

Visual exposure optimizes stimulus encoding in primary visual cortex

Andreea Lazar^{1,2,&,*}, Chris Lewis^{1,&,\$}, Pascal Fries¹, Wolf Singer^{1,2,3,#}, Danko Nikolić^{2,3,#}

1 Ernst Strüngmann Institute for Neuroscience, Frankfurt, Germany

2 Max Planck Institute for Brain Research, Frankfurt, Germany

3 Frankfurt Institute of Advanced Studies, Frankfurt, Germany

&, # These authors contributed equally

* Correspondence: andreea.lazar@brain.mpg.de

\$ Present address: Brain Research Institute, University of Zürich, Switzerland

Author Contributions

Conceptualization, A.L., W.S. and D.N.; Methodology, A.L., and D.N.; Investigation, A.L., C.L., and D.N.; Writing – Original Draft, A.L. and D.N.; Writing – Review & Editing, A.L.; C.L. and W.S.; Funding Acquisition, P.F., W.S., and D.N.; Resources, P.F., W.S., and D.N.; Supervision, W.S., and D.N.

Summary

Sensory exposure alters the response properties of individual neurons in primary sensory cortices. However, it remains unclear how these changes affect stimulus encoding by populations of sensory cells. Here, recording from populations of neurons in cat primary visual cortex, we demonstrate that visual exposure enhances stimulus encoding and discrimination. We find that repeated presentation of brief, high-contrast shapes results in a stereotyped, biphasic population response consisting of a short-latency transient, followed by a late and extended period of reverberatory activity. Visual exposure selectively improves the stimulus specificity of the reverberatory activity, by increasing the magnitude and decreasing the trial-to-trial variability of the neuronal response. Critically, this improved stimulus encoding is distributed across the population and depends on precise temporal coordination. Our findings provide evidence for the existence of an exposure-driven 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.

Introduction

Sensory experience alters the response properties of neurons and populations in sensory cortex. In the visual domain, repetitive exposure to oriented stimuli impacts the response strength and selectivity of early visual neurons, with lasting consequences on post-exposure activity (Cooke and Bear, 2010; Cooke et al., 2015; Dragoi et al., 2000; Frenkel et al., 2006; Karmarkar and Dan, 2006; Kirkwood et al., 1996; Meliza and Dan, 2006; Yao and Dan, 2001). Moreover, exposure to visual stimuli with temporal dynamics, i.e. moving bars and sequences of gratings, consolidates sequential firing across neurons in a stimulus-specific manner (Gavornik and Bear, 2014; Xu et al., 2012). While it is becoming increasingly clear that individual neurons in the primary visual cortex are flexible encoders that modify their filter characteristics through visual experience, it remains deeply puzzling

how these cells coordinate and constrain each other's responses, in meaningful ways, in order to extract improved representations of stimulus structure.

The primary visual cortex has traditionally been regarded as a passive filter bank converting sensory input into a sparse code for further feed-forward processing across the visual hierarchy. Yet, stimulus-evoked responses in the primary visual cortex propagate through the local circuitry in wave-like patterns (Benucci et al., 2007; Xu et al., 2007, Grinvald Arieli and Janke) and can persist long after the cessation of stimulation (Benucci et al., 2009; Funayama et al., 2015; Huang et al., 2008; Nikolić et al., 2009). These complex and temporally-extended population responses have been implicated in reward timing (Chubykin et al., 2013; Gavornik et al., 2009; Shuler and Bear, 2006), working memory (Harrison and Tong, 2009; Munneke et al., 2010; Supér et al., 2001) and are known to interact with (Benucci et al., 2009; Funayama et al., 2015; Gavornik and Bear, 2014; Nikolić et al., 2009; Wolff et al., 2017) and modulate the perception of (Brascamp et al., 2007; Fischer and Whitney, 2014; Funayama et al., 2015; Huang et al., 2008; Kahneman et al., 1992) subsequent visual stimulation. These dynamic properties exhibited by early visual neurons together with the exposure-dependent changes of stimulus-responses, suggest a direct involvement of primary visual cortex in the active distributed representation of more complex visual features, thus supporting a more constructive interpretation of primary cortex function (Olshausen and Field, 2005).

Here, we investigate the impact of visual exposure on the distributed response dynamics of populations of neurons, recorded simultaneously in cat area 17. We expand on previous studies on exposure-induced learning, by considering stimuli that are less redundant than oriented bars and gratings from an information perspective and thus better suited to capture aspects of distributed coding. In particular, we examine how brief, repetitive exposure to a large set of abstract visual shapes (letters of the Latin alphabet and Arabic numerals), affects the capacity of a hypothetical downstream decoder to identify the presented stimulus based on population activity. We employ brief stimulus presentations at high contrast, known to induce strong reverberatory population activity (Funayama et al., 2015; Huang et al., 2008; Nikolić et al., 2009) and report exposure-driven changes in single-unit response properties, i.e. stimulus discrimination, firing rate magnitude and variability. Importantly, we find that the performance of a classifier trained to predict stimulus identity, based on short vectors of data, is improved by stimulus exposure over the course of a recording session: early trials have lower performance scores than late trials in a session. By characterizing the correlations between pairs of neurons and the selectivity of low-dimensional stimulus representations extracted via principal component analysis, we show that this optimized encoding of stimulus structure relies on coordinated changes that affect both spatial and temporal aspects of population dynamics.

Results

We recorded the activity of neuronal populations from area 17 of lightly anesthetized cats using silicon-based multi-electrode arrays (5 animals, 11 recording sessions with independent electrode insertions). We applied standard thresholding and spike-sorting techniques in order to isolate action potentials from single neurons and multiunit clusters (Materials and Methods). The receptive fields (RFs) of the recorded units (27-52 units per session, 443 units in total) were located nearby in visual space and could be jointly stimulated by a single luminance stimulus, flashed for 100 milliseconds over a black

background (Figure 1A). We used a large set of abstract shapes (34 uppercase letters and digits) as stimuli, which, due to their inherent structure, differentially activated the population of neurons with differing RFs. All stimuli were presented 50 times, in random order.

Previous studies have shown that brief stimulus presentations at high contrast produce strong persistent activity in primary visual cortex (Funayama et al., 2015; Huang et al., 2008). In our data, the flashed visual stimulus evoked a biphasic population response (example in Figure 1A), with an initial transient component which started immediately after stimulus onset and a delayed reverberatory component which started ~200 ms after stimulus offset. The reverberatory response was remarkably long lasting: 296 out of 443 units (66.8%) fired above the baseline for the entire duration of the trial (one-tail t-test, $p < 0.05$).

We investigated the impact of visual exposure on these complex population responses from area 17, by subdividing each recording session into two or five consecutive trial-blocks and by comparing the *early* trials to the *late* trials in each session (schematic in Figure 1B).

Exposure triggers changes in single unit response properties

We began by measuring the impact of visual exposure on the ability of single units to discriminate visual shapes. For each unit, we calculated the discriminability index d' (Cohen, 1977), which quantifies the difference between the mean responses to different stimuli relative to the standard deviations of those responses across trials. When early and late trials in each session were compared, we found that visual exposure led to a significant increase in average d' across units (Figures 1C and D, 8.38% d' increase, paired t-test, $p = 6.4 \times 10^{-13}$, 100-800 ms, 443 units). The increase was particularly pronounced during the reverberatory part of the population response (10.86% increase for the interval 300-600 ms, paired t-test $p = 2.8 \times 10^{-12}$), while no significant change in d' was found during the early onset transient (2% increase for the interval 0-300 ms, paired t-test $p = 0.1$). Since this analysis compared only two blocks of trials, it obscured the temporal evolution of visual exposure effects. When sessions were split into 5 blocks of trials, we found that the average d' across units improved gradually with exposure and did not reach a saturation point (Figure 1E, 10.9% increase between block 1 and 5, paired t-test $p = 1.3 \times 10^{-7}$, the black line indicates the linear fit, $y = 2.4 + 98.66x$), suggesting that further improvements in d' may be possible with further exposure.

In a subset of sessions, the recordings were performed with 32-channel laminar arrays, allowing us to investigate how exposure driven changes in single-unit d' varied as a function of cortical depth (7 sessions, 289 units). We applied current source density (CSD) analysis by calculating the second spatial derivative of recorded voltages and estimated the location of the earliest current sinks which occur in thalamo-recipient layers 4 and 6 (example in Figure S1 A). We split the units based on their laminar location into infragranular (IG), granular (G) and supragranular (SG) units. We found that the exposure driven increase in stimulus discrimination was significant at all laminar depths (Figures 1F and S1; paired t-test, $p = 0.03$, SG; $p = 0.002$, G; $p = 1 \times 10^{-9}$, IG). Interestingly, the peak d' was significantly higher for SG and G units compared to IG units (one tail t-test, peak

SG>IG, $p = 0.002$ early trials, $p = 0.021$ late trials; peak G>IG, $p = 0.00007$ early trials, $p = 0.0001$ late trials).

Next we investigated how changes in single-unit d' relate to changes in other single-unit response properties. An increase in stimulus discrimination could be explained by an increased difference between mean responses to different stimuli, or by a reduction in response variability. We found that the mean firing rate across all 443 recorded units, increased significantly with visual exposure for the reverberatory part of the response (Figure 2A, 5.2% increase, paired t-test $p=7 \times 10^{-9}$, 100-800 ms; 6.2% increase, paired t-test $p=2 \times 10^{-8}$, 300-600 ms), while a significant decrease in firing rates was observed for the transient response (4% decrease, paired t-test $p=0.001$, 0-300 ms). In individual sessions, visual exposure led to a significant increase in mean firing rates in 9 out of 11 sessions (Figure S2, paired t-test, $p<0.05$). Over the course of the reverberatory response, the mean firing rate was correlated with the mean discriminability index, i.e. higher firing was associated with higher discrimination (Figure 2B, z-scored data from 11 sessions for the interval 300-600 ms, Spearman $r = 0.58$ early trials; $r = 0.36$ late trials). Interestingly, similar firing rates were associated with higher d' values for late compared to early trials in a session (regression lines $y = 0.47x - 0.35$ early trials; $y = 0.4x + 0.36$ late trials, Figure 2B), suggesting that the increased firing rate could not entirely account for the improvement in discrimination. When sessions were split into 5 blocks of trials we found a gradual increase in firing rate with visual exposure (Figure 2C, 5% increase between block 1 and 5, paired t-test $p = 1 \times 10^{-6}$, linear fit $y = 3.36x + 97.81$)

The variability of single unit responses across trials was quantified by a simple measure, the Fano factor, calculated as the spike-count variance divided by the spike-count mean (Fano, 1947). In agreement with previous findings (Churchland et al., 2010), we observed decline in neuronal variability following stimulus onset by an average 14.7% across all units (the baseline to trough drop in variability ranged from 6.4% to 37% across sessions). Comparing early and late trials in a session, we found that visual exposure led to a significant decrease in response variability (Figure 2D, 2.5% decrease, paired t-test $p=0.013$, 0-300 ms; 3.5% decrease, paired t-test $p=8 \times 10^{-5}$, 300-600 ms; 3.3% decrease, paired t-test $p=0.0001$, 100-800 ms). In addition, over the course of the reverberatory response, the mean firing rate variability across units was negatively correlated with the mean discriminability index (Figure 2E, z-scored data from 11 sessions for the interval 300-600 ms, Spearman $r = -0.30$ early trials; $r = -0.38$ late trials). In addition, similar variability levels were associated with enhanced discrimination for late trials compared to early trials in a session (regression lines $y = -0.29x - 0.33$ early trials; $y = -0.38x + 0.3$ late trials). When sessions were split into 5 blocks of trials we found that firing variability decreased gradually with visual exposure (Figure 2F, regression line $y = -0.40x + 99.76$ for the interval 300-600 ms).

Population encoding of visual shapes improves with visual exposure

We next investigated how visual exposure affects the capacity of a hypothetical downstream decoder to identify stimuli based on the output of primary visual cortex populations. We addressed this question directly by employing a decoding approach. To this end, average unit responses within time bins varying between 10-400 ms were

converted into activity vectors and a Bayesian classifier was trained to determine stimulus identity using a 100-fold cross-validation procedure (schematic in Figure 3A, see Materials and Methods for details). A separate classifier was trained for each time bin (e.g. 43 classifiers were trained for a 1100 ms long trial, with 50 ms time bins and 25 ms overlap between adjacent bins). In all sessions, classifiers trained on short 50 ms integration windows performed significantly above chance level (Figure 3B, data grouped by animal, chance level = 2.94% for 34 stimuli). In agreement with previous studies employing briefly flashed visual stimuli (Funayama et al., 2015; Nikolić et al., 2009; Volgushev et al., 1995), the population vectors of the reverberatory activity allowed for better classification of the stimuli than the initial transient response, indicating that intrinsic circuit computations improved the separability of stimulus specific population responses. Classification performance was also enhanced by the repetitive presentation of the set of visual shapes and this effect was confined to reverberatory activity. The time courses of performance along the trial varied between animals, but were similar for early and late trials within each session (individual performance profiles in Figure S3) and thus could be quantified by the area under the curve (AUC). When the experimental sessions from each animal were pooled together (Figure 3B), both the exposure driven increase in AUC for performance (average increase 33%, range 14-64%, t-test $p < 0.00006$) and the increase in peak performance (average increase 27.7%, range 13-59.6%, t-test $p < 0.0003$), were significant in all five animals. In individual sessions, visual exposure led to a significant increase in AUC for performance in 9 out of 11 sessions (average increase of 26%, t-test $p < 0.03$, Figure S3). Absolute peak performance scores ranged from 8 to 49.5% correct for early trials and 16.5 to 59.6% correct for late trials and increased significantly in 8 out of 11 experimental sessions (36% average increase, t-test $p < 0.016$), while no session showed a significant decrease.

The integration time window affected the difference in decoding performance between early and late trials (Figure 3C). Decoders that counted spikes over intermediate integration windows (50-200 ms) had large performance scores and showed significant differences in performance AUC between early and late trials (t-test, $p < 0.05$), while very short (10 ms) and very long (400 ms) integration windows, resulted in lower performance scores and reduced improvements in performance AUC (t-test, $p > 0.05$). Additionally, we investigated the impact of task difficulty on decoding performance, by varying the number of stimuli being decoded from 2 to 32 (Figure 3D). We found that classification of 8 or more visual stimuli led to significant differences between early and late peak performance (t-test, $p < 0.05$), whereas classification of 2 or 4 stimuli (t-test, $p > 0.05$) did not, as fewer stimuli led to ceiling effects (e.g. 2 class problems had early peak performance scores over 90% in 4 out of 11 sessions).

Previous studies have indicated that remarkably little stimulus-exposure can modify the response properties of primary visual cortex neurons. For example, as few as 100 repetitions of a moving light spot altered the internal dynamics of V1 neurons to such an extent that it enabled post-exposure cue-triggered recall (Xu et al., 2012). Given the large number of stimuli in our set, in most of our recordings, individual stimuli were repeated only 50 times. However, we tested the effects of a larger number of repetitions in a control session that consisted of 5100 trials (150 trials per stimulus). In this case, we observed an increase in performance AUC that continued past 50 repetitions per stimulus (Figure S4), suggesting that the observed effects can gain further strength with longer exposure.

Spike-count correlations are reduced by visual exposure

Having established that visual exposure benefits stimulus encoding in primary visual cortex, we next investigated how exposure altered the structure of population activity. A comprehensive analysis of the structure of high-dimensional distributions is difficult. As a first approach we examined how repeated exposure affects shared trial-to-trial response co-fluctuations between units: for each pair of simultaneously recorded units and each stimulus, we calculated the correlation coefficient between spike-count responses across trials. This measure, termed spike-count correlation or noise correlation, is known to impact the amount of information in a population code (Averbeck et al., 2006). Spike-count correlations are usually estimated over longer time intervals in order to avoid the occurrence of non-Gaussian distributions of spike-counts. Therefore, we calculated correlations over a 300 ms interval (300-600 ms after stimulus onset) corresponding to the period in the trial with high classification performance. We found that noise correlations were on average positive (Figure 4A, mean SC 0.095 ± 0.0015 SEM and 0.075 ± 0.0014 SEM, early and late trials respectively), indicating the existence of shared variability throughout the population. In addition, repeated exposure significantly reduced the strength of spike-count correlations (21% decrease, paired t-test, $p < 10^{-17}$) over the course of the session. In order to investigate how this reduction in shared variability depends on the stimulus preference of individual neurons, we calculated spike-count correlations as a function of signal correlations (Figure 4B). Signal correlations were computed as the correlations between the mean responses of neurons to the various stimuli. Consistent with previous studies (Cohen and Maunsell, 2009; Kohn and Smith, 2005), we found that spike-count correlations were highest for pairs of neurons with similar stimulus preferences (positive signal correlations) and lowest for pairs of neurons with opposing stimulus preferences (negative signal correlations). In our data, the effect of repeated exposure on spike-count correlations was present irrespective of the stimulus preference of the two neurons and it was stronger for neurons with opposing preferences (78% decrease, two-tailed t-test, $p < 10^{-9}$, signal correlations < -0.1 ; 9% decrease, two-tailed t-test, $p < 10^{-6}$, signal correlations > 0.1). Interestingly, attention has also been shown to decrease the strength of spike-count correlations, with a stronger effect on units with opposing stimulus preferences (Cohen and Maunsell, 2009).

From a decoding perspective, ignoring spike-count correlations can lead to a significant loss in performance (Averbeck et al., 2006; Graf et al., 2011). We found that a support vector machine with quadratic features trained on trial-shuffled data and tested on original data (spike counts over the same 300-600 ms interval, see Materials and Methods), performed more poorly compared to a decoder trained on the original data, with access to the intact correlation structure (Figure 4C, mean performance $22.5\% \pm 1.32$ SEM intact vs $19.7\% \pm 1.22$ SEM scrambled, early trials; $26.7\% \pm 1.44$ SEM intact vs $23.2\% \pm 1.37$ SEM scrambled, late trials). Interestingly, the performance loss was significant for both the early and late trials in a session (21.6 % decrease, paired t-test, $p = 0.04$ early trials; and 25.9 % decrease, paired t-test, $p = 0.045$ late trials), suggesting that while repeated exposure decreased the overall level of spike-count correlations in the data, a portion of spike-count correlations present in both early and late trials contributed positively to the population code. This finding is in line with a recent report in awake macaques showing that, when stimuli are structured, spike-count correlations are stimulus specific (Bányai et al., 2017). In addition, the performance of the classifier was significantly higher for late trials compared to early trials in a session both for the intact data and the scrambled data (21.4 % increase,

paired t-test, $p = 0.032$ intact data; and 18.5 % increase, paired t-test, $p = 0.008$ scrambled data). The increase in performance in the scrambled data is suggestive of changes in response properties that are independent of the correlation structure, e.g. the single neuron properties studied in the previous section.

Population responses cluster in low-dimensional projections

To gain insight into how exposure driven changes in pairwise correlations manifest themselves at the level of full population vectors we took a projection approach. The firing rates of a set of n neurons (n , number of units recorded in each session) at a particular time can be represented as a point in a n -dimensional vector space. We mapped these high-dimensional data into a low-dimensional projection space via principle component analysis (PCA). We tracked the response vectors corresponding to different visual stimuli at multiple time points along the trial (50 ms time bins, see Materials and Methods). Over the course of the trial, the population responses to different stimuli (Figure 5A, letters “A”, “B” and “C” and Figure 6a, 34 letters and digits), clustered together in the space defined by the first two principle components in a stimulus specific manner. After stimulus onset, the responses gradually segregated into specific subspaces. We found that these subspaces were more differentiable following repeated stimulus exposure: late trials showed enhanced segregation of stimulus-specific responses compared to early trials in a session (Figures 5A and 6A upper vs. lower panels; individual sessions in Figure S5). To quantify this segregation of clusters, we calculated for each data point the ratio between the Euclidean distance to its cluster center (defined as the average of all points evoked by the same stimulus) and the distance to the center of the data (defined as the average of all points irrespective of stimulus condition). This ratio (R) is close to zero if the points belonging to each stimulus cluster are well segregated in space and close to one if the data is spread randomly. For each session we picked the time point in the trial corresponding to the peak classification performance and calculated mean R across trials for each of the 34 clusters, i.e. 34 stimuli (Figure 5B). A scatter plot of R values corresponding to early and late trials for 34 stimuli in 11 sessions revealed that R decreased with stimulus exposure, i.e. stimulus clusters were better segregated for the late trials in a session. The change in R was significant (11% average decrease, paired one-tailed t-test, $p = 0.0029$, Figure 5C and D).

Next we determined whether the composition of principle components comes primarily from signal correlations (correlations between neurons in their mean responses to multiple stimuli), or from a combination of signal and spike-count correlations (correlations between neurons in their responses to repetitions of a single stimulus). We removed spike-count correlations by shuffling the data across trials, separately for each stimulus condition. Note that the temporal bin used for this analysis is much shorter (50 ms) compared to the bin used in the previous section for calculating pair-wise correlations (300 ms). In order to compare responses with identical distributions, we projected the original data in the principle component space determined with and without shuffling (compare Figure 6 A, B, see Materials and Methods for details). Such a shuffling, explicitly tests whether spike-count correlations are helpful in determining meaningful principle components, which allow for a good segregation of stimulus-specific responses. We quantified the quality of the segregation with the help of a Bayesian decoder trained on data projected in the principle component space determined with or without prior shuffling and tested on unshuffled data

(20-fold validation procedure, see Materials and Methods). In the space described by the first two principle components, the effect of trial shuffling was substantial (2 PCs, Figure 6C left, 6D, 21.3 % decrease AUC, paired t-test $p = 0.0001$, 28.9 % decrease peak performance $p = 0.0416$ for early trials; 24.6 % decrease AUC, paired t-test $p = 0.0025$, 23.1 % decrease peak performance $p = 0.018$ for late trials). In the space described by the first ten principle components, the detrimental effect of trial shuffling was even stronger (10 PCs, Figure 6C right, 6d, 50.3 % decrease AUC, paired t-test $p = 0.00018$, 51 % decrease peak performance $p = 0.00013$ for early trials; 54.5 % decrease AUC, paired t-test $p = 0.00007$, 49.8 % decrease peak performance $p = 0.00089$ for late trials). The effect of trial shuffling on the AUC was significantly stronger for late trials compared to early trials in a session (10 PCs, paired t-test $p = 0.00063$). This suggests that knowledge about the structure of spike-count correlations in early and, even more so, in late trials can significantly enhance decoding performance in low-dimensional projections. Interestingly, only a small drop in the amount of total explained variance was observed after trial shuffling (10 PCs, 5.3 % decrease, paired t test $p = 0.000007$ for early trials; 4.7 % decrease, paired t-test $p = 0.000007$ for late trials) in spite of a strong decrease in decoding performance (>50%). In addition, no significant differences were found in the amount of variance explained by the first 10 PCs between early and late trials in a session (paired t-test, $p = 0.3696$) in spite of a significant increase in peak performance (23.14 % decrease, paired t-test, $p = 0.0132$). The fact that the same number of PCs are required to explain the same amount of variance for both early and late signals, suggests that the advantages in performance come primarily from stimulus-specific constraints on the correlation structure of the population response. Such constraints affect both the correlation structure derived from the mean population responses to multiple stimuli (signal correlations, Figure 6C and 6D, blue and yellow bars), as well as the complete correlation structure including that derived from responses to repetitions of a single stimulus (signal and spike-count correlations, Figure 6C and 6D, gray and red bars).

Exposure optimizes both variant and invariant aspects of population dynamics

The segregation of evoked responses into localized stimulus-specific subspaces varied substantially over the course of the trial and peaked at different moments in time for different experimental sessions (Bayesian classifier Figure 3B and S3). We therefore examined the temporal dynamics of the population activity and investigated whether the evoked response to a stimulus was stable (invariant) or variable (variant) over the course of the trial, i.e. whether different time bins along the trial encode information about a stimulus in a similar or different manner. To this end, we concatenated the population activity vectors over three consecutive time bins of 50 ms each (Figure 7A). We calculated stimulus decoding-performance for the concatenated data and a control, where data were temporally scrambled. We found that temporally scrambled data had significantly lower performance scores compared to the original concatenated data (35% decrease, paired t-test $p = 6.78 \times 10^{-6}$ for early trials; 39% decrease, paired t-test $p = 1.93 \times 10^{-5}$ for late trials). The significant penalty incurred by scrambling data in the temporal domain implies that the encoding of visual stimuli varies at a fast time-scale, i.e. the representation of stimulus S at time t is different from that at time t+50 or t+100 ms. Taken together with the results from the previous sections, this suggests that the structure of the population response, in both the spatial and the temporal domain, is essential for a high quality encoding. Interestingly, for temporally scrambled data, the late trials preserved a significant advantage over the

early trials in a session, albeit reduced (Figure 7b, original concatenated data, 24% increase for late trials, paired t-test $p = 0.0087$; temporally scrambled data 15% increase for late trials, paired t-test $p = 0.0028$), suggesting that temporally invariant representations also benefit from stimulus exposure. Finally, a Bayesian classifier trained on 3 concatenated consecutive 50 ms bins significantly outperformed a classifier trained on a single large 150 ms bin (18% relative decrease, paired t-test $p = 3.14 \times 10^{-5}$, early trials; 23% decrease, paired t-test $p = 3.39 \times 10^{-7}$, late trials) emphasizing again the richness of information present in the temporal domain. Interestingly, the relative contribution of temporal structure to classification performance was stronger for late than early trials in a session, indicating a potential refinement of temporal information with stimulus exposure.

Discussion

We have demonstrated that repetitive exposure to briefly flashed visual shapes improves the discriminative capacity of primary visual cortex populations. Specifically, a Bayesian classifier trained to decode stimulus identity based on population vectors during brief temporal windows performed better during late trials as compared to early trials in a session. Classification performance was positively correlated with the mean firing rate and negatively correlated with the firing rate variability, both during the trial and over the course of a session. However, neither of these measures could fully account for the increase in classification performance: for sessions in which the firing rate didn't change across the session, decoding performance was still higher for late than for early trials. Rather, stimulus discriminability originated primarily from an increase in selectivity that was apparent both at the level of individual units, as quantified by d-prime, and at the level of the population, as captured by the low-dimensional projections.

Additionally, we found that stimulus exposure decreased the overall strength of correlated variability across trials, i.e. pairwise spike-count correlations were reduced with exposure. Such a reduction in shared variability can benefit stimulus encoding even in the absence of firing rate changes (Averbeck et al., 2006) and has been previously reported during attention and perceptual learning tasks (Cohen and Maunsell, 2009; Gu et al., 2011; Ni et al., 2018). Alternatively, an enhancement in signal, rather than a reduction in variability, could be the main driver behind exposure-driven changes in stimulus selectivity. The enhanced segregation of stimulus-specific clusters which we observed in the low-dimensional projections of our data, suggested a combination of both effects: stimulus-exposure reduced the cluster radius, indicative of a reduction in across-trial variability, and increased the overall spread of the data, and implicitly the distance between cluster centers, indicative of an increase in signal. Interestingly, in both early and late trials, the removal of spike-count correlations through trial shuffling, decreased classification performance, as well as the segregation of clusters in the PCA space, strongly suggesting that the intact correlation structure was advantageous for stimulus discrimination. Overall, the effect of stimulus exposure on spike-count correlations was complex: while the strength of correlations was reduced in late trials, knowledge about their structure was beneficial for stimulus discrimination both for the early and late trials in a session.

We found that a brief, high contrast stimulus resulted in a stereotyped, biphasic response, consisting of a high amplitude transient followed by a delayed reverberatory response. Decoding performance diverged from the stereotyped dependence on rate dynamics, with accuracy peaking approximately 300 ms after stimulus offset, during the reverberatory

component of the response. Additionally, performance tended to remain high for the entire duration of the trial, suggesting that a considerable amount of stimulus-specific information was maintained by the dynamic population response. Sustained and information-rich sensory responses, persisting beyond the period of sensory stimulation, have been reported previously in a number of different sensory modalities and species. In vivo whole-cell recordings from mouse primary auditory cortex during an oddball paradigm showed that excitatory neurons and parvalbumin-positive inhibitory interneurons exhibited a delayed response component, approximately 300 ms post-stimulus, which was modulated by stimulus content and carried signatures of deviance detection (Chen et al., 2015). In the primary auditory cortex of awake marmosets, preceding stimuli suppressed or facilitated responses to succeeding stimuli for durations greater than one second (Bartlett and Wang, 2005). In mouse primary sensory cortex, the early sensory response (<50 ms) to a single brief whisker deflection encoded stimulus information, while the later activity (50-400 ms) was shown to drive the subjective detection (Sachidhanandam et al., 2013). Notably, 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 (Funayama et al., 2015). The delayed activity exhibited high orientation selectivity and influenced the evoked response to subsequent inputs in an orientation-selective manner. Moreover, the influence had behavioral consequences: in a psychophysics task in which human subjects were asked to report the direction of motion of a drifting grating, their response latencies were modulated by a preceding matching or non-matching grating flash, presented 0.5 seconds earlier (Funayama et al., 2015). In a separate study, a simultaneous change in both stimulus and background gave rise to delayed activity in macaque V1 (Huang et al., 2008). The magnitude of the delayed response varied with the size of the background and was strongly correlated with the perception of a visual aftereffect (~300 ms post-stimulus) demonstrated through human psychophysics. Finally, in human EEG, information about a previously presented visual stimulus could be decoded from an impulse response, even in the absence of lingering delay activity (activity-silent states), for long intervals after stimulus presentation (Wolff et al., 2017). Taken together, these studies highlight a propensity for primary visual cortex to maintain sensory information, far beyond the temporal intervals required by the traditional feed-forward model of the ventral stream.

Instead, the findings outlined above and the results presented in the current manuscript are compatible with a dynamic coding framework for recurrent processing (Buonomano and Maass, 2009). 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). According to this theory, efficient recurrent processing relies on two simple requirements: (i) stimulus responses must persist beyond the duration of the stimulus, establishing a brief memory of recent events that can be integrated with novel incoming information (fading-memory property) and (ii) the temporal evolution of network states in response to different stimuli must result in reproducible stimulus-specific trajectories (separability property). Dynamic changes in hidden states through exposure to a specific set of inputs can presumably optimize the memory and separability properties exhibited by a recurrent circuit, by altering the network's stimulus-response mapping. In computational models, changes in hidden states via local experience-dependent plasticity rules were shown to explain numerous experimental

findings on cortical variability (Hartmann et al., 2015) and to increase the performance of recurrent networks on memory and prediction tasks (Lazar et al., 2009). Similarly, in our data, we found that experience-dependent changes gradually optimized the encoding of stimuli by primary sensory cortex populations. Given that the experiments were performed under anesthesia, the changes described here are likely to involve unsupervised “automatic” mechanisms, independent of attention and conscious control. However, it is also possible the anesthesia changes the dynamic regime of the cortex and further work is necessary to determine if the effects reported here persist in the waking state.

The network of connections responsible for the observed reverberation of visual responses and the functional changes underlying the marked increase in stimulus discrimination with stimulus exposure are still unknown. Given the presence of feedforward and feedback thalamo-cortical interactions (Hubel and Wiesel, 1959; Pei et al., 1994), we cannot exclude the possibility that exposure driven changes in primary cortex responses originate from interactions with subcortical structures. In fact, early vision 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 (Phillips and Singer, 1974; Singer and Phillips, 1974). However, in our data, stimulus exposure resulted in a stimulus-specific increase in response selectivity, suggests that experience dependent changes affecting the local recurrent interactions in primary visual cortex and/or the long-distance recurrent interactions with higher cortical areas, are more likely responsible for the observed effects. The complex recurrent dynamics of the early visual system arise on the backbone of an intricate connectivity structure (Gilbert and Wiesel, 1992; Stettler et al., 2002), which has been refined during development to capture the statistical properties of the visual environment (Barlow, 1987; Berkes et al., 2011; Helmholtz, 1867; Löwel and Singer, 1992; Singer and Treutner, 1976). The structure and the synaptic weights of both local and long-range connectivity in the visual cortex reflect regularities present in visual scenes and thus are likely to serve as implicit knowledge for the processing of sensory evidence. As our data suggest, visual exposure appears to shape the internal network dynamics in a manner that results in a refinement of sensory coding, akin to perceptual learning (Crist et al., 2001; Gilbert et al., 2009; Li et al., 2004; Sagi and Tanne, 1994; Schoups et al., 2001; Vogels and Orban, 1985; Yan et al., 2014; Seitz and Watanabe, 2009; Watanabe et al., 2001). As such, the primary visual cortex appears to optimize its processing of visual information as a function of prior experience via its specifically wired, reverberating network in order to provide a highly selective representation of familiar stimuli during the late response phase. Within the primary visual cortex this optimization process was apparent across all laminae and seemed to favor the supragranular and granular compartments.

From a functional perspective, the immediate responses of primary visual cortex neurons to feed-forward thalamic input carry information about simple visual features such as orientation, spatial frequency and motion (Hubel and Wiesel 1962). In contrast, the delayed responses that originate from recurrent interactions, carry information about more global aspects of scene organization (Ito and Gilbert, 1999; Lamme and Roelfsema, 2000) and are modulated by stimulus history (Benucci et al., 2009; Funayama et al., 2015; Gavornik and Bear, 2014; Huang et al., 2008; Nikolić et al., 2009; Volgushev et al., 1995), task context (Gilbert and Li, 2012; Li et al., 2004), behavioral state (Bradley et al., 2003), reward expectation (Chubykin et al., 2013; Gavornik et al., 2009; Shuler and Bear, 2006) and

sensory exposure (Cooke and Bear, 2010; Dragoi et al., 2000; Frenkel et al., 2006; Gavornik and Bear, 2014; Karmarkar and Dan, 2006; Xu et al., 2012; Yao and Dan, 2001). The complexity of these responses indicates that the primary visual cortex may contribute, at least partially, to functions traditionally attributed to higher order cortical areas (Olshausen and Field, 2005). Our study adds further support to the idea that neurons in the primary visual cortex are flexible encoders that alter their responses through visual experience. In particular, we have provided compelling evidence that exposure to a large set of visual shapes optimizes the population encoding in early visual cortex, resulting in a more efficient readout of stimulus-specific information. These findings suggest that the efficient visual discrimination of familiar stimuli can be partially achieved through separation of neuronal representations at the earliest cortical stage of sensory processing.

Materials and Methods

Electrophysiological recordings and data processing

Data was recorded from five adult cats under general anesthesia during terminal experiments in two separate laboratories. The general methods have been described thoroughly in previous publications (Ni et al., 2016; Nikolić et al., 2009). 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 N₂O:O₂ (70/30%) and halothane (0.5%–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 was recorded with multiple silicon-based 16-channel probes from the Center for Neural Communication Technology at the University of Michigan (each probe consisted of 4 shanks, 3 mm long, 200 μm distance, 4 contact points each, 1,250 μm² area, 0.3–0.5 MΩ impedance at 1 kHz). To extract multi-unit activity, signals were amplified 10,000 and filtered between 500 and 3500 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 N₂O:O₂ (60/40%) and isoflurane (0.6%–1.0%). After verifying narcosis, Vecuronium (0.25mg/kg/h i.v.) was added for paralysis. Data was collected via multiple 32-contact probes (100 μm inter-site spacing, ~1 MΩ at 1 kHz; NeuroNexus or ATLAS Neuroengineering) and amplified (Tucker Davis Technologies, FL). Signals were filtered with a passband of 700 to 7000 Hz and a threshold was set interactively to retain multi-unit activity.

Spike sorting

The sorting of the recorded multi-units was performed offline via custom software that computed principal components of spike waveforms in order to reduce dimensionality and grouped the resulting data using a density-based clustering algorithm (DBSCAN). Both the well-isolated cells and the remaining multi-units, were included in the analysis.

Visual stimuli

Stimuli consisted of 34 shapes: 26 letters (A–Z) and 8 digits (0–7). They were white on black background and spanned approximately 5–7 degrees of visual angle. In all sessions, we recorded 50 repetitions per stimulus; more than 200 repetitions per stimulus were recorded in session 042414. Trials were 1200 ms long with stimulus onset at 500 ms; stimulus onset was at 600 ms in session col10c05.

Current source density analysis

In 3 cats (7 sessions, 289 units), the recordings were performed with 32-channel linear arrays (100 micron spacing). Local field potentials to moving grating stimuli presented at maximal contrast were recorded either immediately before or immediately after the sessions with letters and digits. In these data, we applied current source density (CSD) analysis using a standard algorithm (Pettersen et al., 2006) 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 4Cα (Mitzdorf and Singer, 1979).

Stimulus classification

MATLAB and the Statistics Toolbox (The MathWorks, Inc.) were used for data analysis. An instantaneous Naïve 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 cross-validation by randomly subsampling the data ($k-1\%$ training, 1% test, $k = 100$ repetitions). 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 $1/\text{number of stimuli} = 1/34$.

A SVM decoder with quadratic kernels was used in a similar manner for the computations shown in Figure 4. Here, we followed the analysis suggested in Averbeck et al., 2006 and tested whether a decoder would perform better or worse if not given access to the correlation structure present in the data. To this end, we trained our decoder on trial-shuffled data (shuffling across trials within stimulus condition), tested on original data and compared this scenario with one in which the decoder was trained and tested on the original data.

Principle component analysis (PCA)

PCA was performed in order to visualize the high dimensional activity vectors corresponding to different visual stimuli. PCA was computed independently on 50 ms time bins of normalized spike count vectors for several relevant time windows in the trial (-100 ms pre-stimulus activity, +100 ms, +300, +500 and +700 ms after stimulus onset). For Figure 5, PCA was computed on the trials corresponding to letters A, B and C. For Figure 6, PCA was computed on the complete set of trials (34 stimuli).

In several cases, we made modifications to the data handled by the decoder in order to test how the readout quality would be affected by those changes. For Figure 6, PCA was performed either on the original data (PCA original) or on trial-shuffled data (PCA shuffled). Independent trials of original data were then projected in these two PCA spaces. For both

projections, a Bayesian decoder was employed to separate the representations of population responses to different stimuli in the space described by either the first two or the first ten principal components. A k-fold validation method was employed (k=20): 19 parts of data were used to construct the PCA original and PCA shuffled spaces and to train the classifier. One part of data (original, not-shuffled) was used to test the performance of the classifier in the two projection spaces. This procedure was repeated 20 times, to obtain a reliable average.

For Figure 7, data was concatenated over three consecutive 50 ms time bins, which increased the dimensionality of the data to three times the number of units. In order to test the specificity of temporal dynamics, the performance of a Bayesian decoder trained on concatenated data was compared to that of a decoder trained on temporally scrambled data (temporal scrambling over three consecutive time bins) and with that of a decoder trained on spike counts calculated over 150 ms.

Supplemental Information

The supplemental information for this article includes five figures.

Acknowledgements

Special thanks go to Thomas Wunderle and Jianguang Ni for technical assistance during recordings. DN and AL acknowledge grant support by the Deutsche Forschungsgemeinschaft (DFG NI708/5-1). PF acknowledges grant support by DFG (SPP 1665), EU (FP7-604102-HBP), and LOEWE (NeFF). WS acknowledges the Reinhart Koselleck grant of the German Research Foundation and the European Union's 7th Framework Programme (FP7/2007-2013 Neuroseeker).

References

- Averbeck, B.B., Latham, P.E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* 7, 358.
- Bányai, M., Lazar, A., Klein, L., Klon-Lipok, J., Singer, W., and Orbán, G. (2017). Stimulus content shapes cortical response statistics. *BioRxiv* 238774.
- Bartlett, E.L., and Wang, X. (2005). Long-Lasting Modulation by Stimulus Context in Primate Auditory Cortex. *J. Neurophysiol.* 94, 83–104.
- Benucci, A., Frazor, R.A., and Carandini, M. (2007). Standing Waves and Traveling Waves Distinguish Two Circuits in Visual Cortex. *Neuron* 55, 103–117.
- Benucci, A., Ringach, D.L., and Carandini, M. (2009). Coding of stimulus sequences by population responses in visual cortex. *Nat. Neurosci.* 12, 1317.
- Brascamp, J.W., Knapen, T.H.J., Kanai, R., Ee, R. van, and Berg, A.V. van den (2007). Flash suppression and flash facilitation in binocular rivalry. *J. Vis.* 7, 12–12.
- Buonomano, D.V., and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nat. Rev. Neurosci.* 10, 113–125.

Chen, I.-W., Helmchen, F., and Lütcke, H. (2015). Specific Early and Late Oddball-Evoked Responses in Excitatory and Inhibitory Neurons of Mouse Auditory Cortex. *J. Neurosci. Off. J. Soc. Neurosci.* 35, 12560–12573.

Chubykin, A.A., Roach, E.B., Bear, M.F., and Shuler, M.G.H. (2013). A Cholinergic Mechanism for Reward Timing within Primary Visual Cortex. *Neuron* 77, 723–735.

Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13, 369.

Cohen, J. (1977). *Statistical power analysis for the behavioral sciences* (Academic Press).

Cohen, M.R., and Maunsell, J.H.R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* 12, 1594.

Cooke, S.F., and Bear, M.F. (2010). Visual Experience Induces Long-Term Potentiation in the Primary Visual Cortex. *J. Neurosci.* 30, 16304–16313.

Cooke, S.F., Komorowski, R.W., Kaplan, E.S., Gavnornik, J.P., and Bear, M.F. (2015). Visual recognition memory, manifested as long-term habituation, requires synaptic plasticity in V1. *Nat. Neurosci.* 18, 262.

Dragoi, V., Sharma, J., and Sur, M. (2000). Adaptation-Induced Plasticity of Orientation Tuning in Adult Visual Cortex. *Neuron* 28, 287–298.

Fano, U. (1947). Ionization Yield of Radiations. II. The Fluctuations of the Number of Ions. *Phys. Rev.* 72, 26–29.

Fischer, J., and Whitney, D. (2014). Serial dependence in visual perception. *Nat. Neurosci.* 17, 738.

Frenkel, M.Y., Sawtell, N.B., Diogo, A.C.M., Yoon, B., Neve, R.L., and Bear, M.F. (2006). Instructive Effect of Visual Experience in Mouse Visual Cortex. *Neuron* 51, 339–349.

Funayama, K., Minamisawa, G., Matsumoto, N., Ban, H., Chan, A.W., Matsuki, N., Murphy, T.H., and Ikegaya, Y. (2015). Neocortical Rebound Depolarization Enhances Visual Perception. *PLOS Biol.* 13, e1002231.

Gavnornik, J.P., and Bear, M.F. (2014). Learned spatiotemporal sequence recognition and prediction in primary visual cortex. *Nat. Neurosci.* 17, 732.

Gavnornik, J.P., Shuler, M.G.H., Loewenstein, Y., Bear, M.F., and Shouval, H.Z. (2009). Learning reward timing in cortex through reward dependent expression of synaptic plasticity. *Proc. Natl. Acad. Sci.* 106, 6826–6831.

Gilbert, C.D., and Li, W. (2012). Adult visual cortical plasticity. *Neuron* 75, 250–264.

Graf, A.B.A., Kohn, A., Jazayeri, M., and Movshon, J.A. (2011). Decoding the activity of neuronal populations in macaque primary visual cortex. *Nat. Neurosci.* 14, 239.

Gu, Y., Liu, S., Fetsch, C.R., Yang, Y., Fok, S., Sunkara, A., DeAngelis, G.C., and Angelaki, D.E. (2011). Perceptual Learning Reduces Interneuronal Correlations in Macaque Visual Cortex. *Neuron* 71, 750–761.

- Harrison, S.A., and Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458, 632.
- Hartmann, C., Lazar, A., Nessler, B., and Triesch, J. (2015). Where's the Noise? Key Features of Spontaneous Activity and Neural Variability Arise through Learning in a Deterministic Network. *PLOS Comput. Biol.* 11, e1004640.
- Huang, X., Levine, S., and Paradiso, M.A. (2008). Rebounding V1 activity and a new visual aftereffect. *J. Vis.* 8, 25–25.
- Hubel, D.H., and Wiesel, T.N. (1959). Receptive fields of single neurones in the cat's striate cortex. *J. Physiol.* 148, 574–591.
- Ito, M., and Gilbert, C.D. (1999). Attention Modulates Contextual Influences in the Primary Visual Cortex of Alert Monkeys. *Neuron* 22, 593–604.
- Kahneman, D., Treisman, A., and Gibbs, B.J. (1992). The reviewing of object files: Object-specific integration of information. *Cognit. Psychol.* 24, 175–219.
- Karmarkar, U.R., and Dan, Y. (2006). Experience-Dependent Plasticity in Adult Visual Cortex. *Neuron* 52, 577–585.
- Kirkwood, A., Rioult, M.G., and Bear, M.F. (1996). Experience-dependent modification of synaptic plasticity in visual cortex. *Nature* 381, 526–528.
- Kohn, A., and Smith, M.A. (2005). Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci. Off. J. Soc. Neurosci.* 25, 3661–3673.
- Lamme, V.A.F., and Roelfsema, P.R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci.* 23, 571–579.
- Lazar, A., Pipa, G., and Triesch, J. (2009). SORN: A Self-Organizing Recurrent Neural Network. *Front. Comput. Neurosci.* 3.
- Li, W., Piëch, V., and Gilbert, C.D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nat. Neurosci.* 7, 651.
- Meliza, C.D., and Dan, Y. (2006). Receptive-Field Modification in Rat Visual Cortex Induced by Paired Visual Stimulation and Single-Cell Spiking. *Neuron* 49, 183–189.
- Mitzdorf, U., and Singer, W. (1979). 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.
- Munneke, J., Heslenfeld, D.J., and Theeuwes, J. (2010). Spatial working memory effects in early visual cortex. *Brain Cogn.* 72, 368–377.
- Ni, A.M., Ruff, D.A., Alberts, J.J., Symmonds, J., and Cohen, M.R. (2018). Learning and attention reveal a general relationship between population activity and behavior. *Science* 359, 463–465.
- Ni, J., Wunderle, T., Lewis, C.M., Desimone, R., Diester, I., and Fries, P. (2016). Gamma-Rhythmic Gain Modulation. *Neuron* 92, 240–251.
- Nikolić, D., Häusler, S., Singer, W., and Maass, W. (2009). Distributed Fading Memory for Stimulus Properties in the Primary Visual Cortex. *PLOS Biol.* 7, e1000260.

Olshausen, B.A., and Field, D.J. (2005). How close are we to understanding v1? *Neural Comput.* **17**, 1665–1699.

Pei, X., Vidyasagar, T.R., Volgushev, M., and Creutzfeldt, O.D. (1994). Receptive field analysis and orientation selectivity of postsynaptic potentials of simple cells in cat visual cortex. *J. Neurosci. Off. J. Soc. Neurosci.* **14**, 7130–7140.

Pettersen, K.H., Devor, A., Ulbert, I., Dale, A.M., and Einevoll, G.T. (2006). 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.

Phillips, W.A., and Singer, W. (1974). Function and interaction of on and off transients in vision I. *Psychophysics. Exp. Brain Res.* **19**, 493–506.

Sachidhanandam, S., Sreenivasan, V., Kyriakatos, A., Kremer, Y., and Petersen, C.C.H. (2013). Membrane potential correlates of sensory perception in mouse barrel cortex. *Nat. Neurosci.* **16**, 1671–1677.

Shuler, M.G., and Bear, M.F. (2006). Reward Timing in the Primary Visual Cortex. *Science* **311**, 1606–1609.

Singer, W., and Phillips, W.A. (1974). Function and interaction of on and off transients in vision II. *Neurophysiology. Exp. Brain Res.* **19**, 507–521.

Supèr, H., Spekreijse, H., and Lamme, V.A.F. (2001). A Neural Correlate of Working Memory in the Monkey Primary Visual Cortex. *Science* **293**, 120–124.

Volgushev, M., Vidyasagar, T.R., and Pei, X. (1995). Dynamics of the orientation tuning of postsynaptic potentials in the cat visual cortex. *Vis. Neurosci.* **12**, 621–628.

Wolff, M.J., Jochim, J., Akyürek, E.G., and Stokes, M.G. (2017). Dynamic hidden states underlying working-memory-guided behavior. *Nat. Neurosci.* **20**, 864.

Xu, S., Jiang, W., Poo, M., and Dan, Y. (2012). Activity recall in a visual cortical ensemble. *Nat. Neurosci.* **15**, 449.

Xu, W., Huang, X., Takagaki, K., and Wu, J. (2007). Compression and Reflection of Visually Evoked Cortical Waves. *Neuron* **55**, 119–129.

Yao, H., and Dan, Y. (2001). Stimulus Timing-Dependent Plasticity in Cortical Processing of Orientation. *Neuron* **32**, 315–323.

Figures and figure legends




Figure 1. Visual exposure enhances stimulus discrimination by primary sensory neurons. (A) Cluster of receptive fields of simultaneously recorded multi-units (rectangles) relative to the location of a visual stimulus and peristimulus time histogram of population responses to stimulus "A". Responses display a transient component following stimulus presentation and a long-lasting reverberatory component. (B) In each session, 34 visual stimuli were presented 50 times each, in random order. Sessions were split in 2 or 5 consecutive blocks of trials and analyzed separately. (C) Discriminability index (d') for early (black) and late (red) trials in a session (2 trial-blocks, 11 sessions, 443 units). (D) Box plots showing effect of visual exposure on the distribution of d' values across 443 units, for the 0-300 ms, 300-600 ms and 100-800 ms time-intervals. (E) Gradual increase in discrimination power with of visual exposure (5 trial-blocks, average d' over the 300-600 ms interval). (F) Laminar analysis in a subset of sessions (289 units): box plots show changes in d' across units from supragranular (SG), granular (G) and infragranular (IG) laminae (complete d' profiles in Figure S1).




Figure 2. Visual exposure increases firing rates and decreases firing rate variability.

(A) Average firing rates for early and late trials (matched baseline) showed significant increases in firing rates with visual exposure for the intervals 100-800 ms and 300-600 ms. **(B)** Normalized firing rates for 6 consecutive 50 ms temporal bins (300-600 ms) against the corresponding d' scores, superimposed for all 11 recording sessions. Similar firing rate levels resulted in higher d' scores in late trials (red) compared to early trials (black) in a session. **(C)** Gradual increase in firing rates with visual exposure for the interval 300-600 ms (5 exposure blocks). **(D)** Response variability (Fano factor) for early and late trials (matched baseline) showed significant decreases in variability with visual exposure for the intervals 100-800 ms and 300-600 ms. **(E)** Normalized Fano factor scores plotted against d' scores (similar to B). Similar FF levels resulted in higher d' scores in late trials (red) compared to early trials (black). **(F)** Gradual decrease in response variability with visual exposure for the interval 300-600 ms (5 session blocks).




Figure 3. Visual exposure increases stimulus classification performance. (A) Illustration of decoding method: spike-count vectors over fixed time bins were used to train a Naïve Bayes classifier in order to estimate the probability of a given stimulus ‘S’ given the response ‘R’ (see Materials and Methods for details). (B) Stimulus decoding performance was substantially above chance level (dotted line) for the reverberatory part of the population response (50 ms time bins, 25 ms sliding window). Grouped data by animal showed significant increases in performance peak and AUC in each individual animal: early trials (black) had lower performance scores compared to late trials (red) in a session. (C) Performance AUC using the integration bin size denoted on the x-axis for early and late trials. Performance is highest for intermediate integration intervals. (D) Peak classification performance as a function of task difficulty (number of stimuli being classified), relative to chance level (dotted lines). Differences between early and late trials are more pronounced for harder classification problems.




Figure 4. Visual exposure reduces spike-count correlations. (A) Mean spike-count correlations across trials as a function of signal correlations, which here measure the similarity of two neurons' responses across stimulus conditions. Spike-counts were calculated over a 300 ms interval corresponding to the reverberatory part of the evoked response (300-600 ms, late trials (red) have lower correlation scores compared to early trials (black) in a session. Margins: Standard error of the mean. **(B)** Performance of a SVM with quadratic kernel trained with (original data, black and red) or without (shuffled data, gray and orange) intact correlation structure and tested on original data.




Figure 5. Visual exposure enhances the segregation of population responses into stimulus-specific clusters. (A) Early and late evolution of population responses to three visual stimuli (letters A, B and C) in the projection space described by the first two principle components, based on 50 ms time bins. Each marker represents the population activity (spike count in 50 ms) from a single trial to one of these three visual stimuli. Individual stimuli form segregated clusters, 300 ms after stimulus onset. The segregation is more pronounced for late trials compared to early trials in the experiment. **(B)** Scatter of average R values for all clusters in all sessions. For each cluster condition, the average R calculates the mean distance from all points to the cluster center divided by the distance to the center of all data. Late trials have lower R values compared to early trials in a session, suggesting an increased segregation of data points into clusters. **(C)** The average 11% drop in mean R is significant. **(D)** When calculated separately for each session, the drop in R is significant in 6 sessions.




Figure 6. Clustering of population responses is sensitive to the correlation structure exhibited in the data. (A) Early and late evolution of population responses to 34 visual stimuli (letters and digits) in the projection space described by the first two principle components, based on 50 ms time bins. **(B)** Trial-scrambling has a weak effect in the space described by the first two principle components. **(C)** Decoding the representations of stimuli in the space described by the first two (left), or ten (middle) principle components, reveals the importance of intact trial-specific population dynamics (34 stimuli, 11 experimental sessions). **(D)** Decoding performance as a function of the number of principle components for early and late trials, original and scrambled data (34 stimuli, 11 experiments). Note that for the original data, the advantage in classification performance does not originate from a difference in the amount of total variance explained.




Figure 7. Visual exposure enhances both temporally specific and temporally invariant aspects of stimulus encoding. (A) Schematic representation of scrambling procedure: in each session, the performance of a Naïve Bayes decoder trained on the spike-counts of N units in three consecutive 50 ms time bins (red or magenta) is compared to temporally scrambled data (blue) and to spike-counts over a single 150 ms time bin (black). (B) Over the course of the trial, decoding performance (three consecutive 50 ms time bins) is higher for late trials (red) compared to early trials (magenta) in a session (average over 11 sessions, shaded areas indicate s.e.m.). Temporally scrambled data has lower performance scores, but preserves a significant difference between early (dark blue) and late trials (light blue). (C) Differences in performance AUC are significant between early and late trials for all three conditions (11 sessions). The performance of a classifier trained on three consecutive 50 ms bins is significantly higher than that of a classifier trained on a single 150 ms bin, suggesting that responses are temporally specific. In addition, the improvement in classification performance based on temporally scrambled data suggests that also temporally invariant aspects of stimulus encoding benefit significantly from visual exposure.