

Visual exposure enhances stimulus encoding and persistence in primary cortex

Andreea Lazar^{a,b,1,2} , Christopher Lewis^{a,c,1}, Pascal Fries^a, Wolf Singer^{a,b,d,2,3}, and Danko Nikolic^{a,b,d,e,3}

^aErnst Strüngmann Institute (ESI) for Neuroscience in Cooperation with Max Planck Society, 60528, Frankfurt, Germany; ^bDepartment of Neurophysiology, Max Planck Institute for Brain Research, 60528 Frankfurt, Germany; ^cLaboratory of Neural Circuit Dynamics, Brain Research Institute, University of Zürich, 8057 Zürich, Switzerland; ^dFrankfurt Institute for Advanced Studies, 60438 Frankfurt, Germany; and ^eevocenta GmbH, 45886 Gelsenkirchen, Germany

Contributed by Wolf Singer, September 3, 2021 (sent for review March 18, 2021; reviewed by Mark F. Bear and Wolfgang Maass)

The brain adapts to the sensory environment. For example, simple sensory exposure can modify the response properties of early sensory neurons. How these changes affect the overall encoding and maintenance of stimulus information across neuronal populations remains unclear. We perform parallel recordings in the primary visual cortex of anesthetized cats and find that brief, repetitive exposure to structured visual stimuli enhances stimulus encoding by decreasing the selectivity and increasing the range of the neuronal responses that persist after stimulus presentation. Low-dimensional projection methods and simple classifiers demonstrate that visual exposure increases the segregation of persistent neuronal population responses into stimulus-specific clusters. These observed refinements preserve the representational details required for stimulus reconstruction and are detectable in postexposure spontaneous activity. Assuming response facilitation and recurrent network interactions as the core mechanisms underlying stimulus persistence, we show that the exposure-driven segregation of stimulus responses can arise through strictly local plasticity mechanisms, also in the absence of firing rate changes. Our findings provide evidence for the existence of an automatic, unguided optimization process that enhances the encoding power of neuronal populations in early visual cortex, thus potentially benefiting simple readouts at higher stages of visual processing.

visual exposure | stimulus persistence | primary visual cortex | self-organization

A key property of cortical circuits is their capacity to reorganize structurally and functionally with experience (1–3). In primary visual cortex, adaptive reorganization is well documented during development (4–7) and growing evidence indicates that sensory responses continue to adapt in adulthood (8–13). The continual refinement of sensory neurons based on the statistics of the sensory environment is at odds with the traditional view of the primary visual cortex as a collection of static filters or feature detectors, passively converting sensory input into a sparse code for further feedforward processing across the visual hierarchy (14). In fact, considerable evidence suggests that primary visual cortex does not statically encode the environment but has rich spatial and temporal dynamics. For example, sensory-evoked activity propagates through the local network in wavelike patterns (15–17), displays a high degree of temporal structure (18), and can persist long after the cessation of stimulation (19–22). These rich dynamic properties exhibited by early visual neurons suggest an active involvement of primary visual cortex populations in the coordinated representation of visual stimuli. Most strikingly, repetitive visual exposure can alter the strength and selectivity of neuronal responses in the primary visual cortex, leaving a lasting mark on postexposure activity in both awake and anesthetized animals (23, 24). Yet, it remains unclear how such changes affect the joint encoding of stimuli across neuronal populations and ultimately the information transmitted to downstream areas.

Given that primary neurons adapt their responses as a function of repeated exposure, one compelling hypothesis is that exposure-driven changes are coordinated across neuronal

populations to collectively improve the representation and maintenance of recently experienced stimuli. Here, we test this hypothesis by investigating the impact of visual exposure on the persistent population response of neurons in cat area 17 to brief, structured stimulation. We employ a large set of abstract stimuli (letters of the Latin alphabet and Arabic numerals) that provide a rich variety of spatial conjunctions across low-level features and are well suited to capture aspects of distributed coding. We find five main signatures of functional reorganization. First, visual exposure optimizes stimulus maintenance in primary visual cortex by increasing the magnitude and decreasing the variability of neuronal responses that persist after stimulus offset. Second, these changes are associated with neural recruitment, a broadening of the dynamic range neurons employ to respond to stimuli, and an enhancement of stimulus-specific tiling of neuronal responses. Third, refinement of individual responses results in increased stimulus encoding at the population level; i.e., a simple hypothetical downstream decoder increases its accuracy in identifying recent stimuli from brief snippets of population activity. Fourth, the exposure-driven enhancements in stimulus persistence maintain the representational structure of stimuli, resulting in improved stimulus reconstruction. Fifth, exposure strengthens patterns in postexposure spontaneous activity. Finally, modeling demonstrates that exposure-driven enhancements in stimulus

Significance

Experience shapes sensory responses, already at the earliest stages of cortical processing. We provide evidence that, in the primary visual cortex of anesthetized cats, brief repetitive exposure to a set of simple, abstract stimuli expands the range and decreases the variability of neuronal responses that persist after stimulus offset. These refinements increase the stimulus-specific clustering of neuronal population responses and result in a more efficient encoding of both stimulus identity and stimulus structure, thus potentially benefiting simple readouts in higher cortical areas. Similar results can be achieved via local plasticity mechanisms in recurrent networks, through self-organized refinements of internal dynamics that do not require changes in firing amplitudes.

Author contributions: A.L., W.S., and D.N. designed research; A.L., C.L., P.F., and D.N. performed research; A.L. analyzed data; and A.L., C.L., W.S., and D.N. wrote the paper.

Reviewers: M.F.B., Massachusetts Institute of Technology; and W.M., Technische Universität Graz.

The authors declare no competing interest.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹A.L. and C.L. contributed equally to this work.

²To whom correspondence may be addressed. Email: andreea.lazar@esi-frankfurt.de or wolf.singer@brain.mpg.de.

³W.S. and D.N. contributed equally to this work.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2105276118/-DCSupplemental>.

Published October 18, 2021.

persistence can arise from recurrent network interactions via local, unsupervised plasticity mechanisms.

Results

We used silicon multielectrode arrays to record the simultaneous activity of neuronal populations in area 17 of five lightly anesthetized adult cats (*Felis catus*; mean age, 2.7 y; range, 1 to 5 y; two females). We applied standard spike-sorting techniques to isolate the action potentials of 112 single-unit and 331 multiunit clusters (*Materials and Methods*). The receptive fields (RFs) of the recorded units (27 to 52 units per session, 11 recording sessions with independent electrode positions, 443 units in total) were located nearby in visual space and were jointly stimulated by a single luminance stimulus, flashed for 100 ms over a black background (example trials in Fig. 1A). Short stimulus presentations at high contrast can produce strong persistent responses in the primary visual cortex (20, 21). In our data, the flashed stimuli evoked a biphasic population response, composed of a transient, low-latency component (≈ 50 ms) and a prolonged, persistent component (example trials and average firing rates across all 443 units from five cats in Fig. 1A). In total, 235 of 443 units (53%) fired above the baseline for the entire duration of the trial and the prestimulus baseline of at least 20%, paired *t*-test *P* values ranged between 0 and 0.0198).

The anesthetic protocol used here, consisting of intravenous suffentanil supplemented by minimal concentrations of isoflurane, was intended to model stable cortical dynamics, absent of strong fluctuations between “up” and “down” states. The stability and quality of the recordings were quantified using the power spectrum of the local field potential (LFP) and comparing the shapes of early and late spikes, neither of which exhibited

any systematic change over the exposure interval (*SI Appendix, Fig. S1*).

Visual Exposure Enhances Stimulus Persistence. How does brief visual exposure to structured stimuli affect the persistence of neuronal population responses in primary visual cortex? To address this question, we presented a large set of alphanumeric stimuli (34 uppercase letters and digits) in random order (1,700 trials in total, 50 trials per stimulus) and compared stimulus responses across either two or five consecutive trial blocks (stimulus order within each block was random and consecutive trials corresponded to different stimuli; schematic in Fig. 1B).

We measured stimulus discriminability, also known as Cohen’s d' (25), by calculating, for each individual unit and each 50-ms time window within the trial, the spread of the mean responses to different stimuli relative to the SDs of those responses across trials (definition of d' for 34 stimuli in *Materials and Methods*). We found that visual exposure led to a substantial increase in average d' across units (two blocks, 8.38% increase for the interval 100 to 800 ms, paired *t* test, $P = 6.4e-13$, $t = -7.43$, $df = 442$ and 10.86% increase for the interval 300 to 600 ms, paired *t* test, $P = 2.8e-12$, $t = -7.19$, $df = 442$; profile of average d' along the trial in Fig. 1C). The increase in d' with visual exposure was gradual and did not reach a saturation point (five blocks, 10.9% increase between blocks 1 and 5, paired *t* test, $P = 1.3e-07$; the black line indicates the linear fit, $y = 2.4x + 98.66$, linear trend was significant at $P = 0.009$; Fig. 1D), suggesting that further improvements may be possible with further exposure.

An improvement in stimulus discriminability is likely to be associated with an increase in neuronal response amplitude or a decrease in response variability. We found that visual exposure resulted in an increase in the amplitude of neuronal responses

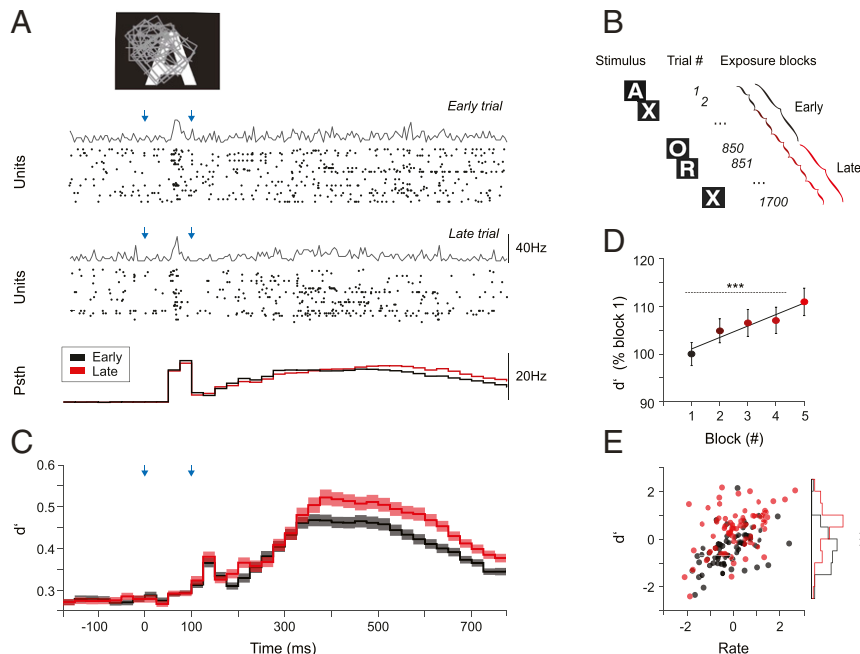


Fig. 1. Visual exposure protocol and simultaneous recordings of neuronal population activity from cat area 17. (A) Cluster of receptive fields of simultaneously recorded multiunits (rectangles) relative to the location of a visual stimulus. Shown is population activity in example trials from one session. After a brief stimulus presentation (100 ms, on and off timing marked in blue arrows), neuronal responses display a short transient followed by a persistent reverberatory component. *Bottom* peristimulus time histogram (Psth) shows mean population firing rates across 443 recorded units in early (black) and late (red) trials. (B) Exposure protocol: In each session 34 visual stimuli were presented in random order (1,700 trials in total). Sessions were split in either two or five consecutive blocks of trials and analyzed separately. (C) Average stimulus discriminability (d') over the course of the trial (d' calculated per unit; 443 units; 34 stimuli; shaded area indicates the SE from the mean). The effect of visual exposure on stimulus discriminability is significant for the persistent (300 to 600 ms), but not the transient (0 to 300 ms) part of the evoked response. (D) The increase in d' is gradual and does not saturate for the exposure interval considered here (five blocks; interval 300 to 600 ms). (E) Similar firing-rate levels result in higher discriminability in late trials (red) compared to early trials (black) in a session. $***P < 0.001$.

that persisted after stimulus presentations (6.2% increase between early and late trials for the 300- to 600-ms interval; paired t test, $P = 2e-08$, $t = -5.71$, $df = 442$; Fig. 1A and *SI Appendix*, Fig. S2). Interestingly, similar firing-rate levels resulted in higher d' values for late trials in a session, suggesting that modulations in firing rates alone cannot explain the observed improvements in stimulus discriminability (Fig. 1E). Moreover, the same conclusion was reinforced by the observation that, in two animals, the mean population firing rate was unchanged by stimulus exposure, despite substantial improvements in stimulus encoding (cats 3 and 5; see related stimulus decoding performance in *SI Appendix*, Fig. S5). It is known that firing variability is reduced by stimulus onset (26). Here we found that visual exposure further reduced variability throughout the trial (3.48% decrease between early and late trials for the interval 300 to 600 ms; paired t test, $P = 8e-05$, $t = 3.98$, $df = 442$; *SI Appendix*, Fig. S2). Laminar analysis revealed that exposure-driven changes in response amplitude, variability, and d' were significant for all compartments (*SI Appendix*, Fig. S3).

Exposure Increases the Dynamic Range and Stimulus Clustering of Neuronal Responses. Given the common assumption that higher response selectivity corresponds to more stimulus information, we considered the possibility that the observed exposure-driven enhancement in stimulus discriminability may be associated with an increase in response sparseness or selectivity. We assessed both the population sparseness for each stimulus and the stimulus selectivity of each unit, separately for early and late trials in each session (response period 300 to 600 ms; *Materials and Methods*). Sparseness was estimated as one minus the fraction of simultaneously recorded units that responded to each stimulus, and, conversely, the selectivity of each unit was estimated as one minus the proportion of stimuli it responded to. We found that both measures decreased with visual exposure (sparseness, paired t test across stimuli and sessions, $P = 6.8e-04$, $t = 3.4$, $df = 373$; values z scored per session; Fig. 2A; selectivity, paired t test, $P = 0.006$, $t = -2.7$, $df = 442$; not shown).

When units were sorted based on their change in firing-rate amplitude and grouped into quartiles, we found that the units that strongly decreased their firing rates with exposure showed increased selectivity, but decreased d' (compare Fig. 2B and C). Conversely, the units that increased their firing rates with exposure became less selective and increased their d' values. Interestingly, we found that exposure recruited more units to stimuli (reduced sparseness) and that recruited units increased the dynamic range of their firing-rate responses (the difference between the strongest and the weakest response across stimuli; *Materials and Methods* and Fig. 2D). The increased dynamic range was highly significant across units (paired t test, $P = 7.5e-17$, $t = -8.6$, $df = 442$).

Reducing the high-dimensional population response via principal component analysis (PCA), we found that exposure increased the segregation of responses, revealing stimulus-specific clusters in low-dimensional projections (examples in Fig. 3A and B; multiple sessions in *SI Appendix*, Fig. S4). Better segregation was quantified as reduced cluster radius and increased intercluster distance. The cluster radius, the mean Euclidean distance in the first two principal components of all cluster points (stimulus trials) to the cluster center (average response), decreased significantly with exposure (10.4% decrease, paired t test, $P = 1.25e-15$, $t = 8.36$, $df = 373$; Fig. 3C). In addition, the cluster distance, the mean distance between the cluster center and all other centers, increased with exposure (15.3% increase, paired t test, $P = 4.7e-25$, $t = -11.13$, $df = 373$; Fig. 3D). Overall, the clustering index, calculated per session as one minus the mean ratio between the cluster radius and the cluster distance, increased with exposure (paired t test, $P = 0.0058$, $t = -3.48$, $df = 10$; Fig. 3E).

In sum, these results suggest an exposure-driven refinement of stimulus encoding. This refinement does not occur through

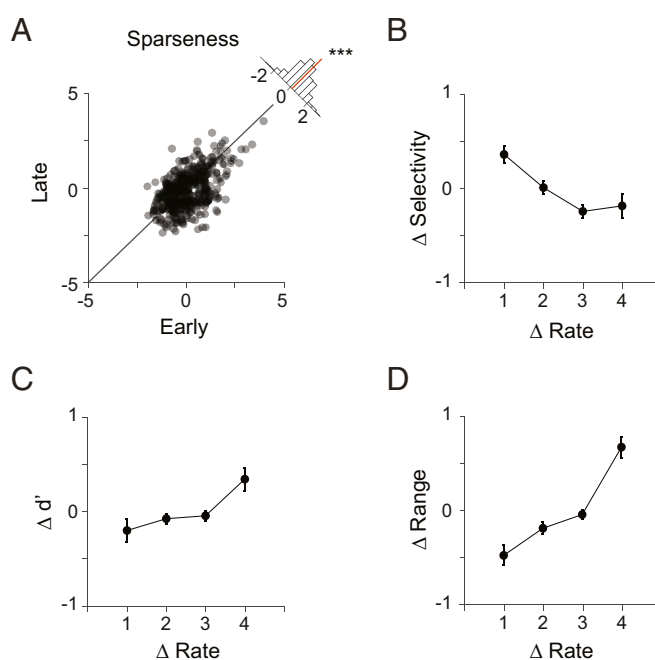


Fig. 2. Exposure-driven refinements in stimulus encoding. (A) Response sparseness decreases significantly with visual exposure, suggesting neural recruitment. Individual markers correspond to individual stimulus conditions (11 sessions, 34 stimuli). (B) Change in response selectivity as a function of change in firing rate with exposure. Units are sorted by amplitude of rate change (late-early) and grouped into quartiles (marked on x axis). Values on y axis are z scored per session. (C) Change in d' as a function of change in firing rate with exposure. Units that increase their firing rates with visual exposure also increase their d' values, but are accompanied by a loss in selectivity (compare to B). (D) Change in response range as a function of change in firing rate with exposure. Positive gains in response range are associated with an increase in firing rates. All measures are calculated over the 300- to 600-ms time interval in the trial. Error bars in B–D indicate SEs from the mean. *** $P < 0.001$.

increased selectivity of units, or population sparseness, but rather through the recruitment of more responsive units into the population response, an expansion of the dynamic range of units, and enhanced stimulus-specific clustering of population responses.

Exposure Enhances Readout Performance. We next sought to investigate the extent to which exposure-driven changes in neuronal responses affect the capacity of a hypothetical downstream decoder to identify visual stimuli based on the primary visual cortex output.

We trained independent Bayesian classifiers to perform time-resolved decoding of stimulus identity based on the population activity vector across the trial, i.e., the spike count in each time bin (instantaneous decoders; schematic in Fig. 4A). Visual exposure led to increased classification performance (time course in Fig. 4B; chance level = 2.94% for 34 stimuli, 50-ms time bins, 100-fold validation procedure; see *Materials and Methods* for details). The magnitude of the increase was substantial given the modest changes in firing rate and variability observed for individual units. Peak accuracy across sessions ranged from 8 to 49.5% correct for early trials and 16.5 to 59.6% correct for late trials and increased significantly in every animal (average increase 27.7%, range 13 to 59.6% increase, t test, all P values < 0.001). To quantify stimulus persistence, we analyzed the time course of classification performance within the trial by calculating the area under the curve (AUC) (100 to 800 ms). We found that visual exposure led to a strong and significant increase in performance AUC in every animal (average increase 33%, range 14 to 64%, t test, all P values < 0.001 ; individual performance profiles in *SI Appendix*, Fig. S5).

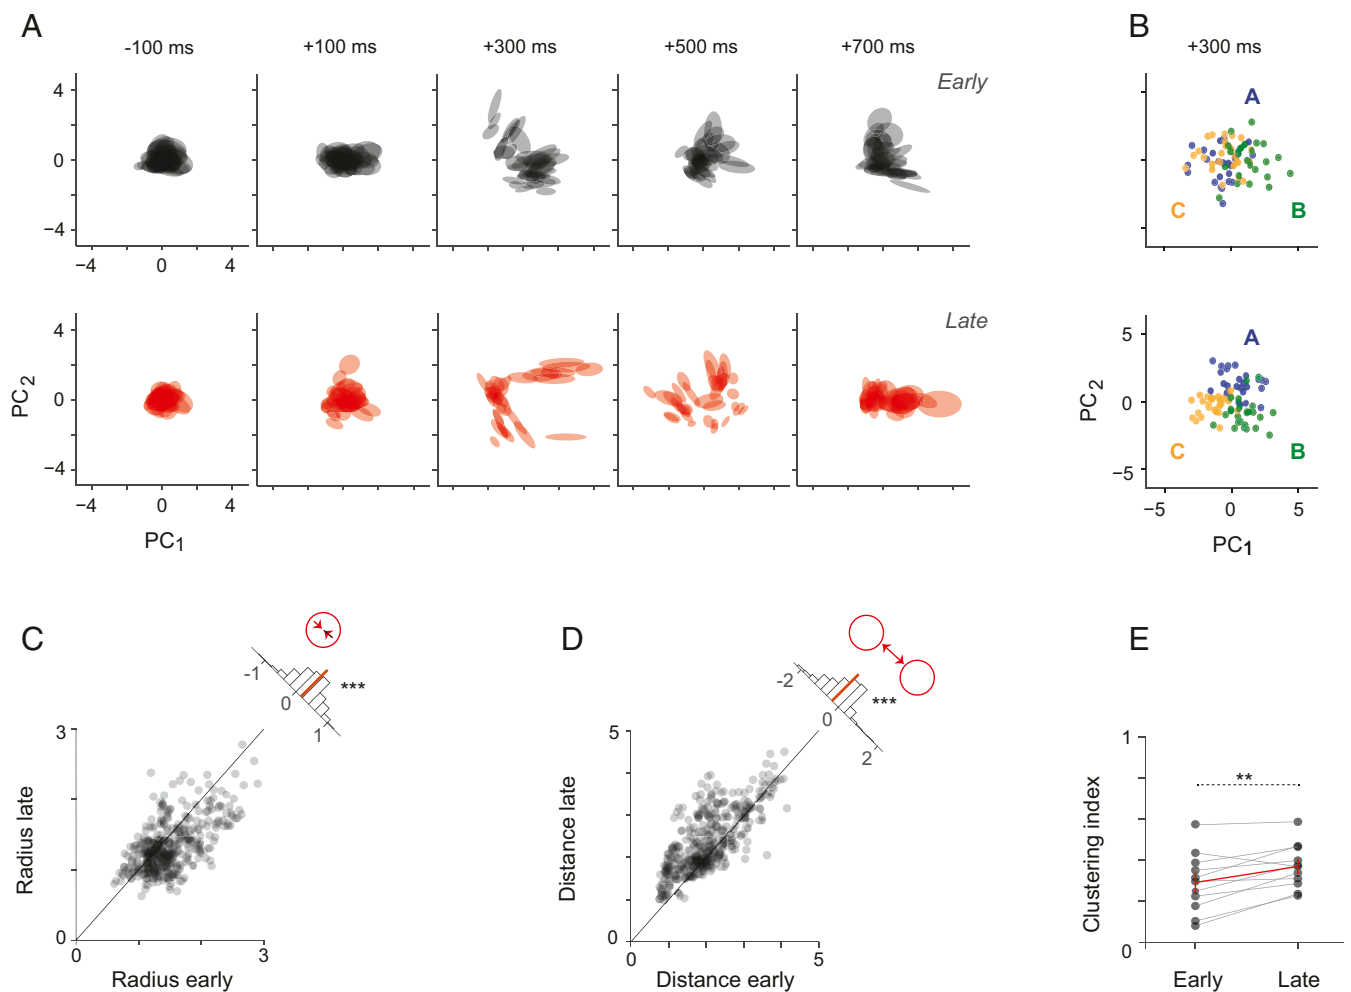


Fig. 3. Exposure increases stimulus-specific clustering and segregation of the population responses. (A) Evolution of population responses to 34 visual stimuli over the course of the trial (example session, 50-ms spike-count vectors). (B) Early and late population responses (50-ms spike-count vectors) to three stimuli (letters A, B and C) in the space defined by the first two principal components. Each marker represents a single trial. For each cluster, an ellipse circumscribes the data points within one SD from the mean. The stimulus-specific clusters segregate ≈ 300 ms after stimulus onset. The segregation is more pronounced for late trials in A (red) compared to early trials (black) in a session. (C) Scatter of cluster radius values in early and late trials, for all stimulus clusters in all sessions (50-ms spike counts, 300 ms after stimulus offset). Inset histogram shows the distribution of differences in cluster radius between early and late trials; the mean of the distribution that is significantly different from 0 is marked in red. (D) Scatter of mean distances from each cluster center to all others, in all sessions (same time window as in C). Inset histogram shows the distribution of differences in cluster distance between early and late trials. (E) The clustering index increases significantly across recording sessions with visual exposure. ** $P < 0.01$, *** $P < 0.001$.

Given the large number of stimuli in our set, individual stimuli were typically repeated only 50 times. Previous studies have indicated that remarkably little stimulus exposure is required to modify the response properties of neurons in the primary visual cortex (24). However, to test the effect of more repetitions, we acquired a session of 5,100 trials (150 trials per stimulus). In this control session, the performance AUC continued to increase past 50 repetitions per stimulus (SI Appendix, Fig. S6), suggesting that additional exposure continues to enhance stimulus encoding.

As expected, the bin size used to integrate spike counts affected the difference in decoding performance between early and late trials (Fig. 4B). Decoders that counted spikes over intermediate integration windows (50 to 200 ms) had high performance and showed significant improvements between early and late trials (performance AUC, t test, $P < 0.05$), while very short (10 ms) and very long (400 ms) integration windows resulted in lower performance and reduced improvement (AUC, t test, $P > 0.05$). Additionally, the effect of visual exposure on decoding performance varied with the task difficulty, i.e., the number of stimuli being decoded (Fig. 4C). We found that exposure improved peak

performance when classifying eight or more stimuli (t test, $P < 0.05$), but not fewer (t test, $P > 0.05$). This is likely due to ceiling effects as peak performance scores for two class problems were beyond 90% for early trials in 4 of 11 sessions.

The segregation of evoked responses into stimulus-specific clusters varied substantially over the course of the trial (Fig. 3A) and peaked at different moments in time for different animals (SI Appendix, Fig. S5). We therefore examined how stimulus-specific information varied with trial time by considering three additional decoding configurations (schematic in Fig. 4A). First, we trained “time-invariant” decoders on activity vectors pooled across five consecutive 50-ms time bins. Note that such decoders have five times more data points for the same N -dimensional space, compared to the instantaneous decoders. Second, we trained “aggregate” decoders on concatenated activity vectors corresponding to five consecutive temporal bins. Aggregate decoders map a five times larger dimensional space compared to the instantaneous decoders. Finally, we trained “scrambled” decoders on temporally scrambled data across five temporal bins. Scrambled decoders have the same number of points and

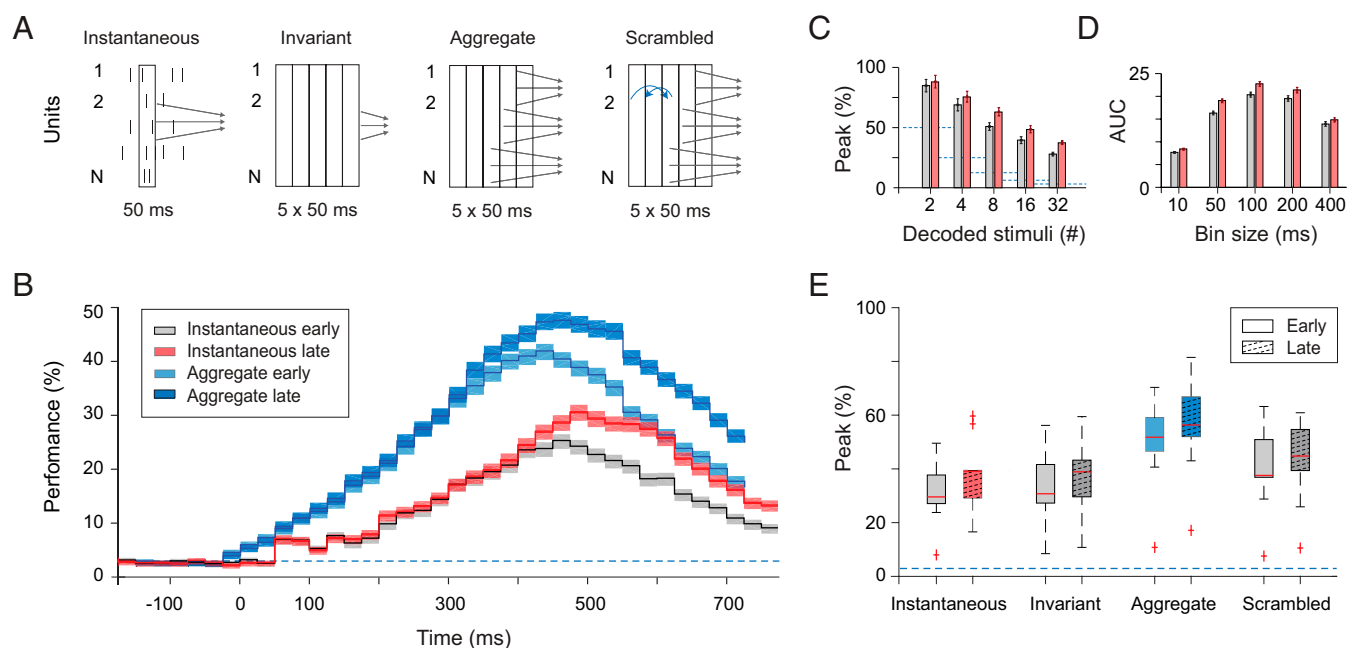


Fig. 4. Exposure enhances stimulus decoding performance. (A) Types of Bayesian decoders: instantaneous (N -dimensional readout trained and tested on brief 50-ms spike counts, independently for each time bin), invariant (N -dimensional readout trained on population spike counts from five consecutive 50-ms time-bins; quantifies information that is consistent across trial time), aggregate ($5 \times N$ -dimensional readout integrates information from five consecutive time bins), and scrambled ($5 \times N$ -dimensional readout integrates across five time bins after temporal scrambling). (B) Stimulus classification performance over the course of the trial for instantaneous and aggregate decoders (chance level indicated by dashed line; individual sessions in *SI Appendix, Fig. S5*). (C) Peak classification performance of instantaneous decoder as a function of task difficulty (number of stimuli being classified), relative to chance level (dashed lines). Differences between early and late trials are more pronounced for harder classification problems. (D) Dependence of performance of instantaneous decoder, quantified by the AUC, on the integration bin size denoted on the x axis, for early and late trials. Performance is highest for intermediate integration intervals. (E) Peak classification performance for the Bayesian decoders described schematically in A. All four types are significantly improved, suggesting that both variant and invariant aspects of stimulus encoding are enhanced by visual exposure.

space dimensions as the aggregate decoders but map an altered space where within-trial correlations between neurons have been disrupted through scrambling. Interestingly, all three decoders showed significant changes with visual exposure. The aggregate decoder performed significantly better than both the invariant and temporally scrambled decoders, suggesting that information was contained not only in the instantaneous structure of the spike-count vector but also in its trajectory (in the sequence of state vectors during trial).

Finally, we considered the impact of visual exposure on the portion of trial-to-trial variability shared between units. Consistent with previous studies (27, 28), spike-count correlations (SCCs) were highest for pairs of units with similar stimulus preferences (positive signal correlations) and lowest for pairs of units with opposing stimulus preferences (negative signal correlations). Visual exposure reduced the strength of SCCs (21% decrease, paired t test, $P = 1e-17$), and the reduction was strongest for units with opposing preferences (78% decrease, two-tailed t test, $P = 1e-09$, signal correlations < -0.1 ; 9% decrease, two-tailed t test, $P = 1e-06$, signal correlations > 0.1 ; *SI Appendix, Fig. S7A*). Ignoring SCCs can decrease decoding performance (29, 30). We found that a support vector machine with quadratic features trained on trial-shuffled data and tested on original data performed worse than a decoder trained on the original data with intact correlation structure (two-way ANOVA; shuffling led to a 10.42% decrease, $P = 0.006$, early trials and 13.02% decrease, $P = 3.7e-09$, late trials; exposure led to a 17.96% increase, $P = 1.2e-07$ for original data and 14.54% increase, $P = 0.0003$ for shuffled data; *SI Appendix, Fig. S7B*). Shuffling reduced performance for both the early and late trials, suggesting that while repeated exposure decreased the overall level of SCCs in the data, a portion of SCCs present in both

early and late trials contributed positively to the population code. Indeed, the fact that SCCs decreased most for units with opposing stimulus preferences might reflect competition between stimulus-specific ensembles, such that correlations are stable between units of similar preference and reduced between units of opposing ensembles.

Exposure Enhances Stimulus Reconstruction. The alphanumeric stimuli are structurally more complex than oriented gratings but less complex than natural scenes. Such a large stimulus set of intermediate complexity is highly suitable for reconstruction, i.e., recreating the luminance pattern of stimuli from their evoked neuronal responses. While stimulus decoding techniques have been applied to many visual cortical areas, stimulus reconstruction has been attempted rarely and not, to our knowledge, in the context of visual exposure.

We performed stimulus reconstruction separately on early and late trials to quantify the impact that visual exposure had on encoding. To reconstruct each stimulus, we trained Bayesian decoders to predict the luminance of individual stimulus patches (576 decoders corresponding to 24×24 image patches) based on the population activity recorded after stimulus offset (schematic in Fig. 5A; luminance values were between 0 and 1; independent train and test trials, 20-fold validation scheme; 11 sessions). The reconstructed stimuli were noisy on single trials (examples from one session, 50-ms spike counts, 400 ms after stimulus offset; Fig. 5B, *Left*) and became considerably more accurate when averaged across 10 test trials (Fig. 5B, *Right*). Exposure increased stimulus reconstruction accuracy, calculated as one minus the difference in pixel luminance between the reconstructed image and the original image (scatter of 34 stimuli from 11 sessions; paired t test, $P = 1.13e-09$, $t = -6.24$, $df = 373$; Fig. 5C, *Left*). This

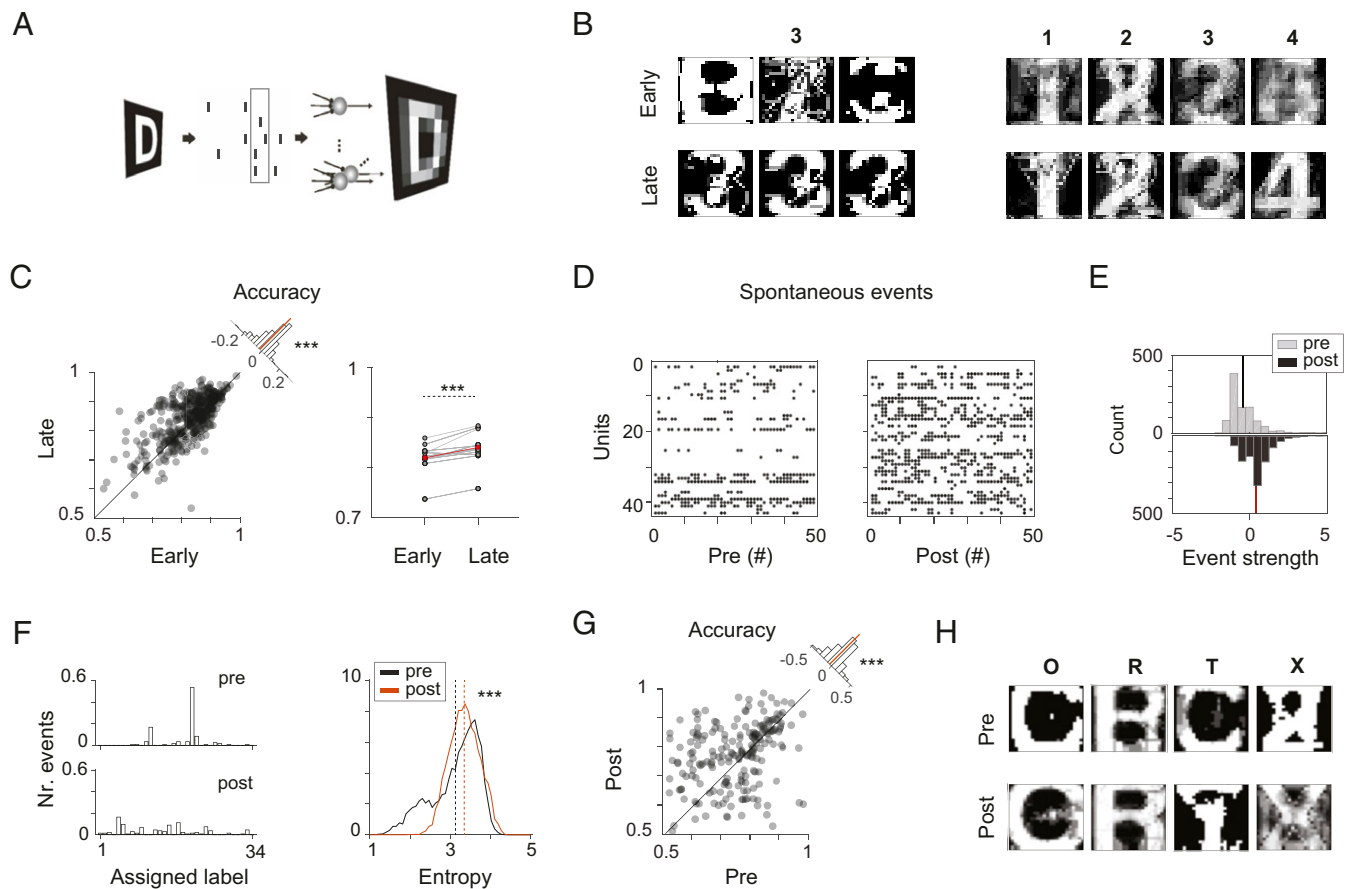


Fig. 5. Improved stimulus reconstruction with visual exposure. (A) Schematic representation of stimulus reconstruction technique: A 24×24 array of Bayesian classifiers is trained to predict pixel luminance from population spike-count vectors (50 ms) recorded 400 ms after stimulus offset. Test trials are omitted from the training set. (B) Single-trial examples for reconstruction of a stimulus (number “3”) and examples of average reconstructions (across 10 test trials, numbers “1” to “4”) for early and late trials in an example session. (C) Scatter depicts the reconstruction accuracy of all stimuli from 11 sessions, early vs. late trials (Left). Changes in reconstruction accuracy are significant (Inset histogram, paired *t* test; mean accuracy per session in Right plot, paired *t* test). (D) Example of 50 spontaneous events detected pre- and postvisual exposure in a session. (E) Spontaneous events become stronger postexposure (events detected pre- and postexposure in 7 sessions). (F) A Bayes classifier trained on evoked activity assigns stimulus labels to spontaneous events. The assigned labels are more uniformly distributed postexposure; entropy (measured in bits) is significantly higher postexposure (Right plot; permutation test). (G) Stimulus reconstruction accuracy based on spontaneous events improves postexposure. (H) Examples of single-letter reconstructions based on spontaneous events occurring pre- and poststimulus exposure (24×24 luminance patches; training on evoked data). *** $P < 0.001$.

improvement was significant also when calculated across sessions (paired *t* test, $P = 8.2e-04$, $t = -4.71$, $df = 10$; Fig. 5C, Right). These results suggest that the exposure-driven enhancements in classification performance are representational in nature and reflect improved encoding of stimulus content. The fact that the structure of multiple stimulus shapes (34 letters and digits) can be reconstructed from relatively small populations of recorded units (27 to 52 units per session) and improves with exposure speaks to the impressive encoding capacity and flexibility of the primary visual cortex.

Structured Postexposure Spontaneous Activity. Exposure-driven changes in evoked activity are often associated with accompanying changes in spontaneous activity. The stimulus reconstruction technique described above allowed us to probe for lasting representational changes in the structure of spontaneous neuronal activity. To this end, in seven exposure sessions from three cats, we recorded and analyzed spontaneous activity (20 blank trials) before and after visual stimulation.

We isolated spontaneously occurring strong-activation events, defined as periods when population activity exceeded mean activity by more than one SD (50-ms spike counts). In total, 961 events were detected during preexposure spontaneous activity and 980

during postexposure (example events corresponding to pre- and postexposure data from one session; Fig. 5D). We found that the strength of spontaneous events increased significantly after visual exposure (event strength was pooled across sessions after being standardized for mean and variance; *t* test, $P = 3.9e-89$, $t = -21$, $df = 1,939$; Fig. 5E).

To quantify structural differences in the pre- and postexposure events, we assigned a stimulus “label” to each event using decoders trained on evoked activity (example label assignments for pre- and postexposure events from one session in Fig. 5F). The assigned labels were more uniformly distributed across stimulus conditions postexposure, as indicated by an increase in entropy (estimated based on 1,000 bootstraps of 50 events, $P = 1e-172$; Fig. 5F, Right). We next generated an image, using the same reconstruction technique as above, for each spontaneous event. Both the reconstruction and the identity decoders were trained on evoked activity from the entire session (no separation of early and late trials), as we wanted to assess exposure-driven changes in the structure of spontaneous activity, not changes in the reconstruction or decoding techniques. We pooled images across events based on the assigned class label to obtain mean stimulus reconstructions (reconstruction examples in Fig. 5H). The accuracy of stimulus reconstruction

improved following exposure (paired t test, $P = 5.3e-04$, $t = -3.52$, $df = 204$; Fig. 5G). Since the assignment of a class label and the reconstruction of an image for the same spontaneous event reflect related content, it is not surprising that some degree of stimulus reconstruction is possible based on spontaneous activity. However, the improvement in postexposure reconstruction accuracy suggests that the changes in spontaneous events are structured and correspond to the experienced visual content.

Self-Organized Recurrent Networks for Stimulus Persistence.

Finally, we sought to demonstrate that a simple self-organized recurrent network, endowed with local plasticity mechanisms for learning and homeostasis, can qualitatively reproduce the exposure-driven enhancements in stimulus encoding, while maintaining stable activity levels.

We considered a simple network composed of 250 deterministic McCulloch and Pitts threshold neurons, 80% excitatory and 20% inhibitory. For simplicity, the recurrent interactions were assumed to arise from a single pool of randomly connected neurons, not a multilayered recurrent network. The connections between the excitatory units were sparse and followed simple topological constraints; i.e., nearby neurons had increased probability for a connection (details in *Materials and Methods*). A subset of neurons received luminance input from a 6×6 array of stimulus image patches (schematic in Fig. 6A).

Recurrent neural networks naturally exhibit a memory of recent inputs, so information about brief stimuli can be retrieved with some delay from stimulus offset (31, 32). To match the empirical data, we strengthened the persistent recurrent responses after stimulus offset through response facilitation, which has been previously implicated in both stimulus persistence

and learning (33, 34) (fixed interval for facilitation marked by shaded area in Fig. 6C).

Visual exposure consisted of 50 brief presentations of 10 stimuli (alphabet letters A to J) in random order. The network self-organized through synaptic plasticity, while homeostatic plasticity maintained the network output at a stable level (example of connectivity changes through plasticity in Fig. 6B, details in *Materials and Methods*).

We considered how exposure-driven learning may interact with several different implementations of response facilitation. Facilitation was implemented by changing the excitatory–inhibitory balance of incoming synaptic gains per unit by either 1) lowering inhibition or 2) increasing excitation. Alternatively, the activity of excitatory units was increased by 3) adding a random input drive and 4) changing neuronal excitability via intrinsic plasticity. The first two methods resulted in persistent unit responses after stimulus offset. Unsupervised learning during visual exposure led to improved stimulus encoding (data for an example run with lowered inhibition in Fig. 6C–E). Similar to the empirical results, the enhancement in stimulus encoding could be captured in low-dimensional projections of the data (Fig. 6E). The last two methods, variable input and changes in neuronal excitability via intrinsic plasticity, resulted in similar, persistent responses after stimulus offset. However, they did not lead to improved stimulus encoding through learning, suggesting that the intrinsic network interactions, which were more severely disrupted by these two methods, played a critical role in the optimization process (Fig. 6F).

Discussion

We found that repeated exposure to briefly flashed visual shapes improves stimulus encoding in primary visual cortex. Visual

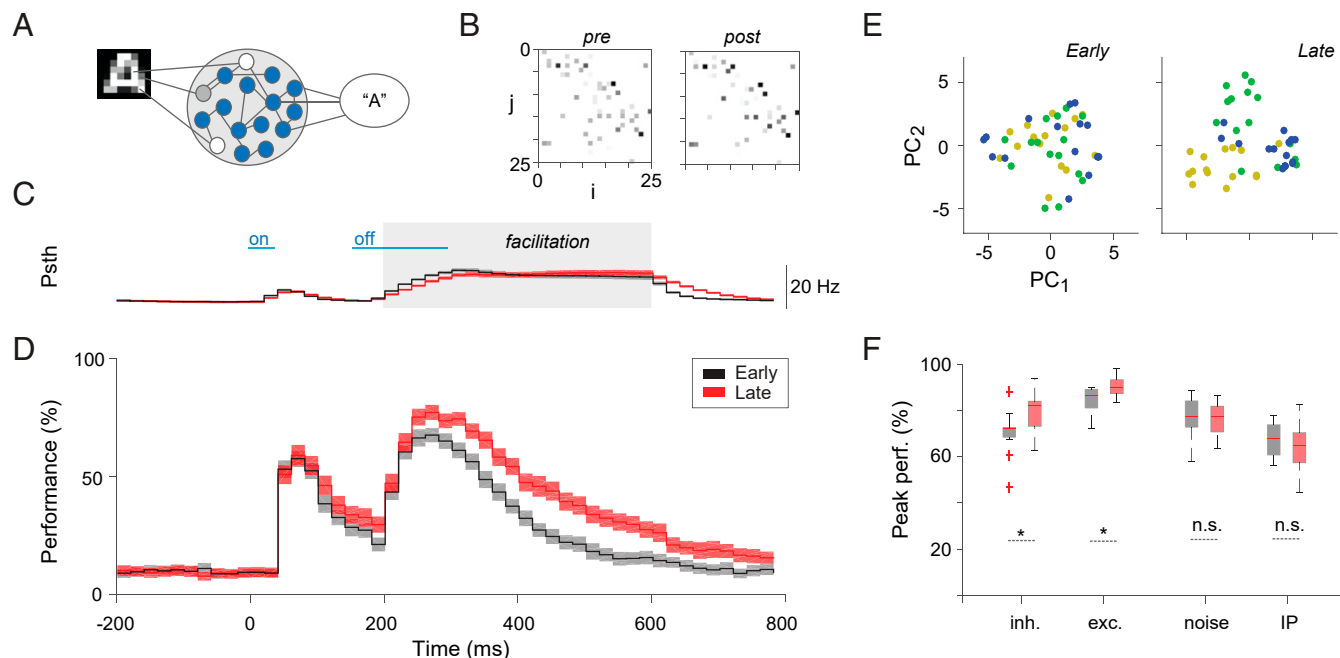


Fig. 6. SORN. (A) Recurrent neural network with subset of neurons receiving pixel luminance inputs. (B) Network topology: Initial connections are sparse and random with nearby connections more likely to occur. Shown is an example of connections between 25 excitatory units out of the 200 total, before and after learning in one simulation. Visual exposure resulted in self-organized changes in synaptic gains. (C) Example of network excitatory response to a brief stimulus. Response facilitation, implemented here through a decrease in synaptic inhibition, is marked by a gray area. (D) Example of classification performance improvement with stimulus exposure via unsupervised local plasticity rules (shaded areas indicate SEs from the mean across 10 simulations). Note that here the firing rates are similar in early and late trials, due to the homeostatic nature of intrinsic plasticity, and thus the enhancement in encoding comes from a reorganization of internal network dynamics. (E) Exposure results in an enhanced clustering of trials belonging to different stimulus conditions in a low-dimensional projection space. (F) Facilitation can boost or impede unsupervised learning, depending on implementation. Facilitation via a decrease in synaptic inhibition or an increase in synaptic excitation leads to a boost in classification performance. Facilitation via nonstructured excitatory input (noise) or threshold modulation via intrinsic plasticity interferes with learning. exc., excitatory; inh., inhibitory. * $P < 0.05$; n.s., not significant.

exposure altered poststimulus population activity in a manner that enhanced both the decoding of stimulus identity and the reconstruction of visual stimuli.

These improvements were associated with neuronal recruitment, an increase in the dynamic range of neuronal responses, and stimulus-specific clustering of population responses. The manner in which exposure enhanced the segregation of population responses into low-dimensional, stimulus-specific clusters suggested two main effects. First, stimulus responses became less variable across trials and more stereotyped, shrinking the radius of individual clusters. Second, responses to different stimuli became more distinct, increasing the distance between clusters. Using decoders we found that the information about stimulus identity was temporally specific; i.e., different time bins in the trial differed in their mapping, and visual exposure improved both variant and invariant aspects of stimulus encoding. Interestingly, we also observed a reduction in covariability, similar to what has been previously documented in the context of attention and perceptual learning (27, 35, 36). In our data, the effect of exposure on spike-count correlations was complex: The strength of correlations was reduced with exposure, but knowledge about their structure remained beneficial for stimulus discrimination.

The brief stimulus presentations employed here resulted in a stereotyped, biphasic neuronal response in primary visual cortex, consisting of a high-amplitude transient followed by a delayed persistent response. Stimulus decoding performance diverged from the expected dependence on firing rate, with accuracy peaking not on the response transient, but 200 to 400 ms after stimulus offset. Sustained and information-rich sensory responses persisting beyond the period of sensory stimulation have been reported previously, not only under anesthesia, but also in various sensory modalities and species in awake behaving animals. In the primary auditory cortex of awake marmosets, preceding stimuli suppressed or facilitated responses to succeeding stimuli for more than 1 s (37). In awake mice, the early sensory responses to a single brief whisker deflection encoded stimulus information, while the later activity appeared to drive the subjective detection (38). In the primary visual cortex of awake mice, an oriented flashing light induced a biphasic membrane voltage response that consisted of an early, transient depolarization and a delayed, slow depolarization (20). The delayed activity exhibited high orientation selectivity and influenced the evoked response to subsequent inputs in an orientation-selective manner. In awake macaques, a simultaneous change in both stimulus and background gave rise to a delayed V1 response that varied with the size of the background and correlated with the perception of a visual aftereffect (21). In human electroencephalography, information about a previously presented visual stimulus persisted even in the absence of delayed activity (activity-silent states) and could be decoded from an impulse response, long after stimulus presentation (39).

The persistence of stimulus information observed in our data and supported by the studies mentioned above highlights a propensity for the primary visual cortex to maintain sensory information, far beyond the temporal intervals required by the traditional feedforward model of the ventral stream. Instead, these findings are compatible with a dynamic coding framework for recurrent computation (31, 32, 40). In this framework, the cortical response to a stimulus emerges from an interaction between the input signals and the internal dynamical state of the network, including the ongoing activity (active states), but also the time-dependent properties of neurons and synapses (hidden states). Efficient recurrent processing relies on two simple requirements: 1) Stimulus responses must persist beyond the duration of the stimulus, establishing a brief memory of recent events (fading-memory property), and 2) the temporal evolution of network states in response to different stimuli must result in reproducible stimulus-specific trajectories (separability

property). Both the memory and separability properties exhibited by a recurrent circuit can be optimized through plasticity by altering the network's stimulus-response mapping (40).

The self-organized recurrent model presented here builds on previous computational work showing that experience-dependent plasticity increases a recurrent network's performance on memory and prediction tasks (41), while explaining numerous experimental findings on cortical variability (42). We introduced response facilitation to boost the network's response to brief stimulation and considered its effect on learning in several implementations. We found that shifts in excitatory-inhibitory synaptic gains led to strong persistent responses after stimulus offset and increased stimulus decoding performance through exposure. In contrast, an increase in excitatory drive or neuronal excitability led to strong persistent responses after stimulus offset, but did not result in improved stimulus decoding performance through exposure, suggesting that the intrinsic network interactions play an essential role during learning.

Notably, various other computational models have shown that the dynamics and performance of recurrent neural networks can be optimized via brain-inspired plasticity mechanisms. For example, spike-frequency adaptation was shown to expand the memory exhibited by recurrent circuits (43) and different forms of biologically plausible synaptic learning rules have been employed to enhance computational performance of recurrent networks in an unsupervised fashion (44–47). Furthermore, several studies made direct attempts to link learning in recurrent networks to optimization of state space dynamics (48, 49) or metalearning (50). While neither of these recurrent models tried to explain how persistent responses after brief stimulation can interact with learning, they provide valuable insights into the various means by which refinements in internal network dynamics result in improved output performance.

The precise anatomical connectivity responsible for the observed reverberation of visual responses and the functional changes underlying exposure-driven improvements in stimulus discrimination are still unknown. Given the presence of both strong feedforward and extensive feedback thalamocortical interactions (51, 52), we cannot exclude the possibility that exposure-driven changes in primary cortex responses originate from interactions with subcortical structures. In fact, studies have shown that slowly decaying inhibitory postsynaptic potentials in the lateral geniculate nucleus can maintain stimulus-specific information for up to 300 ms and can modulate subsequent responses to reoccurring contours (53, 54). However, these effects were short lasting, while in our data repeated exposure to stimuli resulted in an increase in stimulus encoding across numerous trials and was associated with a strengthening of activation patterns in postexposure spontaneous activity. Interestingly, the activation of N-methyl D-aspartate (NMDA) receptors in cortical layer 4, which receives the densest thalamocortical input, does not appear to be necessary for stimulus-selective response potentiation in V1 (55). An alternative is that exposure-dependent changes primarily affect local recurrent interactions within primary cortex and/or the long-distance recurrent interactions with higher cortical areas.

Vision depends on integrating the current sensory input in light of previous experience (56). This integration is achieved through the rich recurrent dynamics of the early visual system, which arise on the backbone of structural connectivity (57, 58). The connectivity of sensory areas is believed to capture the statistics of the environment (4, 7, 56, 59), a process that would improve processing for expected stimuli (60–62). However, exposure-driven changes in the primary sensory cortex of adult subjects suggest a complex, occasionally divergent pattern of results. For example, adaptation classically leads to a reduction in response amplitude (10), and familiar stimuli can evoke reduced responses compared to novel stimuli (63). However, repetitive exposure can

also alter the receptive fields and tuning of neurons in V1 (13), imprint responses to recent stimulus trajectories (24), increase the magnitude of responses to familiar sequences, and signal predictions to missing elements (23). These intricate changes in response amplitudes with visual exposure appear to depend on many factors, among which are the frequency and duration of stimulation, the structure and complexity of the image set, the state of the animal, and the precise signal measured. Regardless, the core findings outlined here do not rely heavily on changes in firing rates to familiar stimuli, but rather on the prolonged maintenance of stimulus information and the increase in discriminability with visual experience. We interpret these results as an accumulation of evidence that optimizes the encoding of a stimulus set akin to learning new stimulus statistics.

State changes are known to vary under anesthesia and can change the cortical response to stimulation. For example, deep anesthesia gives rise to alternating up and down states with distinct dynamic profiles that can have an impact on sensory coding. We did not observe such strong variations in our recordings, which were performed in a modified anesthetic protocol to mimic awake-like brain dynamics. Although a milder form of such variation occurs also during wakefulness, we found no systematic change in brain state that could trivially explain our results. Likewise, brief stimuli are known to generate robust responses irrespective of cortical state (64), which may explain the stability of sensory responses and brain states across recording sessions. Given that the experiments were performed under anesthesia, the reported exposure-driven changes in activity must involve “automatic” mechanisms, independent of attention and conscious control. Further work is necessary to determine to which extent these effects generalize to the waking state, where higher cortical areas with reciprocal connections to V1 as well as subcortical regions, such as the superior colliculus, thalamus, and cerebellum, are likely to play an important role in shaping V1 plasticity. In particular, top-down enhancement of task-relevant stimulus features and suppression of irrelevant ones, the level of attention, motivation, and reward expectation are all likely to guide learning-induced changes in V1.

Our study provides compelling evidence that repetitive visual exposure optimizes sensory processing in primary visual cortex, resulting in a better readout of stimulus-specific information. These findings suggest that the reliable visual discrimination of familiar stimuli can be partially achieved through separation of neuronal representations at the earliest cortical stage in the visual hierarchy. Future work should establish how these changes impact the transformation of sensory signals in the visual hierarchy, manifest at higher visual areas, and interact with behavioral states, such as attention or perception.

Materials and Methods

Electrophysiological Recordings and Data Processing. Data were recorded from five adult cats (*F. catus*; mean age, 2.7 y; range, 1 to 5 y; two females) under general anesthesia during terminal experiments in two separate laboratories. The cats were bred internally, were housed together with other cats in small groups, and experienced normal vision during development. All procedures complied with the German law for the protection of animals and were approved by the regional authority (Regierungspräsidium Darmstadt).

For one of the cats, anesthesia was induced by intramuscular injection of ketamine (10 mg/kg) and xylazine (2 mg/kg) followed by ventilation with $N_2O:O_2$ (70/30%) and halothane (0.5 to 1.0%). After verifying the depth of narcosis, pancuronium bromide (0.15 mg/kg) was added for paralysis. Stimuli were presented binocularly on a 21-inch computer screen (HITACHI CM813ET) with 100-Hz refresh rate. To obtain binocular fusion, the optical axes of the two eyes were first determined by mapping the borders of the respective receptive fields and then aligned on the computer screen with adjustable prisms placed in front of one eye. Data were recorded with multiple 16-channel silicon probes from the Center for Neural Communication Technology at the University of Michigan (each probe consisted of four shanks, 3 mm long, 200 μ m distance, four contact points each, 1,250 μ m² area, 0.3 to 0.5 M Ω impedance at 1 kHz). To extract multiunit

activity, signals were amplified 1,000 times and filtered between 500 and 3,500 Hz.

For four of the cats, anesthesia was induced by intramuscular injection of ketamine (10 mg/kg) and medetomidine (0.02 mg/kg) followed by ventilation with $N_2O:O_2$ (60/40%) and isoflurane (0.6 to 1.0%). After verifying narcosis, vecuronium (0.25 mg \times kg⁻¹ \times h⁻¹ intravenous) was added for paralysis. Data were collected via multiple 32-contact probes (100 μ m intersite spacing, \approx 1 M Ω at 1 kHz; NeuroNexus or ATLAS Neuroengineering) and amplified (Tucker Davis Technologies). Signals were filtered with a passband of 700 to 7,000 Hz and a threshold was set to retain multiunit activity. Thresholds remained fixed during data collection.

Visual Stimuli. Stimuli consisted of 34 shapes: 26 letters (A to Z) and 8 digits (0 to 7). They were white on black background and spanned 5 to 7° of visual angle.

A total of 1,700 trials (50 trials per stimulus) were recorded in every session. More than 6,800 trials (200 trials per stimulus) were recorded in one of the sessions to test whether longer exposure leads to further improvements in stimulus encoding (SI Appendix, Fig. S6).

Stimuli were presented in random order; i.e., consecutive trials corresponded to different stimulus conditions.

Data Analysis. Data were processed and analyzed using custom code written in MATLAB (MathWorks). Statistical significance was calculated via two-tailed paired *t* tests that compared early and late trials in each session, and *P*, *t*, and *df* values were reported in each case. Results across all 443 units (5 cats), pulled together, are reported in Figs. 1 and 2. Figs. 3–5 focused on population activity and combined data across all 11 sessions (5 cats). The exposure improvements in stimulus encoding, as quantified by the Bayesian classifier, were strong and statistically significant in each recorded animal (individual subplots for 5 cats in SI Appendix, Fig. S5).

Spontaneous activity was collected before and after visual exposure in 7 sessions from 3 cats (Fig. 5 D–H). The Fieldtrip toolbox (65) was used for laminar analysis (SI Appendix, Fig. S3). These results were based on 7 recording sessions from 3 cats, for which current source density maps could be calculated.

One-way ANOVA was used to test the significance of longer exposure on stimulus decodability (SI Appendix, Fig. S6) and a two-way ANOVA was used to compare both the impact of exposure and that of trial shuffling on stimulus decoding performance (SI Appendix, Fig. S7); the corresponding *P*, *F*, and *df* values were reported in each case.

Spike Sorting. Spike sorting of the recorded multiunits was performed offline via custom software that computed principal components of spike waveforms to reduce dimensionality and grouped the resulting data using a density-based clustering algorithm (DBSCAN). Only the well-isolated clusters were considered single units and labeled separately. Spike sorting resulted in 112 single units and 221 multiunits, in total 443 units across all datasets.

Current Source Density Analysis. In three cats (seven sessions, 289 units, 32-channel linear arrays 100 μ m spacing), LFPs to moving grating stimuli presented at maximal contrast were recorded either immediately before or immediately after the sessions with letters and digits. These LFPs were subject to current source density (CSD) analysis using a standard algorithm (66) based on the second spatial derivative estimate of the laminar local field potential time series. This analysis revealed successfully the short-latency current sink in the middle layers for each session, which has been shown to correspond most closely to layer 4 (67).

Neuronal Response Properties. For every unit, the discriminability index *d'*, also known as Cohen's effect size (25), for *n* stimuli, was calculated as

$$d' = \sqrt{\frac{\sum_{i=1}^n (r_i - r)^2}{n}} / \sigma, \quad [1]$$

where r_i is the mean response across trials to stimulus *i*, r is the mean response across all trials, and σ is the common within-population SD, here $\sigma = \sum_{i=1}^n \sigma_i / n$, where σ_i is the SD of responses across trials to stimulus *i*. Note that low single-unit *d'* values across many stimulus conditions ($n = 34$) are expected; see for example ref. 68 for reference.

The Fano factor was computed per unit according to Fano (69),

$$FF = \frac{\sum_{i=1}^n (\sigma_i^2 / r_i)}{n}, \quad [2]$$

where σ , r_i , and n are defined as above.

Response sparseness R describes the response distribution of a population of neurons to a single stimulus. Within each session, the response sparseness of the recorded population of units to each stimulus was calculated as

$$R_i = (1 - f_i/m), \quad [3]$$

where f_i/m is the fraction of units out of the total m that fired above the baseline in response to stimulus i .

Stimulus selectivity quantifies the responsiveness of a unit across a set of stimuli and was defined as in ref. 70. For each unit,

$$A = \left(\sum_{i=1}^n r_i/n \right)^2 / \sum_{i=1}^n (r_i^2)/n, \quad [4]$$

where r_i is the unit response to stimulus i and n is the total number of stimuli. We used this measure in its inverted form S ,

$$S = 1 - A, \quad [5]$$

so that large values of S indicate high selectivity.

The response range G was defined, for each unit, as the difference between the maximum and minimum response across all stimuli:

$$G = \max_i(r_i) - \min_i(r_i). \quad [6]$$

All of the measures defined above were calculated for every 50-ms time interval within the trial. When the reported values refer to larger time windows, they represent averages over several 50-ms intervals.

Stimulus Classification. An instantaneous Naive Bayes decoder was trained and tested on individual time bins of population responses. The size of a bin was 50 ms, unless specified otherwise. We performed k -fold cross-validation by randomly subsampling the data ($k - 1$ data partitions used for training, 1 used for test, k repetitions; $k = 100$). The task of the decoder was to determine the stimulus identity for each test trial, based on the population response in a particular time bin. Chance level was $1/\text{number of stimuli} = 1/34$.

Support vector machines (SVMs) with quadratic kernels were applied using a similar k -fold cross-validation procedure ($k = 100$) for the computations shown in *SI Appendix, Fig. S7*. For each data split, we trained the SVMs on either intact or trial-shuffled data (shuffling across trials within stimulus condition) and tested them on intact data to test whether access to the correlation structure present in the data leads to better performance.

Self-Organizing Recurrent Network. The neural network model was composed of 80% excitatory ($N^E = 200$) and 20% inhibitory units ($N^I = 50$). Connectivity matrices W^{IE} , W^{EI} , and W^{II} were dense, randomly drawn from the interval $[0,1]$ and normalized so that the incoming connections to each neuron summed up to a constant ($\sum_j W_{ij} = 1$). The connections between

excitatory units W^{EE} were random and sparse and followed soft topological constraints ($p^{EE} = 0.1$ was the connection probability for neighboring units; i.e., every 10 consecutive units were considered neighbors; $p^{EE} = 0.01$ was the connection probability for nonneighbors). The threshold values for excitatory (T^E) and inhibitory units (T^I) were drawn from a uniform distribution in the intervals $[0, 0.5]$ and $[0, 0.3]$. The network state at time t was given by two binary vectors $x(t) \in \{0, 1\}^{N^E}$ and $y(t) \in \{0, 1\}^{N^I}$, representing activity of the excitatory and inhibitory units, respectively. Each timestep t corresponded to ≈ 20 ms of real time.

The network evolved using the following update functions:

$$x(t+1) = \theta(W^{EE}(t)x(t) - W^{EI}(t)y(t) + U(t)) - T^E(t) \quad [7]$$

- C. M. Lewis, A. Baldassarre, G. Comitteri, G. L. Romani, M. Corbetta, Learning sculpts the spontaneous activity of the resting human brain. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 17558–17563 (2009).
- W. Li, V. Piëch, C. D. Gilbert, Perceptual learning and top-down influences in primary visual cortex. *Nat. Neurosci.* **7**, 651–657 (2004).
- A. Schoups, R. Vogels, N. Qian, G. Orban, Practising orientation identification improves orientation coding in V1 neurons. *Nature* **412**, 549–553 (2001).
- P. Berkes, G. Orbán, M. Lengyel, J. Fiser, Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* **331**, 83–87 (2011).
- A. Kirkwood, M. C. Rioult, M. F. Bear, Experience-dependent modification of synaptic plasticity in visual cortex. *Nature* **381**, 526–528 (1996).
- C. D. Meliza, Y. Dan, Receptive-field modification in rat visual cortex induced by paired visual stimulation and single-cell spiking. *Neuron* **49**, 183–189 (2006).
- W. Singer, F. Tretter, Unusually large receptive fields in cats with restricted visual experience. *Exp. Brain Res.* **26**, 171–184 (1976).

$$y(t+1) = \theta(W^{IE}(t)x(t) - W^{II}(t)y(t) - T^I). \quad [8]$$

The Heaviside step function θ constrained the network activation at time t to a binary representation: A neuron fired if the total drive it received was greater than its threshold.

The stimulus set was composed of 10 digits, 6×6 pixels each. Every fifth excitatory unit received input from one corresponding image pixel; i.e., 36 units were input units, 164 units were reservoir units. The input $U(t)$ varied as a function of time (blue marking in Fig. 6): Initially it represented the luminance of the stimulus at a particular pixel location ("on" response, two time steps), and later it represented half the luminance of the reversed stimulus image at the same location ("off" response, seven time steps). For learning, we utilized a simple additive spike-timing-dependent plasticity (STDP) rule that increased (or decreased) the synaptic weight W^{EE} by a fixed amount $\eta_{STDP} = 0.001$ whenever unit i is active in the time step following (or preceding) activation of unit j :

$$\Delta W_{ij}^{EE}(t) = \eta_{STDP}(x_i(t)x_j(t-1) - x_i(t-1)x_j(t)). \quad [9]$$

In addition, synaptic normalization was used to proportionally adjust the values of incoming connections to a neuron so that they summed up to a constant value $c_E = c_I = 1$:

$$\Delta W_{ij}^{EE}(t) = c_E \left(W_{ij}^{EE}(t) / \sum_j W_{ij}^{EE}(t) \right) \quad [10]$$

$$\Delta W_{ij}^{EI}(t) = c_I \left(W_{ij}^{EI}(t) / \sum_j W_{ij}^{EI}(t) \right). \quad [11]$$

To stabilize learning, we used a homeostatic intrinsic plasticity (IP) rule that spread the activity evenly across units, by modulating their excitability using a learning rate $\eta_{IP} = 0.001$. At each time step, an active unit increased its threshold, while an inactive unit lowered its threshold by a small amount, such that on average each excitatory neuron fired with the target firing rate $\mu_{IP} = 0.1$:

$$\Delta T_i^E = \eta_{IP}(x_i(t) - \mu_{IP}). \quad [12]$$

Response facilitation was applied for a fixed time interval in each trial (20 time steps, shaded area in Fig. 6). Four different implementations were considered: 1) The incoming synaptic inhibition was lowered ($m_I = 0.5$), 2) the incoming synaptic excitation was increased ($m_E = 1.5$), 3) an additional noisy input was fed to all excitatory units (uniformly distributed in the interval $[0, 0.01]$), and 4) the target firing rate of excitatory units set via intrinsic plasticity was increased ($\eta_{IP} = 0.2$).

The self-organizing recurrent network (SORN) model was implemented in MATLAB. A similar implementation of the model in Python (with absent W^{II} connections) can be found at <https://github.com/chrhartm> (42).

Data Availability. All study data are in the main text and *SI Appendix*.

ACKNOWLEDGMENTS. We especially thank Thomas Wunderle and Jianguang Ni for technical assistance during recordings. D.N. and A.L. acknowledge grant support by the Deutsche Forschungs-Gemeinschaft (DFG) (NI708/5-1). P.F. acknowledges support by the DFG (SPP 1665) and the European Union (EU) (FP7-604 102-HBP). W.S. acknowledges the Reinhart Kosselleck grant of the German Research Foundation and the EU's 7th Framework Programme (FP7/2007-2013 Neuroseeker).

- S. F. Cooke, M. F. Bear, Visual experience induces long-term potentiation in the primary visual cortex. *J. Neurosci.* **30**, 16304–16313 (2010).
- S. F. Cooke, R. W. Komorowski, E. S. Kaplan, J. P. Gavornik, M. F. Bear, Visual recognition memory, manifested as long-term habituation, requires synaptic plasticity in V1. *Nat. Neurosci.* **18**, 262–271 (2015).
- V. Dragoi, J. Sharma, M. Sur, Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron* **28**, 287–298 (2000).
- M. Y. Frenkel *et al.*, Instructive effect of visual experience in mouse visual cortex. *Neuron* **51**, 339–349 (2006).
- U. R. Karmarkar, Y. Dan, Experience-dependent plasticity in adult visual cortex. *Neuron* **52**, 577–585 (2006).
- H. Yao, Y. Dan, Stimulus timing-dependent plasticity in cortical processing of orientation. *Neuron* **32**, 315–323 (2001).
- B. A. Olshausen, D. J. Field, How close are we to understanding v1? *Neural Comput.* **17**, 1665–1699 (2005).

15. A. Arieli, A. Sterkin, A. Grinvald, A. Aertsen, Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. *Science* **273**, 1868–1871 (1996).
16. A. Benucci, R. A. Frazor, M. Carandini, Standing waves and traveling waves distinguish two circuits in visual cortex. *Neuron* **55**, 103–117 (2007).
17. W. Xu, X. Huang, K. Takagaki, J. Y. Wu, Compression and reflection of visually evoked cortical waves. *Neuron* **55**, 119–129 (2007).
18. M. N. Havenith *et al.*, Synchrony makes neurons fire in sequence, and stimulus properties determine who is ahead. *J. Neurosci.* **31**, 8570–8584 (2011).
19. A. Benucci, D. L. Ringach, M. Carandini, Coding of stimulus sequences by population responses in visual cortex. *Nat. Neurosci.* **12**, 1317–1324 (2009).
20. K. Funayama *et al.*, Neocortical rebound depolarization enhances visual perception. *PLoS Biol.* **13**, e1002231 (2015).
21. X. Huang, S. Levine, M. A. Paradiso, Rebounding V1 activity and a new visual after effect. *J. Vis.* **8**, 25 (2008).
22. D. Nikolić, S. Häusler, W. Singer, W. Maass, Distributed fading memory for stimulus properties in the primary visual cortex. *PLoS Biol.* **7**, e1000260 (2009).
23. J. P. Gavornik, M. F. Bear, Learned spatiotemporal sequence recognition and prediction in primary visual cortex. *Nat. Neurosci.* **17**, 732–737 (2014).
24. S. Xu, W. Jiang, M. M. Poo, Y. Dan, Activity recall in a visual cortical ensemble. *Nat. Neurosci.* **15**, 449–455, S1–S2 (2012).
25. J. Cohen, *Statistical Power Analysis for the Behavioral Sciences*. (Academic Press, 1977).
26. M. M. Churchland *et al.*, Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nat. Neurosci.* **13**, 369–378 (2010).
27. M. R. Cohen, J. H. R. Maunsell, Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* **12**, 1594–1600 (2009).
28. A. Kohn, M. A. Smith, Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* **25**, 3661–3673 (2005).
29. B. Averbeck, P. Latham, A. Pouget, Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, 358–366 (2006).
30. A. Graf, A. Kohn, M. Jazayeri, J. A. Movshon, Decoding the activity of neuronal populations in macaque primary visual cortex. *Nat. Neurosci.* **14**, 239–245 (2011).
31. W. Maass, T. Natschläger, H. Markram, Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Comput.* **14**, 2531–2560 (2002).
32. H. Jaeger, “The “echo state” approach to analysing and training recurrent neural networks-with an erratum note” (Tech. Rep. 148, German National Research Center for Information Technology, Bonn, Germany, 2001).
33. J. J. Letzkus, S. B. Wolff, A. Lüthi, Disinhibition, a circuit mechanism for associative learning and memory. *Neuron* **88**, 264–276 (2015).
34. M. B. Pardi *et al.*, A thalamocortical top-down circuit for associative memory. *Science* **370**, 844–848 (2020).
35. Y. Gu *et al.*, Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron* **71**, 750–761 (2011).
36. A. M. Ni, D. A. Ruff, J. J. Alberts, J. Symmonds, M. R. Cohen, Learning and attention reveal a general relationship between population activity and behavior. *Science* **359**, 463–465 (2018).
37. E. L. Bartlett, X. Wang, Long-lasting modulation by stimulus context in primate auditory cortex. *J. Neurophysiol.* **94**, 83–104 (2005).
38. S. Sachidhanandam, V. Sreenivasan, A. Kyriakatos, Y. Kremer, C. C. H. Petersen, Membrane potential correlates of sensory perception in mouse barrel cortex. *Nat. Neurosci.* **16**, 1671–1677 (2013).
39. M. J. Wolff, J. Jochim, E. G. Akyürek, M. G. Stokes, Dynamic hidden states underlying working-memory-guided behavior. *Nat. Neurosci.* **20**, 864–871 (2017).
40. D. V. Buonomano, W. Maass, State-dependent computations: Spatiotemporal processing in cortical networks. *Nat. Rev. Neurosci.* **10**, 113–125 (2009).
41. A. Lazar, G. Pipa, J. Triesch, SORN: A self-organizing recurrent neural network. *Front. Comput. Neurosci.* **3**, 23 (2009).
42. C. Hartmann, A. Lazar, B. Nessler, J. Triesch, Where’s the noise? Key features of spontaneous activity and neural variability arise through learning in a deterministic network. *PLoS Comput. Biol.* **11**, e1004640 (2015).
43. D. Salaj *et al.*, Spike frequency adaptation supports network computations on temporally dispersed information. *eLife* **10**, e65459 (2021).
44. B. Nessler, M. Pfeiffer, L. Buesing, W. Maass, Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity. *PLoS Comput. Biol.* **9**, e1003037 (2013).
45. G. Bellec *et al.*, A solution to the learning dilemma for recurrent networks of spiking neurons. *Nat. Commun.* **11**, 3625 (2020).
46. C. Clopath, L. Büsing, E. Vasilaki, W. Gerstner, Connectivity reflects coding: A model of voltage-based STDP with homeostasis. *Nat. Neurosci.* **13**, 344–352 (2010).
47. Y. Sweeney, C. Clopath, Population coupling predicts the plasticity of stimulus responses in cortical circuits. *eLife* **9**, e56053 (2020).
48. V. Goudar, D. V. Buonomano, Encoding sensory and motor patterns as time-invariant trajectories in recurrent neural networks. *eLife* **7**, e31134 (2018).
49. B. Feulner, C. Clopath, Neural manifold under plasticity in a goal driven learning behaviour. *PLoS Comput. Biol.* **17**, e1008621 (2021).
50. A. Subramoney, F. Scherr, W. Maass, “Reservoirs learn to learn” in *Reservoir Computing*, K. Nakajima, I. Fischer, Eds. (Springer Singapore, 2021), pp. 59–76.
51. D. H. Hubel, T. N. Wiesel, Receptive fields of single neurones in the cat’s striate cortex. *J. Physiol.* **148**, 574–591 (1959).
52. X. Pei, T. R. Vidyasagar, M. Volgushev, O. D. Creutzfeldt, Receptive field analysis and orientation selectivity of postsynaptic potentials of simple cells in cat visual cortex. *J. Neurosci.* **14**, 7130–7140 (1994).
53. W. Singer, W. A. Phillips, Function and interaction of on and off transients in vision. II. Neurophysiology. *Exp. Brain Res.* **19**, 507–521 (1974).
54. W. A. Phillips, W. Singer, Function and interaction of on and off transients in vision. I. Psychophysics. *Exp. Brain Res.* **19**, 493–506 (1974).
55. M. F. Fong *et al.*, Distinct laminar requirements for NMDA receptors in experience-dependent visual cortical plasticity. *Cereb. Cortex* **30**, 2555–2572 (2020).
56. H. von Helmholtz, *Handbuch der Physiologischen Optik* (Voss, 1867).
57. C. D. Gilbert, T. N. Wiesel, Receptive field dynamics in adult primary visual cortex. *Nature* **356**, 150–152 (1992).
58. D. D. Stettler, A. Das, J. Bennett, C. D. Gilbert, Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neuron* **36**, 739–750 (2002).
59. S. Löwel, W. Singer, Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science* **255**, 209–212 (1992).
60. H. Barlow, “Cerebral cortex as model builder” in *Matters of Intelligence*, L. M. Vaina, Ed. (Synthese Library, Springer, Dordrecht, The Netherlands, 1987), pp. 395–406.
61. M. Bányaí *et al.*, Stimulus complexity shapes response correlations in primary visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 2723–2732 (2019).
62. C. Uran *et al.*, Predictive coding of natural images by v1 activity revealed by self-supervised deep neural networks. bioRxiv [Preprint] (2021). <https://doi.org/10.1101/2020.08.10.242958> (Accessed 10 August 2020).
63. M. Garrett *et al.*, Experience shapes activity dynamics and stimulus coding of VIP inhibitory cells. *eLife* **9**, e50340 (2020).
64. K. D. Harris, A. Thiele, Cortical state and attention. *Nat. Rev. Neurosci.* **12**, 509–523 (2011).
65. R. Oostenveld, P. Fries, E. Maris, J. M. Schoffelen, FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Comput. Intell. Neurosci.* **2011**, 156869 (2011).
66. K. H. Pettersen, A. Devor, I. Ulbert, A. M. Dale, G. T. Einevoll, Current-source density estimation based on inversion of electrostatic forward solution: Effects of finite extent of neuronal activity and conductivity discontinuities. *J. Neurosci. Methods* **154**, 116–133 (2006).
67. U. Mitzdorf, W. Singer, Excitatory synaptic ensemble properties in the visual cortex of the macaque monkey: A current source density analysis of electrically evoked potentials. *J. Comp. Neurol.* **187**, 71–83 (1979).
68. P. Berens *et al.*, A fast and simple population code for orientation in primate V1. *J. Neurosci.* **32**, 10618–10626 (2012).
69. U. Fano, Ionization yield of radiations. II. The fluctuations of the number of ions. *Phys. Rev.* **72**, 26–29 (1947).
70. E. T. Rolls, M. J. Tovee, Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *J. Neurophysiol.* **73**, 713–726 (1995).