Rat sensitivity to multipoint statistics is predicted by efficient coding of natural scenes

Riccardo Caramellino[1,#], Eugenio Piasini[2,#], Andrea Buccellato[1], Anna Carboncino[1], Vijay Balasubramanian[2,*], Davide Zoccolan[1,*]

[1]*Visual Neuroscience Lab, International School for Advanced Studies (SISSA), Trieste, 34136, Italy*

[2]*Computational Neuroscience Initiative, University of Pennsylvania, Philadelphia, PA 19104, United States*

[#] these authors contributed equally to this work

* Corresponding authors:

Davide Zoccolan
Visual Neuroscience Laboratory
International School for Advanced Studies (SISSA)
Via Bonomea 265, 34136 Trieste (TS) ITALY
Office: +39 040 3787 701 Fax: +39 040 3787 702
E-mail: zoccolan@sissa.it

Vijay Balasubramanian
Computational Neuroscience Initiative
University of Pennsylvania
3700 Hamilton Walk, Philadelphia PA 19104 (USA)
Office: +1 215 573 0982
E-mail: vijay@physics.upenn.edu

Efficient processing of sensory data requires adapting the neuronal encoding strategy to the statistics of natural stimuli. Humans, for instance, are most sensitive to multipoint correlations that vary the most across natural images. Here we show that rats possess the same sensitivity ranking to multipoint statistics as humans, thus extending a classic demonstration of efficient coding to a species where neuronal and developmental processes can be interrogated and causally manipulated.

It is widely believed that the tuning of sensory neurons is determined by the need to efficiently represent the statistical structure of the signals they must encode[1]. This normative principle, known as efficient coding, has been successful in explaining many aspects of neural processing in vision[2–6], audition[7,8] and olfaction[9], and in accounting for general mechanisms that are shared across sensory modalities, such as adaptation[10] and gain control[11]. At the perceptual level, strong support for the role played by efficient coding as an organizing principle comes from psychophysical studies showing that human sensitivity to visual textures defined by local multipoint correlations depends on the variability of such correlations across natural scenes[12,13]. This suggests that, as a whole, the visual system preferentially encodes features that are less predictable in the environment, and thus are more informative about its state. However, it remains unknown whether this is a general principle underlying perceptual sensitivity to image statistics across species. Understanding this would be a critical step towards determining the neural substrates of efficient coding and the role of postnatal visual experience in giving rise to it, as it would enable further investigation with invasive recording techniques and altered rearing paradigms[14–16].

To address this question, we measured rat sensitivity to visual textures defined by local multipoint correlations. We stayed as close as possible to the paradigm employed in previous human psychophysics experiments[12,13], training the animals to discriminate "white noise" binary visual textures from textures that contained structured noise. White noise textures were generated by sampling each pixel independently, setting it to black or white with equal probability. As such, they contained no spatial correlations. Structured textures, on the other hand, were designed to enable precise control over the type and intensity of the correlations they contained. To generate these textures we built and published a software library[17] that implements the method developed in Victor and Conte (2012)[18]. Briefly, for any given type of multipoint correlation pattern (also termed a *statistic* in what follows), we sampled from the distribution over binary textures that had the desired probability of occurrence of that statistic, but otherwise contained the least amount of structure (i.e., had maximum entropy). The probability of occurrence of the pattern was parametrized by the *intensity* of the corresponding statistic, determined by a parity count of white or black pixels inside tiles of 1, 2, 3 or 4 pixels (termed *gliders*) used as the building blocks of the texture[18]. When the intensity is zero, the texture does not contain any structure – it is the same as white noise (Fig. 1a, left). When the intensity is +1, every possible placement of the glider across the texture contains an even number of white pixels, while a level of -1 corresponds to all

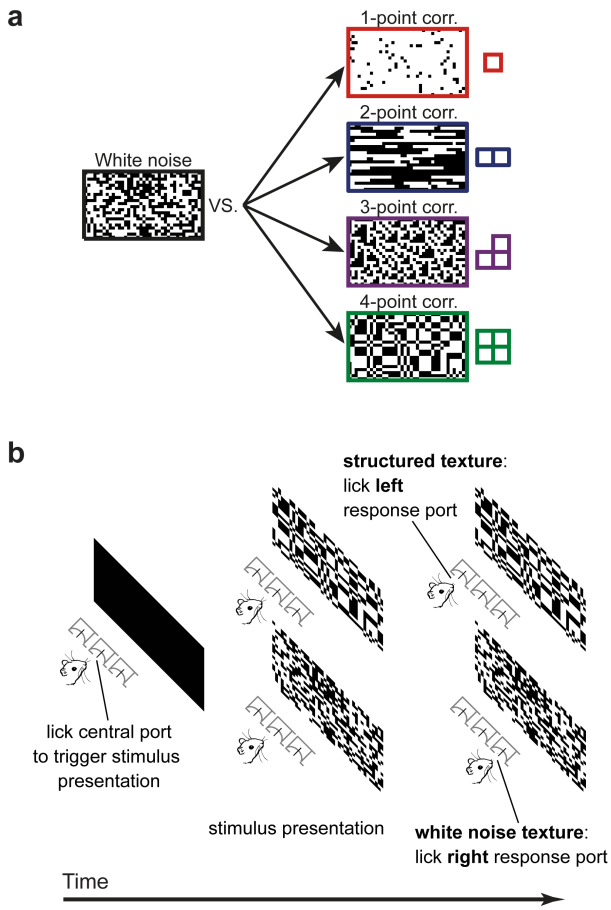

Fig. 1 | Visual stimuli and behavioral task. a, Schematic of the four kinds of texture discrimination tasks administered to the four groups of rats in our study. Each group had to discriminate unstructured binary textures containing white noise (example on the left) from structured binary textures containing specific types of local multipoint correlations among nearby pixels (i.e., 1-, 2-, 3- or 4-point correlations; examples on the right). The textures were constructed to be as random as possible (maximum entropy), under the constraint that the strength of a given type of correlation matched a desired level. The strength of a correlation pattern was quantified by the value (intensity) of a corresponding statistic (see main text), which could range from 0 (white noise) to 1 (maximum possible amount of correlation). The examples shown here correspond to intensities of 0.85 (1- and 2-point statistics) and 0.95 (3- and 4-point statistics). **b**, Schematic representation of a behavioral trial. Left and center: animals initiated the presentation of a stimulus by licking the central response port placed in front of them. This prompted the presentation of either a structured (top) or an unstructured (bottom) texture. Right: in order to receive the reward, animals had to lick either the left or right response port to report whether the stimulus contained the statistic (top) or the noise (bottom).

placements containing an odd number of white pixels. Intermediate intensity levels correspond to intermediate fractions of gliders containing the even parity count. The structure of the glider dictates the appearance of the final texture – a 1-point glider produces textures with different luminance biases; a 2-point glider produces oriented edges; a 3-point glider produces L-shape patterns; and a 4-point glider produces rectangular blocks (see Fig. 1a, right for examples of textures, with the corresponding glider shown on the right).

For each of these 1- to 4-point image statistics, we trained a group of rats to discriminate between white noise and textures containing that statistic with an intensity greater than zero (Fig. 1a). The experiment had four phases (see Methods): 1) an initial learning phase, where we used a single intensity level of the statistic, close to the maximum (between 0.85 and 0.95); 2) an additional training phase, where we introduced a new, lower level of the statistic every time the rat achieved a certain criterion performance on the previous level using a staircase procedure; 3) a main testing phase, where textures were sampled at regular intervals along the intensity level axis and were randomly presented to the rats; and 4) a further testing phase, where rats originally trained with a given statistic were probed with a different one. Each trial of the experiment started with the rat autonomously triggering the presentation of a stimulus by licking the central response port within an array of three (Fig.1b). The animal

then reported whether the texture displayed over the monitor placed in front of him contained the statistic (by licking the left port) or white noise (by licking the right port). The rat received liquid reward for correct choices and was subjected to a time-out period for incorrect ones (Fig. 1b).

The main test phase yielded psychometric curves showing the sensitivity of each animal in discriminating white noise from the structured texture with the assigned statistic (example in Fig. 2a, black dots). To interpret results, we developed an ideal observer model, in which the presentation of a texture with a level of the statistic equal to *s* produces a percept *x* sampled from a truncated Gaussian distribution centred on the actual value of the statistic (*s*) with a fixed standard deviation σ that only depends on the identity of the animal[19,20]. Here, σ measures the "blurriness" in the animal's sensory representation for a particular type of statistic (i.e., the perceptual noise) and, consequently, its inverse 1/σ captures its resolution, or sensitivity. This resolution determines the perceptual threshold for discriminating a structured texture from white noise, because the statistic *s* must be at least a distance ~σ away from 0 in order to be perceived as different from an unstructured image. The support of the truncated Gaussian was constrained to the interval [-1,1], as the value of a statistic is limited to that range by construction. We assumed that each rat had its own bias α, so that the animal's prior probability of white noise presentation is $p(s=0)=\frac{1}{1+e^{-\alpha}}$, and the prior probability of patterned texture is $p(s>0)=\frac{1}{K(1+e^{\alpha})}$ for each of the *K* nonzero levels of the statistic. Given these assumptions, for each percept *x* we can define a decision variable $D(x)=ln\frac{p(s=0|x)}{p(s>0|x)}$, where $p(s=0|x)$ and $p(s>0|x)$ are the posterior probabilities of a given percept *x* being produced by a white noise pattern ($s=0$) or by a structured texture ($s>0$). In our model, the rat always reports the most likely answer, so the answer is "noise" when $D>0$ and "structured texture" otherwise. From $D(x)$ we can derive the psychometric function giving the probability of responding "noise" at any given level of the statistic *s* as

$$p(reportnoise|s)=\frac{\Phi\left[\frac{x^*(\alpha,\sigma)-s}{\sigma}\right]-\Phi\left[\frac{-1-s}{\sigma}\right]}{\Phi\left[\frac{1-s}{\sigma}\right]-\Phi\left[\frac{-1-s}{\sigma}\right]},$$

where $\Phi(x)$ is the standard Normal cumulative density function and $x^*(\alpha,\sigma)$ is the value of *x* such that $D(x^*)=0$. In other words, x^* is the decision boundary used by the animal to divide the perceptual axis into "noise" and "structured texture" regions (see Methods). The psychometric function has a sigmoidal shape, which is characterized by the two free parameters of the model (α and σ) that can be estimated from the behavioral data by maximum likelihood (a psychometric function fitted to the choices of an example rat is shown in Fig. 2a, blue curve). The prior bias α is related to the horizontal offset of the curve, while the sensitivity 1/σ is related to its slope (see Methods for an example which makes this intuition more precise).

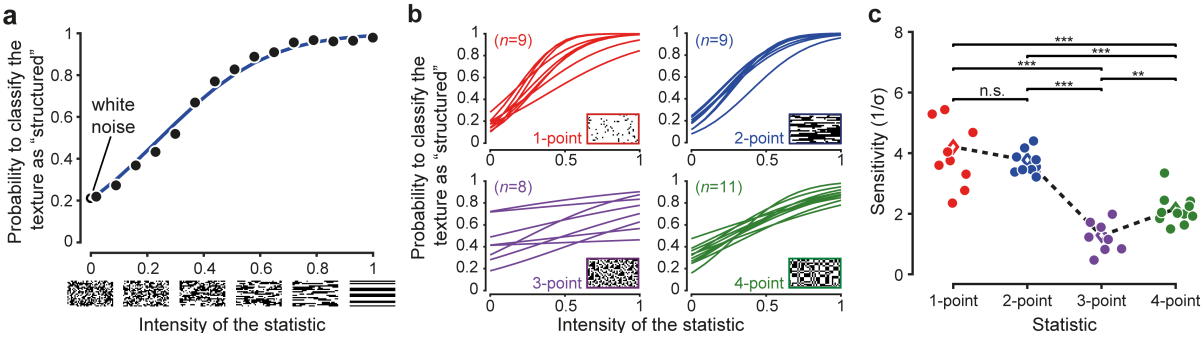

Fig. 2 | Rat sensitivity to multipoint correlations. a, Psychometric data for an example rat trained on 2-point correlations. Black dots: fraction of trials in which a texture with the corresponding intensity of the statistic was correctly classified as "structured". Empty black circle: fraction of trials the rat has judged a white noise texture as containing the statistic. Blue line: psychometric function corresponding to the fitted ideal observer model (see main text). **b**, Psychometric functions obtained for all the rats tested on the four statistics (*n* indicates the number of animals in each group). **c**, Values of the perceptual sensitivity 1/σ to each of the four statistics. Filled dots: individual rat estimates. Empty diamonds: group averages. The dashed line emphasizes the sensitivity ranking observed for the four statistics. Significance markers ** and *** indicate, respectively, $p < 0.01$ and $p < 0.001$ for a two-sample t-test with Holm-Bonferroni correction.

Fitting this model to the behavioral choices of rats in the four groups led to psychometric functions with a characteristic shape, which depended on the order of the multipoint statistic an animal had to discriminate (Fig. 2b). In particular, focusing on the values obtained for σ revealed the presence of a specific ranking among the groups. This is shown in Fig. 2c, where we report the sensitivity 1/σ for the individual rats within each group (dots), as well as their group averages (diamonds). We found that the sensitivity to 1- and 2-point correlations was higher than the sensitivity to both 3-point ($p_1<0.001$ and $p_2<0.001$, two sample t-test with Holm-Bonferroni correction) and 4-point ($p_1<0.001$, $p_2<0.001$) correlations. In addition, the sensitivity to 4-point correlations was higher than that to 3-point correlations ($p<0.01$).

To further validate these findings, we performed additional within-group and within-subject comparisons. To this end, each group of animals was either tested with a new statistic (e.g., all 7 rats originally trained/tested with 3-point correlations were tested with 4-point) or was split into two subgroups, each tested with a different statistic (e.g., 8 of the 11 rats originally trained/tested with 4-point correlations were tested with 3-point, while the remaining 3 animals were tested with 2-point). Results of these additional experiments are reported in Fig. 3, where the average sensitivity of each group when tested on the statistic they learned first (i.e., same data as in Fig. 2c) is shown along with the sensitivity of the same group (or two subgroups) when tested with the new statistic(s) (error bars are SEM). Rats trained on 1- and 2-point statistics (to which animals were found to be more sensitive; see Fig. 2c) performed poorly when asked to generalize to higher-order correlations (compare green and purple dots with red and blue dots on the left half of Fig. 3), while animals trained on the 4-point statistic performed on 2-point correlations as well as rats that were originally trained on those textures (compare blue dots). This shows that the better discriminability of textures containing lower-order correlations is a very robust phenomenon, which is independent of the history of training and is observable within individual subjects. Moreover, now comparing 3- and 4-

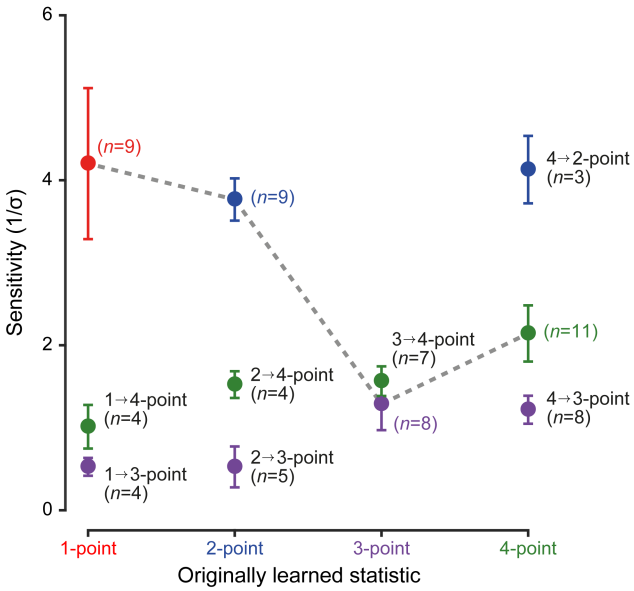

Fig. 3 | Rat sensitivity to multipoint correlations – dependence on training history and within-subject analysis. Colored dots connected by the dashed line show the sensitivities of the four groups of rats to the statistics (indicated in abscissa) they were originally trained on (i.e., same data as the colored diamonds in Fig. 2c). The other dots show the sensitivities obtained when subgroups of rats originally trained on a given statistics (indicated in abscissa) were tested with different statistics. Specifically: 1) out of the nine rats originally trained/tested with 1-point correlations (red dot), four were tested with 3-point (purple dot) and four with 4-point (green dot) correlations (1 rat did not reach this test phase); 2) out of the nine rats originally trained/tested with 2-point correlations (blue dot), five were tested with 3-point (purple dot) and four with 4-point (green dot) correlations; 3) out of the eight rats originally trained/tested with 3-point correlations (purple dot), seven were tested with 4-point (green dot) correlations (1 rat did not reach this test phase); and 4) out of the eleven rats originally trained/tested with 4-point correlations (green dot), eight were tested with 3-point (purple dot) and three with 2-point correlations (blue dot).

points correlations, generalization to 4-point discrimination always led to higher performance than generalization to 3-point (green dots are all above purple dots), independently of the statistic on which the animal was initially trained. In particular, the 7 rats switching from the 3-point to the 4-point statistic showed slightly but significantly higher sensitivity for the latter, despite the extensive training they had received with the former ($p < 0.05$, paired one-tailed t-test). Vice versa, the 8 rats switching from the 4-point to the 3-point statistic displayed significantly higher sensitivity for the former ($p < 0.01$, paired one-tailed t-test). This confirms that the larger discriminability of the 4-point statistic, as compared to the 3-point one, is a small but statistically robust phenomenon within individual subjects.

Overall, our results (Fig. 2 and 3) show that rat sensitivity to multipoint statistics is similar to that observed in humans and matches their variability across natural images[12,13]. As such, our findings are the first demonstration that a rodent species and humans are similarly adapted to process the statistical structure of visual textures. This attests to the fundamental role of the statistics of the natural world in shaping visual processing across species and opens a path towards a causal test of the efficient coding principle through altered-rearing experiments[14–16].

Author contributions. Conceptualization: RC, EP, VB, DZ. Data curation: RC. Formal analysis: RC, EP. Funding acquisition: VB, DZ. Investigation: RC, AC, AB. Methodology: RC, EP, DZ. Project administration: DZ. Resources: DZ. Software: EP. Supervision: VB, DZ. Visualization: RC. Writing – original draft: RC. Writing – review & editing: RC EP VB DZ.

Acknowledgements. We acknowledge the financial support of the European Research Council Consolidator Grant project no. 616803-LEARN2SEE (DZ), the National Science Foundation grant 1734030 (VB), the National Institutes of Health grant R01NS113241(EP) and the Computational Neuroscience Initiative of the University of Pennsylvania (VB).

References

1. Sterling, P. & Laughlin, S. *Principles of Neural Design*. (The MIT Press, 2015).
2. Atick, J. J. & Redlich, A. N. Towards a Theory of Early Visual Processing. *Neural Comput.* **2**, 308–320 (1990).
3. Olshausen, B. A. & Field, D. J. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**, 607–9 (1996).
4. Pitkow, X. & Meister, M. Decorrelation and efficient coding by retinal ganglion cells. *Nat. Neurosci.* **15**, 628–635 (2012).
5. Fairhall, A. L., Lewen, G. D., Bialek, W. & de Ruyter van Steveninck, R. R. Efficiency and ambiguity in an adaptive neural code. *Nature* **412**, 787–792 (2001).
6. Laughlin, S. A Simple Coding Procedure Enhances a Neuron's Information Capacity. *Z. Für Naturforschung C* **36**, 910–912 (1981).
7. Smith, E. C. & Lewicki, M. S. Efficient auditory coding. *Nature* **439**, 978–982 (2006).
8. Carlson, N. L., Ming, V. L. & DeWeese, M. R. Sparse Codes for Speech Predict Spectrotemporal Receptive Fields in the Inferior Colliculus. *PLOS Comput. Biol.* **8**, e1002594 (2012).
9. Teşileanu, T., Cocco, S., Monasson, R. & Balasubramanian, V. Adaptation of olfactory receptor abundances for efficient coding. *eLife* **8**, e39279 (2019).
10. Młynarski, W. & Hermundstad, A. M. Efficient and adaptive sensory codes. *bioRxiv* 669200 (2020) doi:10.1101/669200.
11. Schwartz, O. & Simoncelli, E. P. Natural signal statistics and sensory gain control. *Nat Neurosci* **4**, 819–25 (2001).
12. Hermundstad, A. M. *et al.* Variance predicts salience in central sensory processing. *eLife* **3**, e03722 (2014).
13. Tkačik, G., Prentice, J. S., Victor, J. D. & Balasubramanian, V. Local statistics in natural scenes predict the saliency of synthetic textures. *Proc. Natl. Acad. Sci.* **107**, 18149–18154 (2010).
14. Hunt, J. J., Dayan, P. & Goodhill, G. J. Sparse Coding Can Predict Primary Visual Cortex Receptive Field Changes Induced by Abnormal Visual Input. *PLOS Comput. Biol.* **9**, e1003005 (2013).
15. Matteucci, G. & Zoccolan, D. Unsupervised experience with temporal continuity of the visual environment is causally involved in the development of V1 complex cells. *Sci. Adv.* **6**, eaba3742 (2020).
16. White, L. E. & Fitzpatrick, D. Vision and Cortical Map Development. *Neuron* **56**, 327–338 (2007).
17. Piasini, E. metex – Maximum Entropy TEXtures Version 1.0.0. *https://pypi.org/project/metex* (2021) doi:https://doi.org/10.5281/zenodo.4753198.

18. Victor, J. D. & Conte, M. M. Local image statistics: maximum-entropy constructions and perceptual salience. *JOSA A* **29**, 1313–1345 (2012).
19. Fleming, S. M., Maloney, L. T. & Daw, N. D. The Irrationality of Categorical Perception. *J. Neurosci.* **33**, 19060–19070 (2013).
20. Geisler, W. S. Contributions of ideal observer theory to vision research. *Vision Res.* **51**, 771–781 (2011).

Methods

M.1 Psychophysics experiments

Subjects. A total of 42 male adult Long Evans rats (Charles River Laboratories) were tested in a visual texture discrimination task. Animals started the training at 10 weeks, after 1 week of quarantine upon arrival in our institute and 2 weeks of handling to familiarize them with the experimenters. Their weight at arrival was approximately 300g and they grew to over 600g over the time span of the experiment. Rats always had free access to food but their access to water was restricted in the days of the behavioral training (5 days a week). They received 10–20ml of diluted pear juice (1:4) during the execution of the discrimination task, after which they were also given free access to water for the time needed to reach at least the recommended 50ml/Kg intake per day.
The number of rats was chosen in such a way to yield meaningful statistical analyses (i.e., to have about 10 subjects for each of the texture statistic tested in our study), under the capacity constraint of our behavioral rig. The rig allows to simultaneously test 6 rats, during the course of 1-1.5 hours[21,22]. Given the need for testing four different texture statistics, we started with a first batch of 24 animals (i.e., 6 per statistics), which required about 6 h of training per day. This first batch was complemented with a second one of 18 more rats, again divided among the four statistics (see below for details), so as to reach the planned number of about 10 animals per texture type. The first batch arrived in November 2018 and was tested throughout most of 2019; the second group arrived in September 2019 and was tested throughout most of 2020. In the first batch, 4 animals did not reach the test phase (i.e., the phase yielding the data shown in Fig. 2), because three of them did not achieve the criterion performance during the initial training phase (see below) and one died shortly after the beginning of the study. In the second batch one rat died before reaching the test phase and two more died before the last test phase with switched statistics (i.e., the phase yielding the data of Fig. 3).
All animal procedures were conducted in accordance with the international and institutional standards for the care and use of animals in research and were approved by the Italian Ministry of Health and after consulting with a veterinarian (Project DGSAF 25271, submitted on December 1, 2014 and approved on September 4, 2015, approval 940/2015-PR).

Experimental setup. Rats were trained in a behavioral rig consisting of two racks, each equipped with three operant boxes (a picture of the rig and a schematic of the operant box can be found in previous studies[21,22]). Each box was equipped with a 21.5” LCD monitor (ASUS VEZZHR) for the presentation of the visual stimuli and an array of three stainless-steel feeding needles (Cadence Science), serving as response ports. To this end, each needle was connected to a led-photodiode pair to detect when the nose of the animal approached and touched it (a Phidgets 1203 input/output device was used to collect the signals of the photodiodes). The two lateral feeding needles were also connected to computer-controlled syringe pumps (New Era Pump System NE-500) for delivery of the liquid reward. In each box, one of the walls bore a 4.5 cm-diameter viewing hole, so that a rat could extend its head outside the box, face the stimulus display (located at 30 cm from the hole) and reach the array with the response ports.

Visual stimuli. Maximum-entropy textures were generated using the methods described by Victor and Conte[18]. To this end, we implemented a standalone library and software package that we have since made publicly available as free software[17]. In the experiment, we used white noise

textures as well as textures with positive levels of four different multipoint statistics (Fig. 1a). In the nomenclature established by Victor and Conte[18], these are γ, β, θ and α, corresponding to, respectively, 1-, 2-, 3- and 4-point correlations. It should be noted that, with the exception of the extreme value of the γ statistic ($\gamma = 1$ corresponds to a fully white image), the intensity level of a given statistic does not specify deterministically the resulting texture image. In our experiment, for any intensity level of each statistic, multiple, random instances of the textures were built to be presented to the rats during the discrimination task (see below for more details).
Subjects had to discriminate between visual textures containing one of the 4 selected statistics and white noise. Each texture had a size of 39x22 pixels and occupied the entire monitor (full-field stimuli). The pixels had a dimension of about 2 degrees of visual angle. Given that the maximal resolution of rat vision is about 1 cycle per degree[23,24,25], such a choice of the pixel size guaranteed that the animals could discriminate between neighboring pixels of different color. Textures were showed at full-contrast over the LCD monitors that were calibrated in such a way to have minimal luminance of 0.126±0.004 cd/mm[22](average ± SD across the six monitors), mean luminance of 42±5 cd/mm[22], maximal luminance of 129±5 cd/mm[22]and an approximately linear luminance response curve.

Discrimination task. As shown in Fig. 1b, each rat was trained to: 1) touch the central response port to trigger stimulus presentation and initiate a behavioral trial; and 2) touch one of the lateral response ports to report the identity of the visual stimulus and collect the reward (all the animals were trained with the following stimulus/response association: structured texture $\rightarrow$ left response port; white noise texture $\rightarrow$ right response port). The stimulus remained on the display until the animal responded or for a maximum of 5 seconds, after which the trial was considered as ignored. In case of a correct response the stimulus was removed, a positive reinforcement sound was played and a white (first animal batch) or gray (second batch) background was shown during delivery of the reward. In case of an incorrect choice, the stimulus was removed and a 1-3s time-out period started, during which the screen flickered from middle-gray to black at a rate of 10 Hz, while a "failure" sound was played. During this period the rat was not allowed to initiate a new trial. To prevent the rats from making impulsive random choices, trials where the animals responded in less than 300 or 400 ms were considered as aborted: the stimulus was immediately removed and a brief sound was played. In each trial, the visual stimuli had the same probability (50%) of being sampled from the pool of white noise textures or from the pool of structured textures, with the constraint that stimuli belonging to the same category were shown for at most n consecutive trials (with n varying between 2 and 3 depending on the animal and on the session), so as to prevent the rats from developing a bias towards one of the response ports.
Stimulus presentation, response collection and reward delivery were controlled via workstations running the open source suite MWorks (https://mworks.github.io).

Experimental design. Each rat was assigned to a specific statistic, from 1- to 4-point, for which it was trained in phases I and II and then tested in phase III. Generalization to a different statistic from the one the rat was trained on was assessed in phase IV. Out of the 42 rats, 9 were trained with 1-point statistics, 9 with 2-point, 12 with 3-point and 12 with 4-point. The animals that reached phase III were 9, 9, 8 and 11, respectively, for the four statistics.

Phase I. Initially, rats were trained to discriminate unstructured textures made of white noise from structured textures containing a single high-intensity level of one of the statistics (for 1-point

and 2-point: 0.85; for 3-point and 4-point: 0.95). To make sure that the animals learned a general distinction between structured and unstructured textures (and not between specific instances of the two stimulus categories), in each trial both kinds of stimuli were randomly sampled (without replacement) from a pool of 350 different textures. Since the rats typically performed between 200 and 300 trials in a training session, every single texture was not shown more than once. A different pool of textures was used in each of the five days within a week of training. The same five texture pools were then used again (in the same order) the following week. Therefore, at least 7 days had to pass before a given texture stimulus was presented again to a rat.
For the first batch of rats, we moved to the second phase of the experiment all the animals that were able to reach at least an average performance of 65% correct choices over a set of 500 trials (collected across a variable number of consecutive sessions). Based on this criterion, 2 rats tested with 3-point textures and 1 rat tested with 4-point textures were excluded from further testing.
For the second batch of rats, we decided to admit all the animals to the following experimental phases after a prolonged period of training in the first phase. In fact, we reasoned that, in case some texture statistic was particularly hard to discriminate, imposing a criterion performance in the first phase of the experiment would bias the pool of rats tested with such very difficult statistic towards including only exceptionally proficient animals. This in turn, could lead to an overestimation of rat typical sensitivity to such difficult statistic. On the other hand, the failure of a rat to reach a given criterion performance could be due to intrinsic limitations of its visual apparatus (such as a malfunctioning retina or particularly low acuity). Therefore, to make sure that our result did not depend on including in our analysis some animals of the second batch that did not reach 65% correct discrimination in the first training phase, the perceptual sensitivities were re-estimated after excluding those rats (i.e., after excluding 1 rat from the 2-point, 3 rats from the 3-point and 1 from the 4-point groups). As shown in Supplementary Fig. 1, the resulting sensitivity ranking was unchanged (compare to Fig. 1c) and all pairwise comparisons remained statistically significant (two sample t-test with Holm-Bonferroni correction).

Phase II. In this phase, we introduced progressively lower levels of intensity of each statistic, bringing them gradually closer to the zero-intensity level corresponding to white noise. To this end, we applied an adaptive staircase procedure to update the minimum level of the statistic to be presented to a rat based on its current performance. Briefly, in any given trial, the level of the multipoint correlation in the structured textures was randomly sampled between a minimum level (under the control of the staircase procedure) and a maximum level (fixed at the value used in phase I). Within this range, the sampling was not uniform, but was carried out using a geometric distribution (with the peak at the minimum level), so as to make much more likely for rats to be presented with intensity levels at or close to the minimum. The performance achieved by the rats on the current minimum intensity level was computed every ten trials. If such a performance was higher than 70% correct, the minimum intensity level was decreased by a step of 0.05. By contrast, if the performance was lower than 50%, the minimum intensity level was increased of the same amount.
This procedure allowed the rats to learn to discriminate progressively lower levels of the statistic in a gradual and controlled way. At the end of this phase, the minimum intensity level reached by the animal in the three groups was: 0.21 ± 0.12, 0.2 ± 0.2, 0.70 ± 0.22, and 0.56 ± 0.18 (group average ± SD) for, respectively, 1-, 2-, 3- and 4-point correlations.

Phase III. After the training received in phases I and II, the rats were finally moved to the main test phase, where we measured their sensitivity to the multipoint correlations they were trained on. In each trial of this phase, the stimulus was either white noise or a patterned texture with equal probability. If it was a patterned texture, the level of the statistic was randomly selected from the set $\{0.02, 0.09, 0.16, \ldots, 0.93, 1\}$ (i.e., from 0.02 to 1 in steps of 0.07) with uniform probability. The responses of each rat over this range of intensity levels yielded psychometric curves (see example in Fig. 2a), from which rat sensitivity was measured by fitting the Bayesian ideal observer model described below (Fig. 2b and c).

Phase IV. To verify the sensitivity ranking observed in phase III, we carried out an additional test phase, where each rat was tested on a new statistic, which was different from the one the animal was previously trained and tested on. The two groups of rats that were originally trained with the statistics yielding the highest sensitivity in phase III (i.e., 1- and 2-point correlations; see Fig. 2c) were split in approximately equally-sized subgroups and each of these subgroups was tested with the less discriminable statistics (i.e., 3- and 4-point correlations; leftmost half of Fig. 3). This allowed assessing that, regardless of the training history, sensitivity to 4-point correlations was slightly but consistently higher than sensitivity to 3-point correlations. For the group of rats originally tested with the 3-point statistic, all the animals were switched to the 4-point (third set of points in Fig. 3). This allowed comparing the sensitivities to these statistics at the within-subject level (notably, these rats were found to be significantly more sensitive to the 4-point textures than to the 3-point, despite the extensive training they had received with the latter). For the same reason, most of the rats (8/11) of the last group (i.e., the animals originally trained/tested with the 4-point correlations; last set of points in Fig. 3) were switched to the 3-point statistic, which yielded again the lowest discriminability. A few animals (3/11) were instead tested with the 2-point statistic, thus verifying that the latter was much more discriminable than the 4-point one (again, despite the extensive training the animals of this group had received with the 4-point textures).

Data availability. Experimental data will be made freely and permanently available at [26] upon completion of peer review.

M.2 Ideal observer model

In this section we describe the ideal observer model we used to estimate the sensitivity of the rats to the different textures. The approach is a standard one and is inspired by that in [19]. Because our intention is to use an ideal observer as a model for animal behavior, we will write interchangeably "rat", "animal"and "ideal observer" in the following.

M.2.1 Preliminaries

The texture discrimination task is a two-alternative forced choice (2AFC) task, where the stimulus can be either a sample of white noise or a sample of textured noise, and the goal of the animal is to correctly report the identity of each stimulus. On any given trial, either stimulus class can happen with equal probability. The texture class is composed of K discrete, positive values of the texture. In practice, $K = 14$, and these values are $\{0.02, 0.09, \ldots, 0.93, 1\}$, but we'll use a generic K in the derivations for clarity. The texture statistics are parametrised such that a statistic value of

zero corresponds to white noise. Therefore, if we call s the true level of the statistic, the task is a parametric discrimination task where the animal has to distinguish $s = 0$ from $s > 0$.

M.2.2 Key assumptions

1. each trial is independent from those preceding and following it (both for the generated texture and for the animal's behavior);

2. on any given trial, the nominal (true) value of the statistic is some value s. Because the texture has finite size, the empirical value of the statistic in the texture will be somewhat different from s. We lump this uncertainty together with that induced by the animal's perceptual process, and we say that any given trial results on the production of a *percept* x, sampled from a truncated Normal distribution centered around the nominal value of the statistic and bounded between $a = -1$ and $b = 1$:

$$p(x|s,\sigma,a,b) = \frac{1}{\sigma}\frac{\phi\left(\frac{x-s}{\sigma}\right)}{\Phi\left(\frac{b-s}{\sigma}\right) - \Phi\left(\frac{a-s}{\sigma}\right)}$$

   where $\phi(\cdot)$ is the probability density function of the standard Normal and $\Phi(\cdot)$ is its cumulative density function. Setting the bounds to -1 and 1 allows us to account for the fact that the value of a statistic is constrained within this range by construction. We will keep a and b in some of the expressions below for generality and clarity, and we will substitute their values only at the end.

3. we assume that each rat has a certain prior over the statistic level that we parametrise by the log prior odds:

$$\alpha := \ln \frac{p(s=0)}{p(s>0)}$$

   Where α depends on the rat. More specifically, we assume that each rat assigns a prior probability $p(s = 0) = 1/(1 + e^{-\alpha})$ to the presentation of a noise sample, and a probability of $1/[K(1 + e^{\alpha})]$ to the presentation of a texture coming from any of the K nonzero statistic values. In formulae:

$$p(s) = \frac{\delta_{s,0}}{1+e^{-\alpha}} + \frac{1}{K}\sum_{k=1}^{K}\frac{\delta_{s,s_k}}{1+e^{\alpha}}$$

   where δ is Kronecker's delta, and $s_k > 0, k \in \{1,\ldots,K\}$ are the K possible nonzero values of the statistic. Note that this choice of prior matches the distribution actually used in generating the data for the experiment, except that α is a free parameter instead of being fixed at 0.

4. we assume that the true values of α, σ, a and b are accessible to the decision making process of the rat.

M.2.3 Derivation of the ideal observer

For a particular percept, the ideal observer will evaluate the posterior probability of noise vs texture given that percept. It will report "noise" if the posterior of noise is higher than the posterior of texture, and "texture" otherwise.

More in detail, for a given percept x we can define a decision variable D as the log posterior ratio:

$$D(x) := \ln \frac{p(s=0|x)}{p(s>0|x)} = \ln \frac{p(x|s=0)}{p(x|s>0)} + \ln \frac{p(s=0)}{p(s>0)} \tag{1}$$

With this definition, the rat will report "noise" when $D > 0$ and "texture" otherwise.
By plugging in the likelihood functions and our choice of prior, we get

$$D(x) = \alpha + \ln \left[\frac{1}{\sigma} \frac{\phi(x/\sigma)}{\Phi\left(\frac{b}{\sigma}\right) - \Phi\left(\frac{a}{\sigma}\right)} \right] - \ln \left[\frac{1}{K} \sum_k \frac{1}{\sigma} \frac{\phi\left(\frac{x-s_k}{\sigma}\right)}{\Phi\left(\frac{b-s_k}{\sigma}\right) - \Phi\left(\frac{a-s_k}{\sigma}\right)} \right] \tag{2}$$

Now, remember that *given a value of the percept* x, the decision rule based on D is fully deterministic (maximum a posteriori estimate). But on any given trial we don't know the value of the percept — we only know the nominal value of the statistic. On the other hand, our assumptions above specify the distribution $p(x|s)$ for any s, so the deterministic mapping $D(x)$ means that we can compute the probability of reporting "noise" as

$$p(\text{report noise}|s) = p(D > 0|s) = \int_{x:D(x)>0} p(x|s)\,\mathrm{d}x \tag{3}$$

We note at this point that $D(x)$ is monotonic: indeed,

$$\begin{aligned}
\frac{\mathrm{d}D(x)}{\mathrm{d}x} &= -\frac{x}{\sigma^2} + \\
&- \left[\frac{1}{K} \sum_k \frac{\exp\left[-\frac{(x-s_k)^2}{2\sigma^2}\right]}{\Phi\left(\frac{b-s_k}{\sigma}\right) - \Phi\left(\frac{a-s_k}{\sigma}\right)} \right]^{-1} \frac{1}{K} \sum_k \frac{\exp\left[-\frac{(x-s_k)^2}{2\sigma^2}\right] \left(-\frac{x-s_k}{\sigma^2}\right)}{\Phi\left(\frac{b-s_k}{\sigma}\right) - \Phi\left(\frac{a-s_k}{\sigma}\right)} \\
&= - \left[\frac{1}{K} \sum_k \frac{\exp\left[-\frac{(x-s_k)^2}{2\sigma^2}\right]}{\Phi\left(\frac{b-s_k}{\sigma}\right) - \Phi\left(\frac{a-s_k}{\sigma}\right)} \right]^{-1} \frac{1}{K} \sum_k \frac{\exp\left[-\frac{(x-s_k)^2}{2\sigma^2}\right] \frac{s_k}{\sigma^2}}{\Phi\left(\frac{b-s_k}{\sigma}\right) - \Phi\left(\frac{a-s_k}{\sigma}\right)} \\
&< 0 \text{ for all } x
\end{aligned} \tag{4}$$

where for the last inequality we have used the fact that $a < b$ and therefore $\Phi((b-s_k)/\sigma) > \Phi((a-s_k)/\sigma)$. This result matches the intuitive expectation that a change in percept in the positive direction (i.e., away from zero) should always make it less likely for the observer to report "noise".
Because $D(x)$ is monotonic, there will be a unique value of x such that $D(x) = 0$, and the integration region $x : D(x) > 0$ will simply consist of all values of x smaller than that. More formally, if we define

$$x^* = x^*(\alpha, \sigma) \text{ such that } D(x^*) = 0 \tag{5}$$

we can write

$$
\begin{aligned}
p(\text{report noise}|s) &= \int_a^{x^*} \frac{1}{\sigma} \frac{\phi\left(\frac{x-s}{\sigma}\right)}{\Phi\left(\frac{b-s}{\sigma}\right) - \Phi\left(\frac{a-s}{\sigma}\right)} \mathrm{d}x \\
&= \frac{\Phi\left(\frac{x^*(\alpha,\sigma)-s}{\sigma}\right) - \Phi\left(\frac{a-s}{\sigma}\right)}{\Phi\left(\frac{b-s}{\sigma}\right) - \Phi\left(\frac{a-s}{\sigma}\right)} \\
&= \frac{\Phi\left(\frac{x^*(\alpha,\sigma)-s}{\sigma}\right) - \Phi\left(\frac{-1-s}{\sigma}\right)}{\Phi\left(\frac{1-s}{\sigma}\right) - \Phi\left(\frac{-1-s}{\sigma}\right)}
\end{aligned}
\tag{6}
$$

where in the last passage we have substituted $a = -1$ and $b = 1$.

M.2.4 Example: single-level discrimination case

To give an intuitive interpetation of the results above, consider the case where $K = 1$, so the possible values of the statistic are only two, namely 0 and s_1. In this case,

$$
\begin{aligned}
D^{(1)} &= \alpha + \ln \frac{\Phi\left(\frac{b-s_1}{\sigma}\right) - \Phi\left(\frac{a-s_1}{\sigma}\right)}{\Phi\left(\frac{b}{\sigma}\right) - \Phi\left(\frac{a}{\sigma}\right)} - \frac{2xs_1 - s_1^2}{2\sigma^2} \\
&= \alpha + \beta - \frac{2xs_1 - s_1^2}{2\sigma^2}
\end{aligned}
$$

where

$$
\beta := \ln \frac{\Phi\left(\frac{b-s_1}{\sigma}\right) - \Phi\left(\frac{a-s_1}{\sigma}\right)}{\Phi\left(\frac{b}{\sigma}\right) - \Phi\left(\frac{a}{\sigma}\right)}
$$

so that we can write x^* in closed form:

$$
x^{*(1)} = D^{(1)-1}(0) = \frac{s_1}{2} + \frac{\sigma^2}{s_1}(\alpha + \beta)
$$

which can be read as saying that the decision boundary is halfway between 0 and s_1, plus a term that depends on the prior bias and the effect of the boundaries of the domain of x (but involves the sensitivity too, represented by σ).
Simplifying things even further, if we remove the domain boundaries (by setting $a \to -\infty$ and $b \to +\infty$), we have that $\beta \to 0$. In this case, by plugging the expression above in Equation 6 we obtain

$$
p(\text{report noise}|s) = \Phi\left[\frac{\sigma}{s_1}\alpha - \frac{(s - s_1/2)}{\sigma}\right] \tag{7}
$$

and therefore we recover a simple cumulative Normal form for the psychometric function. By looking at Equation 7 it is clear how the prior bias α introduces a horizontal shift in the psychometric curve, and σ controls the slope (but also affects the horizontal location when $\alpha \neq 0$).

M.3 Fitting the ideal observer model to the experimental data

Independently for each rat, we infer a value of α and σ by maximising the likelihood of the data under the model above. More in detail, for a given rat and a given statistic value s (including 0), we call N_s the number of times the rat reported "noise", and T_s the total number of trials. For a given fixed value of α and σ, under the ideal observer mode the likelihood of N_s will be given by a Binomial probability distribution for T_s trials and probability of success given by the probability of reporting noise in Equation 6,

$$p_s(N_s|\alpha,\sigma) = \binom{T_s}{N_s} p(\text{rep. noise}|s,\alpha,\beta)^{N_s} \left(1 - p(\text{rep. noise}|s,\alpha,\beta)\right)^{T_s - N_s}$$

Assuming that the data for the different values of s is conditionally independent given α and σ, the total log likelihood for the data of the given rat is simply the sum of the log likelihoods for the individual values of N_s

$$\ln p(\{N_{s_k}\}_{k=1}^{K}|\alpha,\sigma) = \sum_{k=1}^{K} \ln p_{s_k}(N_k|\alpha,\sigma)$$

We find numerically the values of α and σ that maximise this likelihood, using Matlab's `mle` function with initial condition $\alpha = 0.1$, $\sigma = 0.4$. Note that evaluating the likelihood for any given value of α and σ requires finding x^*, defined as the zero of Equation 2. We do this numerically by using Matlab's `fzero` function with initial condition $x = 0$.

References

21. Zoccolan, D. Invariant visual object recognition and shape processing in rats. *Behavioural Brain Besearch* 285, 10–33 (2015).
22. Djurdjevic, V., Ansuini, A., Bertolini, D., Macke, J. H. & Zoccolan, D. Accuracy of Rats in Discriminating Visual Objects Is Explained by the Complexity of Their Perceptual Strategy. *Current biology* 28, 1005–1015.e5 (2018).
23. Keller, J., Strasburger, H., Cerutti, D. T. & Sabel, B. A. Assessing spatial vision — automated measurement of the contrast-sensitivity function in the hooded rat. *Journal of Neuroscience Methods* 97, 103–110 (2000).
24. Prusky, G. T., West, P. W. & Douglas, R. M. Behavioral assessment of visual acuity in mice and rats. *Vision Research* 40, 2201-2209 (2000).
25. Prusky, G. T., Harker, K. T., Douglas, R. M. & Whishaw, I. Q. Variation in visual acuity within pigmented, and between pigmented and albino rat strains. *Behavioural Brain Research* 136, 339–348 (2002).
26. Caramellino, R. et al. Data from "Rat sensitivity to multipoint statistics is predicted by efficient coding of natural scenes" 2021. doi:10.5281/zenodo.4762567.