

# The contribution of response correlations to the neural code of V1

Mihály Bányai, Marcell Stippinger, Dávid Szalai, Gergő Orbán

{banyai.mihaly, stippinger.marcell, szalai.david, orban.gergo}@wigner.mta.hu  
Computational Systems Neuroscience Lab, MTA Wigner RCP, Budapest, Hungary

Andreea Lazar, Liane Klein, Johanna Klon-Lipok, Wolf Singer

{andreea.lazar, liane.klein, johanna.klon-lipok, wolf.singer}@brain.mpg.de  
Ernst Strüngmann Institute for Neuroscience in Cooperation with Max Planck Society, Frankfurt, Germany

## Abstract

**Contribution of joint statistics of neuron populations to stimulus encoding can distinguish theories of neural computation. Specifically, probabilistic inference in a hierarchical model of perception predicts the emergence of content-specific modulations in the fine structure of spike count correlations. By recording spiking activity from the V1 of behaving macaques viewing naturalistic and synthetic stimuli we demonstrate that compositional objects elicit correlational structures that are more specific to the identity of the stimulus than stimuli without structured content. Further, we demonstrate that decoding schemes exploiting stimulus-specific pairwise response statistics outperform those relying on marginal statistics, thus showing that joint statistics carry information about the stimulus independently from marginal statistics. To rule out possible simpler explanations of the observed patterns in the correlation structure, we introduce an array of controls. We develop Contrastive Rate Matching to control for firing rate-related changes in correlation magnitudes. Further, we analyze phenomenological models of noise correlations, the raster marginal model and a family of models featuring collective additive and/or multiplicative noise sources. Our results show that stimulus-dependence of noise correlations at the level of V1 reflect high-order structure in the stimulus, is independent of changes in firing rates and cannot be explained by phenomenological accounts.**

**Keywords:** hierarchical perception; response correlations

## Introduction

Computation in the visual cortex is a hierarchical inference process leading from the elementary decomposition of the retinal image to object recognition (Lee & Mumford, 2003). Probabilistic inference of a latent variable in a hierarchical model involves combining information derived from the stimulus (through the likelihood function) and the results of inferences at higher levels (constituting a prior for low-level inferences). Therefore top-down computational effects arising due to the inference of higher-level, compositional variables, such as objects or textures, play a crucial role in the inference of lower-level, elementary variables such as oriented edges. Such effects are observed behaviourally during the statistical learning of visual patterns (Orbán, Fiser, Aslin, & Lengyel, 2008) and the perception of illusory contours (Kanizsa, 1979).

The latter is also demonstrated to have a neural correlate as a feedback effect from V2 to V1, measured in mean neural responses (Lee & Nguyen, 2001), but how top-down effects shape joint response statistics, and specifically and correlations, has not been elucidated.

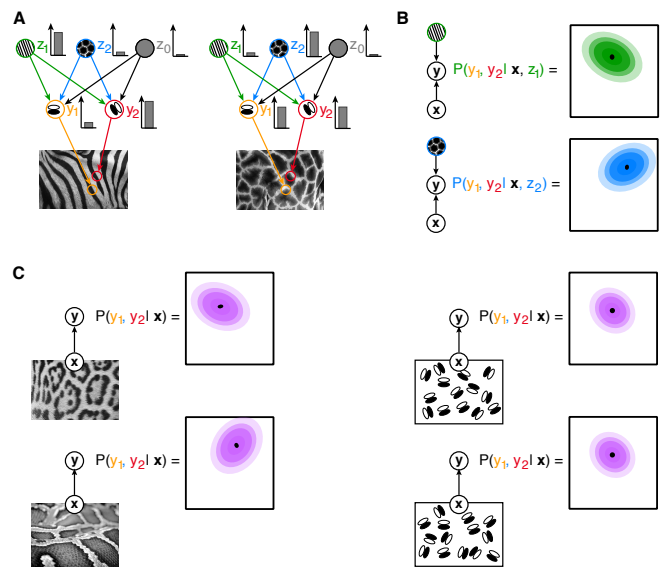


Figure 1: **A.** A hierarchical model of early vision describes how low-level features are combined to create high-level percepts. **B.** Through feedback, the posterior distribution of low-level features changes if the high-level percept is different, both in mean and correlational structure. **C.** Different images that evoke different high-level percepts elicit more different correlational structures in low-level features than different images evoking no high-level percept.

The structure and computational role of neural spike count correlations (SCCs), also called noise correlations, have been controversial. Regularities of SCCs related to stimuli have been investigated mostly focusing on the population mean (Kohn & Smith, 2005; Cohen & Kohn, 2011; Ecker et al., 2010), and found to be modest in awake animals. The dependence of neural correlations on some aspects of stimulus statistics has been demonstrated (Rikhye & Sur, 2015), and top-down effects in SCCs have been established related to fluctuations in attentional state (Ruff & Cohen, 2014; Bondy, Haefner, & Cumming, 2018). In this paper, we aim to demonstrate the stimulus-dependence of the fine structure of SCCs

is modulated by hierarchical perceptual inference.

Neural response variability in the visual cortex has been demonstrated to encode perceptual uncertainty using probabilistic models (Orbán, Berkes, Fiser, & Lengyel, 2016). By extending the model of V1 to take into account the hierarchical computations taking place in the visual cortex, we can make predictions regarding the covariability of neural responses as well. Assuming such a generative model of natural images, the posterior distribution of V1-level variables will depend on the inferred value of higher level variables (Fig 1A), which impose stimulus-specific SCC structures on them (Fig 1B).

Importantly, our model remains agnostic about the network-level mechanism of the SCC modulations. The effect of high-level percepts is top-down in the computational sense, but the feedback from higher-level visual areas to V1 is but one of the possibilities to implement it. Structured correlations have been shown to arise due to lateral connectivity as well (Rosenbaum, Smith, Kohn, Rubin, & Doiron, 2017). In both of these scenarios, the internal dynamics of cortical processing introduces structured patterns in SCCs according to the result of the probabilistic inference of image content (Berkes, Orbán, Lengyel, & Fiser, 2011).

The key prediction of the hierarchical model of perception is that the difference between SCCs in response to different images will be proportional to the difference in the degree to which specific higher-level variables are activated by them, i.e. how specific the higher level statistics are to the contrasted stimuli. Specifically, two images with different recognizable content (expected to elicit different higher-level inferences) will elicit more different SCCs than two similarly different images that don't have any recognizable content, and elicit similar higher-level inferences (Fig 1C).

In order to be behaviourally relevant, content-specific SCCs need to carry information about stimuli towards downstream structures. Similarly to other aspects of neural correlations, whether they can help decoding has been controversial (Averbeck, Latham, & Pouget, 2006). In awake animals, it has been found that correlation-blind decoders are hard to outperform (Berens et al., 2012). However, as decoding parametric grating stimuli only makes use of a low-dimensional manifold of neural activity, it may obscure contributions of higher-order response statistics (Gao & Ganguli, 2015) that could be revealed using naturalistic stimuli. Furthermore, designs with many different stimuli may suffer from data limitations, and recording during passive viewing may lack the complete response structure that could be observable during task engagement. Therefore we predict that SCCs will contribute more to the decodability of structured stimuli than that of structure-free ones.

Finally, we need to assess whether the observed regularities in SCCs can emerge from simple phenomenological models instead of hierarchical perception. Such models have been used to characterise neural correlations in sensory cortices (Okun et al., 2012; Lin, Okun, Carandini, & Harris, 2015), and provide a structured control to our prediction.

## Methods

Briefly, we recorded multiunit spiking activity from the V1 of two macaques performing a visual attention task. The task was used only for maintaining task engagement and only the portion of the trial was used for analysis where stimulus was shown but attentional cue was not presented. In addition to natural images, we designed synthetic stimuli by combining Gabor filters with coefficients sampled from the marginal empirical filter activation distribution calculated from natural images but did not have any higher-order statistical structure (LL-synthetic, Fig 2A). We also generated synthetic images with second-order statistical structure in the Gabor coefficients calculated from filter responses to texture-like photos (HL-synthetic, Fig 2B) to contrast them with structure-free stimuli. In order to be able to estimate correlations precisely, we used 6 or 8 images per session, allowing for about 120 repetitions per stimulus.

A critical component of the analysis concerned the firing rate dependence of correlations since different degrees of specificity in mean responses to different stimuli can fundamentally affect our conclusions. The magnitude of measured cortical SCCs depend significantly on the magnitude of firing rates (Schulz, Sahani, & Carandini, 2015), due to the nonlinear effect of the firing threshold (de La Rocha, Doiron, Shea-Brown, Josić, & Reyes, 2007). We introduce Contrastive Rate Matching (CRM) based on Churchland et al. (2010), where we take the difference in both geometric mean rates (GMR) and SCCs in response to all stimulus pairs in two categories (e.g. natural and synthetic) from all channel pairs, and subsample this data in such way that the resulting data set has identical GMR distributions for both categories, and only assess correlation differences in the subsampled data.

## Results

### Modulation of SCC specificity by stimulus content

We calculated the dissimilarity of the fine structure of SCCs by assessing the absolute difference of SCC matrices of all channel pairs in response to natural and LL-synthetic images and applied CRM to control for the possible confound of category-specific firing rates. Contrasting these conditions, SCCs display significantly higher specificity to natural images than to LL-synthetic ones (Fig 2A) after CRM, in line with our predictions (Fig 1C). To test if it is the difference in high-level content of the images and not changes in spectral content that produces the modulation, we designed synthetic stimuli that featured the minimal complexity that elicits specific responses in V2, synthetic texture images Freeman, Ziemba, Heeger, Simoncelli, and Movshon (2013). We show that the higher stimulus-specificity of SCCs after CRM appears in response to HL-synthetic images as well, corroborating the idea that compositional features represented at higher levels in the computational hierarchy can give rise to stimulus-specificity in SCCs.

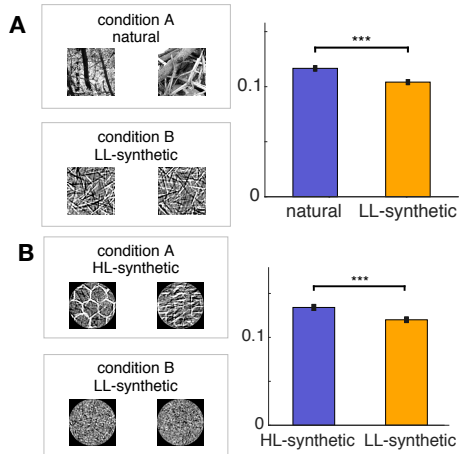


Figure 2: **A.** Stimulus-specificity of SCCs is greater in response to natural stimuli than to synthetic stimuli lacking statistical structure beyond V1-level representation. **B.** Adding higher-order structure to synthetic stimuli by combining Gabor filters to texture patterns reintroduces greater stimulus-specificity of SCCs

### The role of correlations in decoding

After assessing the modulation of correlations by the statistical structure of the stimulus, we investigated whether such differences are relevant for downstream computation in the visual processing hierarchy, and ultimately behaviour as well, by assessing the contribution of correlations to the decodability of stimulus identity. First, we z-transformed the spike count data in order to isolate the contribution of SCCs from those of the mean responses, and decoded the stimulus using a logistic regression model consisting of pairwise products of channel scores as predictors. We show that stimuli are decodable from the scores above chance level for both natural and LL-synthetic images, but with greater precision in the case of the former (Fig 3A). This analysis however only demonstrates that stimulus-related information is present in the correlations, but not that it is not also present in firing rates. So we also compared the performance of two likelihood-based mixture of Gaussians decoders on the untransformed data, one only able to utilise information in mean responses (linear), and the other using information from both mean responses and SCCs (quadratic). We show that natural images are decodable with a higher performance using a quadratic decoder than a linear one (Fig 3B), which difference is not present in the case of LL-synthetic images (Fig 3C). This result provides a confirmation that SCC-specificities convey extra information beyond that carried by the mean of the responses.

### Comparison to phenomenological models

While the results presented in Figs 2 & 3 provide support to hierarchical probabilistic inference in the visual cortex (Fig. 1), it remains open whether the same effect is reproducible with simple mechanisms without the introduction of top-down

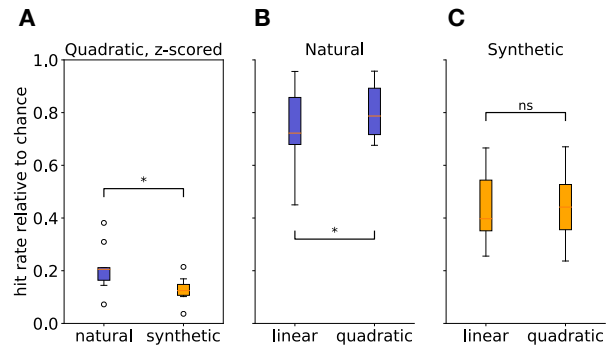


Figure 3: **A.** When removing the mean from responses by z-scoring, a logistic regression decoder with quadratic coefficients can still recover stimulus identity with an above chance, hit rate which is higher for natural images. Vertical axis is (hit rate - chance) / (1 - chance). **B.** When decoding the complete response statistics based on likelihoods from a mixture of Gaussians model, natural images are more identifiable if the model incorporates correlations as well instead of only mean responses. **C.** The contribution of correlations for decodability is absent for structureless synthetic images.

modulation in a perceptual hierarchy. In order to assess this question, we fitted phenomenological models to the measured spike counts, simulated synthetic datasets from them, and calculated the SCC differences in the same way as for the measurement data. First, we used the raster marginal model (RMM) introduced by Okun et al. (2012). This model simulates spike trains by fixing the channel means and the population counts in small time bins. We demonstrate that the SCC specificity measured in response to natural images is extremely unlikely under the RMM model based on 500 simulations, more so than the one in response to LL-synthetic images (Fig 4A). Second, we implemented the affine model introduced by Lin et al. (2015). This model describes response variability as a combination of two global fluctuations, a multiplicative and an additive one. We demonstrate that spike trains synthesized from the affine model (resampled 100 times) show stimulus-specificity in SCCs, but lack the modulation by higher-level stimulus statistics (Fig 4B). These results support our claim that the SCC-specificity modulation is a signature of perceptual computations in the visual cortex.

### Conclusions

We have shown that spike count correlation patterns are more stimulus-specific in response to structured stimuli than structure-free ones in the V1 of behaving monkeys after controlling for the specificity of firing rates. We demonstrated that the decodability of stimulus identity is enhanced by the increased specificity of correlations. We have shown that phenomenological models parametrised by marginal firing statistics and global fluctuations do not reproduce the measured effect, but it is predicted by the hierarchical model of visual perception.

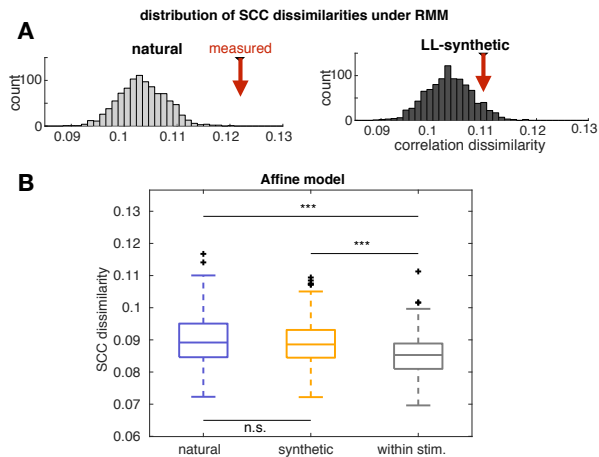


Figure 4: **A.** Raster marginal models, controlling for finite data effects on correlations, parametrised by measured average firing rates and spike co-occurrence counts do not account for measured SCC specificities. **B.** Synthetic data generated from the affine model fitted to measured responses is able to reproduce the stimulus-specificity of SCCs, but not the difference in specificity related to higher-order statistical structure.

## Acknowledgments

This work was supported by a Lendület Award (G.O.), the National Brain Program (NAP-B KTIA NAP 12-2-201, G.O.), the Deutsche Forschungsgemeinschaft (DFG NI 708/5-1, A.L.) and the European Union's 7th Framework Programme (FP7/2007-2013 Neuroseeker, W.S.).

## References

- Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5), 358–366.
- Berens, P., Ecker, A. S., Cotton, R. J., Ma, W. J., Bethge, M., & Tolias, A. S. (2012, August). A fast and simple population code for orientation in primate V1. *Journal of Neuroscience*, 32(31), 10618–10626.
- Berkes, P., Orbán, G., Lengyel, M., & Fiser, J. (2011, January). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, 331(6013), 83–87.
- Bondy, A. G., Haefner, R. M., & Cumming, B. G. (2018). Feedback determines the structure of correlated variability in primary visual cortex. *Nature neuroscience*, 21(4), 598.
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., ... Shenoy, K. V. (2010, February). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience*, 13(3), 369–378.
- Cohen, M. R., & Kohn, A. (2011, June). Measuring and interpreting neuronal correlations. *Nature Neuroscience*, 14(7), 811–819.
- de La Rocha, J., Doiron, B., Shea-Brown, E., Josić, K., & Reyes, A. (2007, August). Correlation between neural spike trains increases with firing rate. *Nature*, 448(7155), 802–806.
- Ecker, A. S., Berens, P., Keliris, G. A., Bethge, M., Logothetis, N. K., & Tolias, A. S. (2010, January). Decorrelated neuronal firing in cortical microcircuits. *Science*, 327(5965), 584–587.
- Freeman, J., Ziemba, C. M., Heeger, D. J., Simoncelli, E. P., & Movshon, J. A. (2013, May). A functional and perceptual signature of the second visual area in primates. *Nature Neuroscience*, 16(7), 974–981.
- Gao, P., & Ganguli, S. (2015). On simplicity and complexity in the brave new world of large-scale neuroscience. *Current opinion in neurobiology*, 32, 148–155.
- Kanizsa, G. (1979). *Organization in vision: Essays on gestalt perception* (Vol. 49). Praeger New York.
- Kohn, A., & Smith, M. A. (2005, April). Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *Journal of Neuroscience*, 25(14), 3661–3673.
- Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am. A*, 20(7), 1434–1448.
- Lee, T. S., & Nguyen, M. (2001, February). Dynamics of subjective contour formation in the early visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 98(4), 1907–1911.
- Lin, I.-C., Okun, M., Carandini, M., & Harris, K. D. (2015). The nature of shared cortical variability. *Neuron*, 87(3), 644–656.
- Okun, M., Yger, P., Marguet, S. L., Gerard-Mercier, F., Benucci, A., Katzner, S., ... Harris, K. D. (2012). Population Rate Dynamics and Multineuron Firing Patterns in Sensory Cortex. *Journal of Neuroscience*, 32, 17108–17119.
- Orbán, G., Berkes, P., Fiser, J., & Lengyel, M. (2016). Neural variability and sampling-based probabilistic representations in the visual cortex. *Neuron*, 92(2), 530–543.
- Orbán, G., Fiser, J., Aslin, R. N., & Lengyel, M. (2008, February). Bayesian learning of visual chunks by human observers. *Proceedings of the National Academy of Sciences of the United States of America*, 105(7), 2745–2750.
- Rikhye, R. V., & Sur, M. (2015). Spatial correlations in natural scenes modulate response reliability in mouse visual cortex. *Journal of Neuroscience*, 35(43), 14661–14680.
- Rosenbaum, R., Smith, M. A., Kohn, A., Rubin, J. E., & Doiron, B. (2017). The spatial structure of correlated neuronal variability. *Nature neuroscience*, 20(1), 107.
- Ruff, D. A., & Cohen, M. R. (2014). Attention can either increase or decrease spike count correlations in visual cortex. *Nature neuroscience*, 17(11), 1591.
- Schulz, D. P., Sahani, M., & Carandini, M. (2015). Five key factors determining pairwise correlations in visual cortex. *Journal of Neurophysiology*, 114(2), 1022–1033.