-test, p=0.011; N=3 animals). (**F**) Mechanism schematic: the recurrent network filters sequential inputs, amplifying some and attenuating others. Error bars: SEM.

Methods

Animals

All experiments were conducted in adherence to NIH and Department of Health and Human Services (HHS) guidelines for animal research and were approved by the Institutional Animal Care and Use Committee (IACUC) at the relevant institutions (optogenetic experiments: NIMH; Neuropixels experiments: University of Washington.)

Optogenetic experiments

Emx1-Cre animals (N=11, https://www.jax.org/strain/005628 (*24*)) of both sexes were used. No systematic differences were observed between males and females. Viral injection and window implants were done at ages 2-7 months. After procedures, animals were singly housed on a reverse 12-hour dark/light cycle.

Neuropixels experiments

Male C57BL/6J mice (N=2), aged two to three months, were used. After procedures, animals were singly housed on a standard 12-hour dark/light cycle.

All animals were put on water schedule five or more days after head-plate surgery, and their weights were carefully monitored to ensure they remained above 85% of their baseline body weight.

Viral injection and cranial window implants

We performed injections and implants as described in (*10*) (optogenetic experiments; NIMH) or (*25*) (2023; Neuropixels; University of Washington). Minor differences in the two sets of approaches are not expected to impact results.

Optogenetic experiments

Briefly, mice were given dexamethasone (3.2 mg/kg) 30 minutes before surgery and anesthetized with 1-3% isoflurane (in 100% O_2). A titanium headplate was affixed using C&B Metabond (Parkell) and a 3 mm craniotomy was made over the primary visual cortex (-3.1 mm ML, +1.5 mm AP from lambda) and a glass optical window chronically inserted. Mice were injected with either a mixture of two AAV viruses to induce GCaMP and stChrimsonR into the cells (mixed virus injections), or a bicistronic virus expressing both GCaMP and stChrimsonR in each cell (*10*). For the mixed injection, the fluorescent calcium indicator, (AAV9-syn-jGCaMP8s; titers: 5.0-10x10^{12} genome copies (GC)/ml) and the soma-targeted opsin (AAV9-hsyn-DIO-stChrimson-mRuby2; titer: 3.0x10^{12} GC/ml) were mixed in phosphate-buffered saline. In the case of the bicistronic virus injections, we used AAV9-hsyn-DIO-jGCaMP8s-p2a-stChrimson (titers: 2.9-5.9x10^{12} GC/mL). For each mouse, 3 to 5 injections (100 nL/min, 200 μm depth) were made to cover a wider area of the cortex. A custom made light-blocking cap was fixed onto the implant to limit ambient light exposure and prevent debris from contacting the window. Experiments began three weeks after these procedures.

Neuropixel experiments

Briefly, mice were induced into anesthesia with 5% isoflurane, and maintained at 2-3%. Carprofen (5mg/kg) and Lidocaine (2mg/kg) were administered for analgesia. A titanium head plate and a 3D-printed chamber were affixed using Metabond. Carprofen was administered at 0.05 mg/ml in the water post-surgery for three days. Five days following recovery, mice were put on water restriction. After two days, they were habituated to head fixation for two days, during which they received up to thirty random 5μL sucrose water 10% rewards. Subsequently, animals were exposed to a set of 20 natural images presented over 1000 trials, followed by a 5μL sucrose water reward for five days. This step was performed to provide a control training condition for another experiment. None of the 20 images in this step were used as part of the experiment mentioned in this work. One or two days before recordings, a 2 x 2 mm craniotomy was performed over the primary visual cortex using a dental drill and stereotaxic techniques. The area was then covered with silicone gel. This process was conducted under anesthesia procedures used for the head plate implant, and Carprofen (5mg/kg) was given for analgesia. A cap was placed on the chamber to prevent debris from getting into the area around the craniotomy.

Retinotopic mapping

Before optogenetic experiments, we determined the location of V1 in the cranial window using a hemodynamic intrinsic imaging protocol previously described in (*26*). Briefly, small visual stimuli were presented while 530 nm light was delivered using a fiber-coupled LED (M350F2; Thorlabs, Newton NJ). Hemodynamic response was calculated as the change in reflectance of the cortical surface between the baseline period and a response window starting 3 ms after stimulus onset. Imaging was done on a Zeiss Discovery stereo microscope with a 1x widefield objective through a green long-pass emission filter, acquired at 2Hz. An average retinotopic map was fit to the cortical responses based on the centroids of the hemodynamic response for each stimulus location.

Two-photon holographic imaging and stimulation

Two photon imaging and stimulation procedures are described in detail in (*10*). Briefly, animals were awake and alert under a 16x water-immersion objective (Nikon; Tokyo, Japan; NA=0.8) manually positioned over the implanted optical window. Imaging was done with a custom-built microscope and controlled by ScanImage software in MATLAB (The Mathworks, Natick, MA). Calcium responses were measured ~100-200 μm below the surface of the pia (L2/3 of V1) with an imaging field of view of 414 x 414 μm. Imaging was performed using 920 nm wavelength light at 15-20 mW and frames acquired at 30 Hz.

Holographic stimulation was performed using a femtosecond pulsed laser (1030nm, Satsuma, Amplitude Laser, or 1040nm, Monaco, Coherent Inc.) A spatial light modulator (SLM) shaped the laser wavefront to create stimulation patterns (10 μm diameter disks; 30, 60, 90, or 120 ms; 8-16mW/target, 500kHz pulse rate). The laser was gated on during horizontal flyback periods and off during the imaging pixel acquisition to allow for approximately simultaneous stimulation and imaging. Reported stimulation power is the average power over on and off periods (i.e. reduced by a factor of 0.3 from the laser power measured without this gating.)

Visual responses were measured in awake, head-fixed mice viewing flashed single natural image frames (40° circular mask with neutral gray background; 120 ms) or full natural movie stimuli (full-field, 2 seconds) presented on an LCD monitor. Mice were given small volumes of water on 20% of trials (3μL).

Two-photon data analysis

We performed motion correction using the CaImAn toolbox (*27*) and cell segmentation using Suite2p (*28*). All data analysis was done in Python (https://www.python.org). For pixel based analyses, we computed $\Delta F/F_0$ ($\Delta F/F0$; F: raw fluorescence intensity; F0: average fluorescence across the 45 imaging timepoints, 1.5 s, prior to stimulus presentation) at every pixel of the image stacks for display across the FOV. The $\Delta F/F_0$ for all time courses was calculated from fluorescent traces output from Suite2p. For full sequence trials (ABC, CBA), averaging windows for each pattern (A, B, C) started at stimulation onset of that pattern and lasted 300 ms; shifted for pattern order so the response window always began on the first frame of stimulation of that pattern (raw time course data without shift is shown in Fig. 2B). Activity of the C pattern in ABC sequence was calculated as [ABC - AB]. To isolate the optogenetic response following visual input, we subtracted off the response to the visual stimulus alone, [(vis+opto)-vis] (Fig. 3). Averaged response interval for individual cells began at optogenetic stimulus onset and ended at stimulus offset (120 ms, Fig. 3E). In Figure 4, the averaging interval for individual cell responses started at stimulus onset and lasted 200 ms (Fig. 4D-E). Response intervals for all pixel based analyses started at stimulus onset and lasted 300 ms (Fig. 2E, 3B). Deconvolution (Fig. 1J) was done using OASIS with an autoregressive constant of 1 (*29*).

Cell pattern selection

For analysis of visual response to flashed natural images, individual cell responses were averaged over a 300 ms interval starting from the stimulus onset. Cells activated by the visual stimulus were defined by responses exceeding a 5% $\Delta F/F_0$ threshold. All cells that met this threshold were included in the optogenetic pattern (11-15 cells per pattern, Supp. Fig. 3EF).

To construct the amplified and attenuated patterns (Fig. 4), we first selected at random (not based on visual responses) the first 15-20 cell pattern in the sequence. Average $\Delta F/F_0$ response to stimulation of that pattern was calculated across a 300ms time interval starting at stimulation onset. Excited off-target cells meeting a minimum threshold of 5% $\Delta F/F_0$ were included in the next pattern for the amplified sequence (15-20 cells per pattern). Conversely, suppressed off-target cell responses below -5% $\Delta F/F_0$ (15-20 cells per pattern) were selected for the attenuated sequence pattern. If more than 20 cells met these criteria, the 20 cells with the largest magnitude response were used. This process was repeated for the two sets of first and second patterns to generate the third pattern for each sequence. Each pattern in both sequences had an equal number of cells. Any cells included in previous patterns were excluded from selection for the subsequent patterns (Supp. Fig. 4D-H).

Electrophysiology

For Neuropixels electrophysiology the mice were head-fixed for the recordings while seated in plastic support with forepaws on a wheel (*30*), with the ability to move body parts other than the head but no ability to locomote. The behavioral state of the animals was monitored using two Basler acA2440-75um cameras at 560x560 resolution. One camera focused on the eye area was used to monitor the pupil diameter, and the other focused on the face was used to detect licks and ensure the animal was not in distress. Animals were awake during the duration of the experiment. Recordings were conducted from all layers of the primary visual cortex using Neuropixel 2.4 probes (*31*) with a sampling rate of 30 kHz. The recording sites were targeted based on the stereotaxic coordinates in the CCF Allen atlas using the Pinpoint system (*25*). The reference for the recordings was set at the tip of the electrode. Three recordings were performed: one from the first animal and two (on subsequent days) from the second animal. Spike sorting was executed using Kilosort 2.5 (https://github.com/MouseLand/Kilosort; RRID: SCR_016422; (*31*)) on each of the four shanks separately.

Visual stimuli were presented on three 60 Hz screens (LG LP097QX1), surrounding the mice and covering 270 x 70 degrees (azimuth x elevation) visual angle. Each trial featured a natural image from a set of two, sourced from the Allen Institute Brain Observatory, and repeated across all screens. These images were equiluminant and underwent histogram equalization for contrast consistency. Images were warped to appear rectilinear from the animals' viewpoint. Each image was displayed 30 times for 50 ms. A reward of 5 μL of 10% sucrose water was administered at 1150 ms post-stimulus offset to maintain engagement. The trials were randomly interleaved with each other and with additional trials from another experiment, which also included natural image presentations and sucrose rewards. The inter-trial interval followed an exponential distribution with a minimum and mean of 2 and 2.6 seconds, respectively.

All electrophysiology analysis was done in Python. Baseline value computed in a 40 ms window ending 5 ms prior to stimulus presentation. Response value was calculated across a 50 ms interval beginning 5 ms after stimulus presentation. Excited cells are those with average response greater than baseline; suppressed cells those with response less than baseline.

References

1. V. Braitenberg, A. Schüz, *Anatomy of the Cortex: Statistics and Geometry* (Springer Science & Business Media, 2013).

2. Y.-J. Liu, M. U. Ehrengruber, M. Negwer, H.-J. Shao, A. H. Cetin, D. C. Lyon, Tracing inputs to inhibitory or excitatory neurons of mouse and cat visual cortex with a targeted rabies virus. *Curr. Biol.* **23**, 1746–1755 (2013).

3. R. J. Douglas, C. Koch, M. Mahowald, K. A. Martin, H. H. Suarez, Recurrent excitation in neocortical circuits. *Science*. **269**, 981–985 (1995).

4. D. Sussillo, M. M. Churchland, M. T. Kaufman, K. V. Shenoy, A neural network that finds a naturalistic solution for the production of muscle activity. *Nat. Neurosci.* **18**, 1025–1033 (2015).

5. S. Fujisawa, A. Amarasingham, M. T. Harrison, G. Buzsáki, Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat. Neurosci.* **11**, 823–833 (2008).

6. D. H. Hubel, T. N. Wiesel, Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* **160**, 106–154 (1962).

7. K. Reinhold, A. D. Lien, M. Scanziani, Distinct recurrent versus afferent dynamics in cortical visual processing. *Nat. Neurosci.* **18**, 1789–1797 (2015).

8. V. Emiliani, A. E. Cohen, K. Deisseroth, M. Häusser, All-Optical Interrogation of Neural Circuits. *J. Neurosci.* **35**, 13917–13926 (2015).

9. A. M. Packer, D. S. Peterka, J. J. Hirtz, R. Prakash, K. Deisseroth, R. Yuste, Two-photon optogenetics of dendritic spines and neural circuits. *Nat. Methods*. **9**, 1202–1205 (2012).

10. P. K. LaFosse, Z. Zhou, N. G. Friedman, Y. Deng, A. J. Li, B. Akitake, M. H. Histed, Bicistronic Expression of a High-Performance Calcium Indicator and Opsin for All-Optical Stimulation and Imaging at Cellular Resolution. *eNeuro*. **10** (2023), doi:10.1523/ENEURO.0378-22.2023.

11. P. K. LaFosse, Z. Zhou, V. M. Scott, Y. Deng, M. H. Histed, Single cell optogenetics reveals attenuation-by-suppression in visual cortical neurons. *bioRxiv* (2023), doi:10.1101/2023.09.13.557650.

12. W. E. Vinje, J. L. Gallant, Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*. **287**, 1273–1276 (2000).

13. E. Froudarakis, P. Berens, A. S. Ecker, R. J. Cotton, F. H. Sinz, D. Yatsenko, P. Saggau, M. Bethge, A. S. Tolias, Population code in mouse V1 facilitates readout of natural scenes through increased sparseness. *Nat. Neurosci.* **17**, 851–857 (2014).

14. Sparse coding of sensory inputs. *Curr. Opin. Neurobiol.* **14**, 481–487 (2004).

15. J. F. O'Rawe, Z. Zhou, A. J. Li, P. K. LaFosse, H. C. Goldbach, M. H. Histed, Excitation creates a distributed pattern of cortical suppression due to varied recurrent input. *Neuron* (2023), doi:10.1016/j.neuron.2023.09.010.

16. M. N. Shadlen, W. T. Newsome, The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* **18**, 3870–3896 (1998).

17. C. van Vreeswijk, H. Sompolinsky, Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*. **274**, 1724–1726 (1996).

18. N. Brunel, Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comput. Neurosci.* **8**, 183–208 (2000).

19. T. Kenet, D. Bibitchkov, M. Tsodyks, A. Grinvald, A. Arieli, Spontaneously emerging cortical representations of visual attributes. *Nature*. **425**, 954–956 (2003).

20. C. Stringer, M. Pachitariu, N. Steinmetz, M. Carandini, K. D. Harris, High-dimensional geometry of population responses in visual cortex. *Nature*. **571**, 361–365 (2019).

21. A. Graves, Generating Sequences With Recurrent Neural Networks. *arXiv [cs.NE]* (2013), (available at http://arxiv.org/abs/1308.0850).

22. Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, Dario Amodei, Learning Models are Few-Shot Learners. *Neural Information Processing Systems (NeurIPS)* (2020), doi:10.48550/arXiv.2005.14165 .

23. M. Abeles, *Corticonics: Neural Circuits of the Cerebral Cortex* (Cambridge University Press, 1991).

24. J. A. Gorski, T. Talley, M. Qiu, L. Puelles, J. L. R. Rubenstein, K. R. Jones, Cortical excitatory neurons and glia, but not GABAergic neurons, are produced in the Emx1-expressing lineage. *J. Neurosci.* **22**, 6309–6314 (2002).

25. Z. Ye, M. S. Bull, A. Li, D. Birman, T. L. Daigle, B. Tasic, H. Zeng, N. A. Steinmetz, Brain-wide topographic coordination of traveling spiral waves. *bioRxiv* (2023), p. 2023.12.07.570517.

26. H. C. Goldbach, B. Akitake, C. E. Leedy, M. H. Histed, Performance in even a simple perceptual task depends on mouse secondary visual areas. *Elife*. **10** (2021), doi:10.7554/eLife.62156.

27. A. Giovannucci, J. Friedrich, P. Gunn, J. Kalfon, B. L. Brown, S. A. Koay, J. Taxidis, F. Najafi, J. L. Gauthier, P. Zhou, B. S. Khakh, D. W. Tank, D. B. Chklovskii, E. A. Pnevmatikakis, CaImAn an open source tool for scalable calcium imaging data analysis. *Elife*. **8** (2019), doi:10.7554/eLife.38173.

28. M. Pachitariu, C. Stringer, M. Dipoppa, S. Schröder, L. F. Rossi, H. Dalgleish, M. Carandini, K. D. Harris, Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *bioRxiv* (2017), p. 30.

29. J. Friedrich, P. Zhou, L. Paninski, Fast online deconvolution of calcium imaging data. *PLoS*

Comput. Biol. **13**, e1005423 (2017).

30. C. P. Burgess, A. Lak, N. A. Steinmetz, P. Zatka-Haas, C. Bai Reddy, E. A. K. Jacobs, J. F. Linden, J. J. Paton, A. Ranson, S. Schröder, S. Soares, M. J. Wells, L. E. Wool, K. D. Harris, M. Carandini, High-Yield Methods for Accurate Two-Alternative Visual Psychophysics in Head-Fixed Mice. *Cell Rep.* **20**, 2513–2524 (2017).

31. N. A. Steinmetz, C. Aydin, A. Lebedeva, M. Okun, M. Pachitariu, M. Bauza, M. Beau, J. Bhagat, C. Böhm, M. Broux, S. Chen, J. Colonell, R. J. Gardner, B. Karsh, F. Kloosterman, D. Kostadinov, C. Mora-Lopez, J. O'Callaghan, J. Park, J. Putzeys, B. Sauerbrei, R. J. J. van Daal, A. Z. Vollan, S. Wang, M. Welkenhuysen, Z. Ye, J. T. Dudman, B. Dutta, A. W. Hantman, K. D. Harris, A. K. Lee, E. I. Moser, J. O'Keefe, A. Renart, K. Svoboda, M. Häusser, S. Haesler, M. Carandini, T. D. Harris, Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science*. **372** (2021), doi:10.1126/science.abf4588.

Supplemental figures

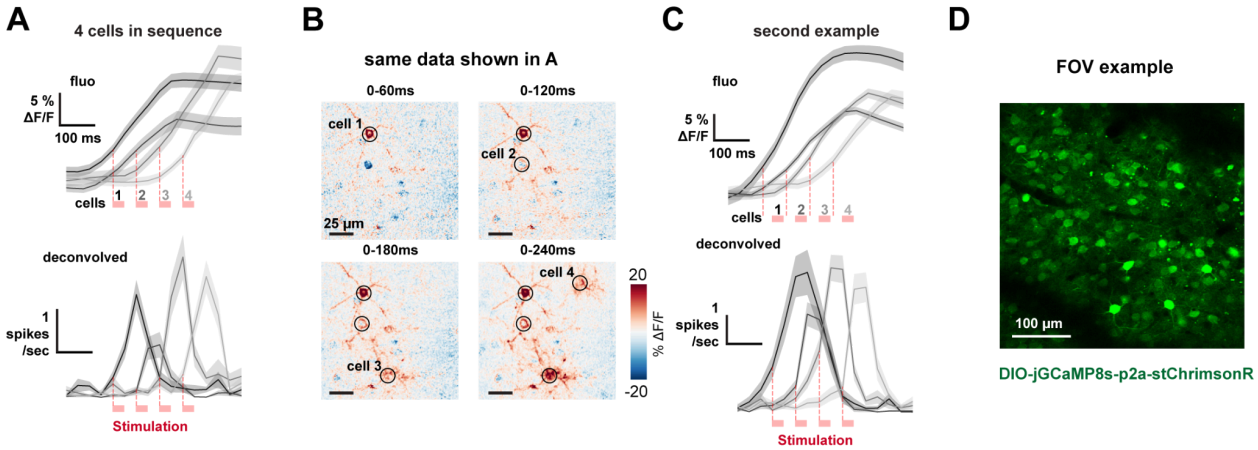

Supplementary Figure 1: Examples of two-photon stimulation timing precision. (**A**) Sequential stimulation of 4 individual cells (30 ms per stim, 30 ms between stim). Fluorescent calcium traces above; deconvolved traces below. (**B**) Response maps from data in A. Stimulation pattern lasts 30 ms each in sequence triggered at times: 0 ms, 60 ms, 120 ms, 180 ms. (**C**) Second example of 4 cell sequential stimulation (30 ms per stim, 30 ms between stim). (**A-C**) Use dual viral injection of AAV9-syn-jGCaMP7s and AAV9-syn-DIO-stChrimsonR (**D**) Example FOV image in L2/3 V1 with single virus injection, DIO-jGCaMP8s-p2a-stChrimsonR; this uses GCaMP8s which has even faster dynamics than GCaMP7s; more characterization in (*10*). Associated with Figure 1.

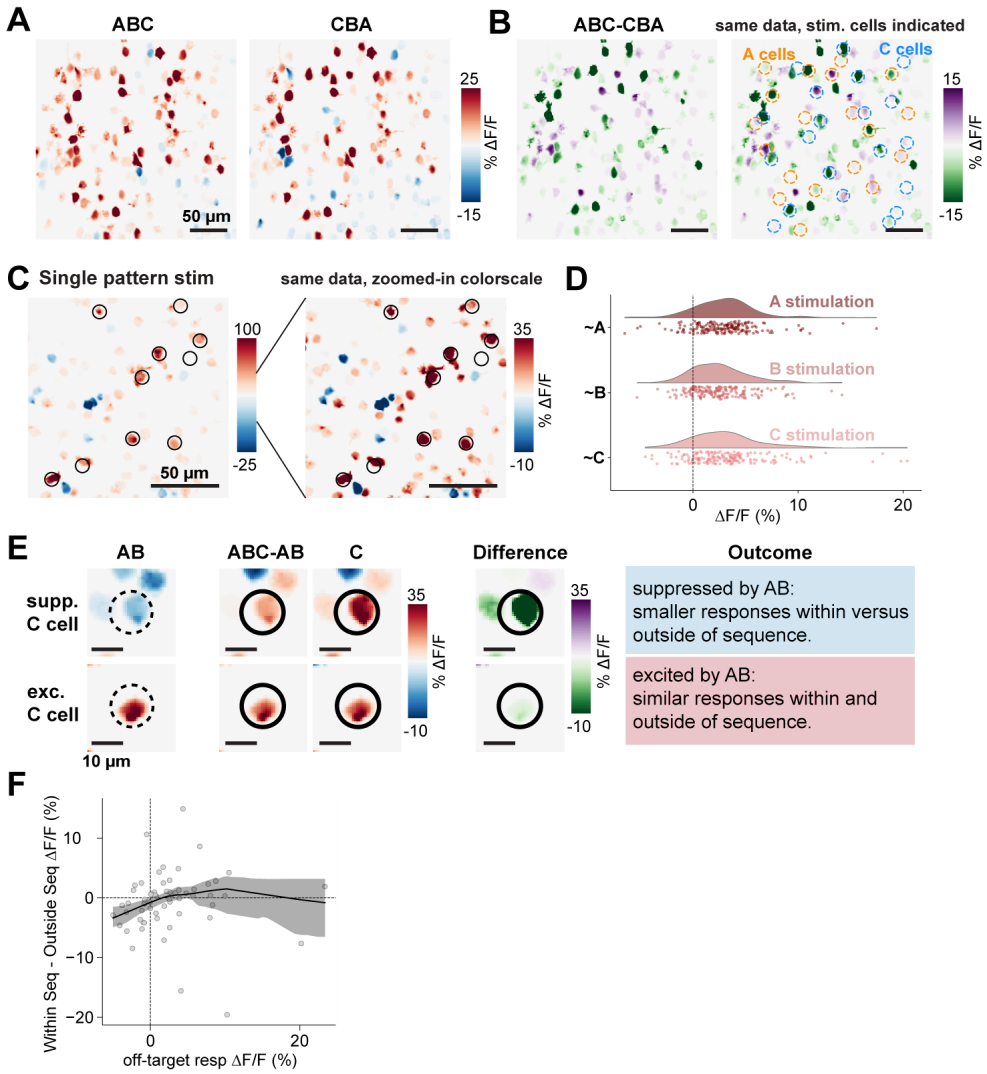

Supplementary Figure 2: Patterned stimulation leads to specific patterned responses in stimulated and off-target cells. (**A**) Responses (ΔF/F) to ABC stimulation and CBA stimulation. (**B**) Difference between ABC and CBA. A cells show a larger mean response than C cells, but effects vary from cell to cell (explored in more detail via matching visual responses in Fig. 3). (**C**) Responses (ΔF/F) to single pattern. Black circles: stimulation targets. Right: same data as left, decreased color scale highlights off-target patterns of excitation and suppression in non-stimulated cells. (**D**) Distribution of mean off-target (non-stimulated) cell responses from A, B and C pattern stimulation. Stimulation generates a range of excitation and suppression in non-stimulated cells. (**E**) Dotted black circle: nonstimulated cell; Solid black circle: stimulated cell. Top: example non-stimulated suppressed cell in response to AB stim. Prior suppression attenuates the stimulated response. Bottom: example non-stimulated excited cell in response to AB stim with same comparison within versus outside the sequence. Prior excitation leaves the incremental response largely unchanged. (**F**) Lowess fit of C cell mean responses within vs. outside of the sequence ((ABC-AB) - C) depending on the off-target modulation from prior stimulation (AB). More suppression from prior stim leads to stronger attenuation of the cell response within the sequence. Associated with Figure 2.

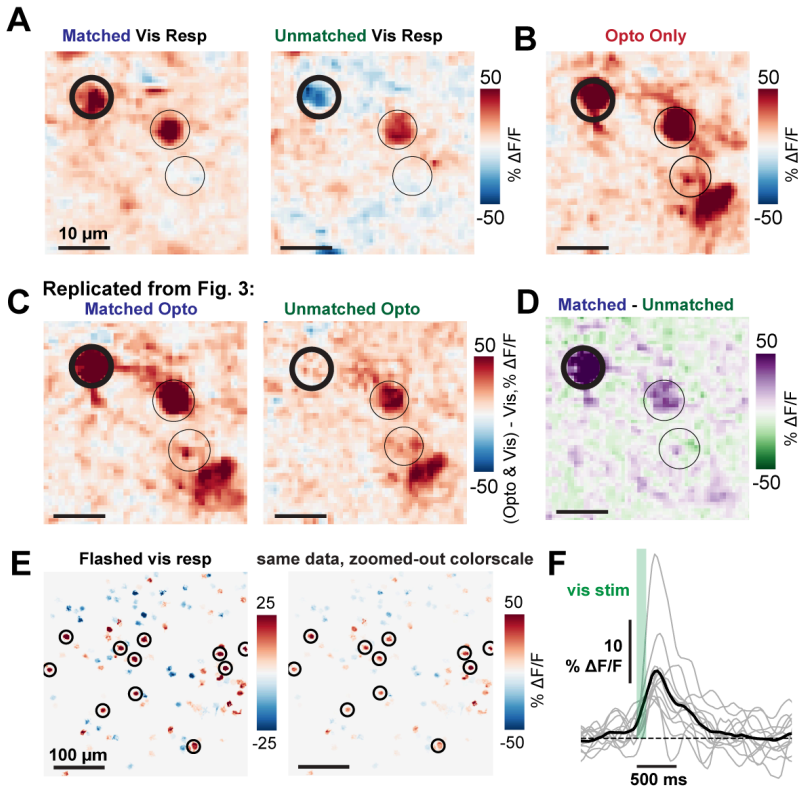

Supplementary Figure 3: Amplification is not due to spike or indicator saturation. (**A**) Visual response of the same cells as shown in Fig. 3. Thick black circle: example single cell excited in matched context while suppressed in unmatched context. (**B**) Optogenetic response of same cells without prior visual stimulus. (**C**) Optogenetic response with matched and unmatched visual context (data replicated from Fig. 3 for comparison purposes). Example cell (thick black line) shows a smaller response with the unmatched context, but starting from a suppressed baseline (panel A), ruling out attenuation being due to indicator saturation. (**D**) Difference between matched opto and unmatched opto. Matched context amplifies optogenetic pattern response relative to the unmatched context. (**E-F**) Cell selection for the natural-image-derived optogenetic pattern. (**E**) Response to flashed single frame; stimulated cells circled in black. Right: same data as left with un-zoomed color scale to show the larger cell responses that were chosen for optogenetic stimulation. (**F**) Timecourses of the circled cells in E. Only cells with a mean response above 5% are included in the pattern (Methods). Associated with Figure 3.

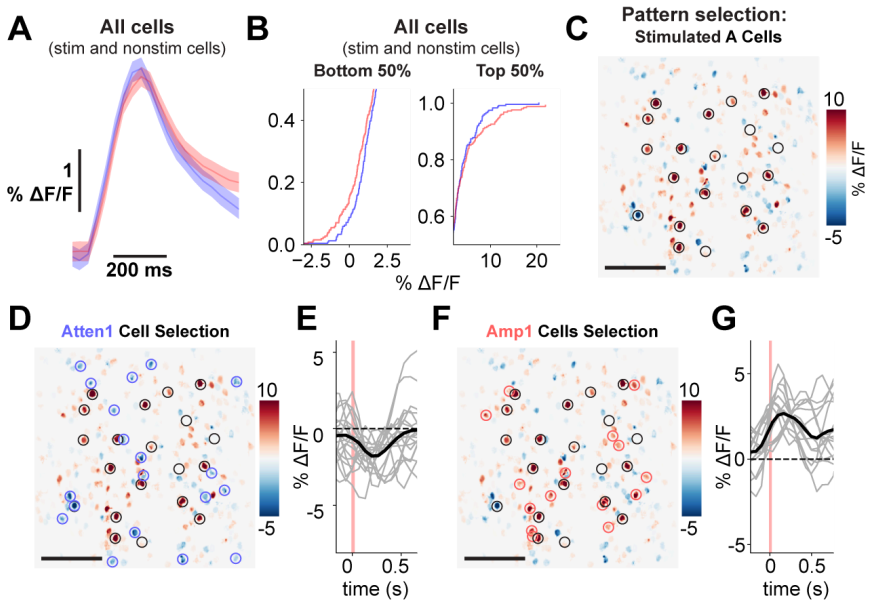

Supplementary Figure 4: Constructed amplified and attenuated sequences have similar mean population responses. In Fig. 4, we show responses to stimulated cells. Here we show both stimulated and off-target (non-stimulated) cells. While the stimulated cells show amplification/suppression (Fig. 4) as expected based on how they were selected, the overall population response mean does not appreciably change. (**A**) Average response (mean +/- SEM) of all cells in the FOV for the attenuated sequence (blue) and amplified sequence (pink). The mean population response for the amplified and attenuated sequences is similar. (**B**) Bottom 50% of cell responses for all cells in the FOV (left) for the attenuated sequence (blue) and amplified sequence (pink). Top 50% of cell responses (right). Some increased responses in the amplified sequence are balanced by an increased number of suppressed cell responses. (**C**) Example pattern; here the first arbitrary pattern (N=20 cells; pattern A) that was used to initiate sequence design (Fig. 4). (**D**) Same data as in C with attenuated pattern neurons shown. Selected suppressed (attenuated) cells circled in blue. This pattern contains cells that are suppressed by prior stimulation. (**E**) Timecourses of individual suppressed cells (from D) to A pattern stimulation. Mean trace (black) shows the mean is suppressed; suppressed cells were selected. (**F**) Same data as in C with amplified pattern cells shown. Amplified pattern: selected cells (red circles) excited by A stimulation. This pattern contains cells that are activated by prior stim. (**G**) Response timecourses of excited cells (F) in response to A pattern stimulation. Mean trace in black. Associated with Figure 4.