HHS Public Access

Author manuscript

Cell. Author manuscript; available in PMC 2020 May 02.

Published in final edited form as:

Cell. 2019 May 02; 177(4): 999–1009.e10. doi:10.1016/j.cell.2019.04.005.

Evolving images for visual neurons using a deep generative network reveals coding principles and neuronal preferences

Carlos R. Ponce[#1,4,*], **Will Xiao**[#2,*], **Peter F. Schade**[#1,*], **Till S. Hartmann**[1], **Gabriel Kreiman**[3], **Margaret S. Livingstone**[1,5]

[1]Department of Neurobiology, Harvard Medical School, Boston, MA 02115, USA.

[2]Department of Molecular and Cellular Biology, Harvard University, Cambridge, MA 02138, USA.

[3]Department of Ophthalmology, Boston Children's Hospital, Harvard Medical School, Boston, MA 02115, USA.

[4]Department of Neuroscience, Washington University School of Medicine, St. Louis MO 63110, USA.

[5]Lead Contact

[#] These authors contributed equally to this work.

SUMMARY

What specific features should visual neurons encode, given the infinity of real-world images and the limited number of neurons available to represent them? We investigated neuronal selectivity in monkey inferotemporal cortex using the vast hypothesis space of a generative deep neural network, avoiding assumptions about features or semantic categories. A genetic algorithm searched this space for stimuli that maximized neuronal firing. This led to the evolution of rich synthetic images of objects with complex combinations of shapes, colors, and textures, sometimes resembling animals or familiar people, other times revealing novel patterns that did not map to any clear semantic category. These results expand our conception of the dictionary of features encoded in the cortex, and the approach can potentially reveal the internal representations of any system whose input can be captured by a generative model.

In Brief

Neurons guided the evolution of their own best stimuli with a generative deep neural network.

Correspondence: Carlos Ponce: crponce@wustl.edu; Margaret Livingstone: mlivingstone@hms.harvard.edu.

[*]Equal contribution

AUTHOR CONTRIBUTIONS

GK and WX conceived of the approach. WX implemented the network algorithm and did all the *in silico* experiments. CP & PS adapted the algorithm to neurophysiology. PS, ML, TH & CP collected and analyzed data. CP, WX & ML wrote the manuscript, all authors revised the manuscript, and ML acquired funding.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Graphical Abstract

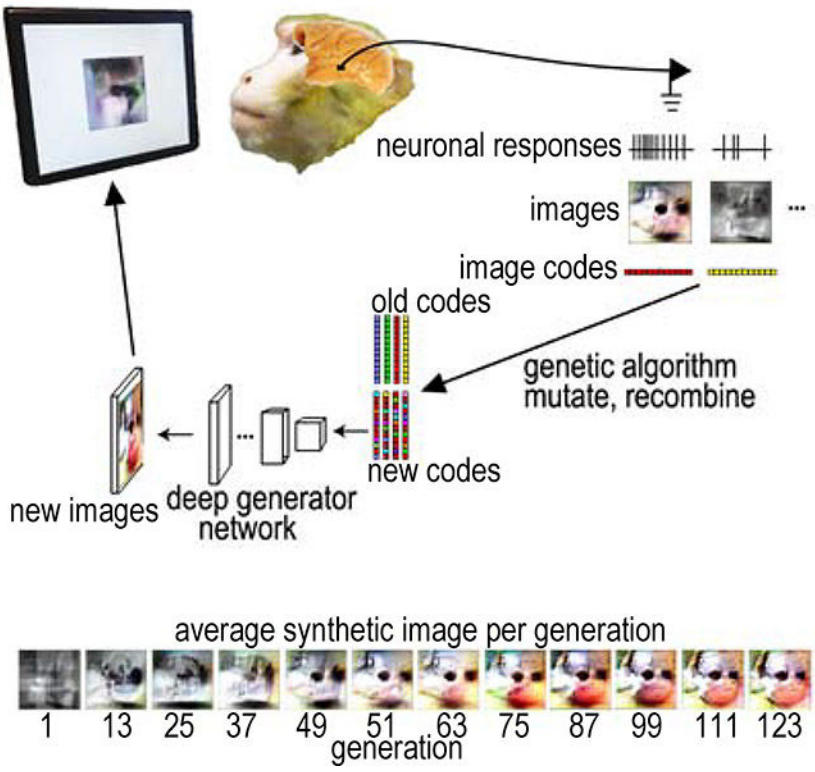

Keywords

neural networks; inferotemporal cortex; Generative Adversarial Network

INTRODUCTION

A transformative revelation in neuroscience was the realization that visual neurons respond preferentially to some stimuli over others (Hubel and Wiesel, 1962). Those findings opened the doors to investigating neural coding for myriad stimulus attributes. A central challenge in elucidating neuronal tuning in visual cortex is the impossibility of testing all stimuli. Even for a small patch of 100 × 100 pixels, there are $\sim 10^{3,010}$ possible binary images, $\sim 10^{24082}$ grayscale images, or $\sim 10^{72247}$ 8-bit color images. Using natural images reduces the problem, but it is still impossible to present a neuron with all possible natural stimuli. Investigators circumvent this formidable empirical challenge by using *ad hoc* hand-picked stimuli, inspired by hypotheses that particular cortical areas encode specific visual features (Felleman and Van Essen, 1987; Zeki, 1973, 1974). This approach has led to important insights through the discovery of cortical neurons that respond to stimuli with specific motion directions (Hubel, 1959), color (Michael, 1978), binocular disparity (Barlow et al., 1967), curvature (Pasupathy and Connor, 1999), and even complex natural shapes such as hands or faces (Desimone et al., 1984; Gross et al., 1972).

Despite the successes using hand-picked stimuli, the field may have missed stimulus properties that better reflect the "true" tuning of cortical neurons. A series of interesting alternative approaches have addressed this question. One approach is to start with hand-picked stimuli that elicit strong activation and systematically deform those stimuli (Chang and Tsao, 2017; Freiwald et al., 2009; Kobatake and Tanaka, 1994); this approach has revealed that neurons often tend to respond even better to distorted versions of the original stimuli (Freiwald et al., 2009; Leopold et al., 2006). Another is spike-triggered averaging of noise stimuli (Gaska et al., 1994; Jones and Palmer, 1987), but this has not yielded useful results in higher cortical areas, because it cannot capture non-linearities. An elegant alternative is to use a genetic algorithm whereby the neuron under study can itself guide its own stimulus selection. Connor and colleagues (Carlson et al., 2011; Yamane et al., 2008) pioneered this approach to study selectivity in macaque V4 and IT. Our method extends and complements this approach in order to investigate the tuning properties of inferior temporal cortex (IT) neurons in macaque monkeys.

Here we use a novel combination of a pre-trained deep generative neural network (Dosovitskiy and Brox, 2016) and a genetic algorithm to allow neuronal responses to guide the evolution of synthetic images. By training on more than one million images from ImageNet (Russakovsky et al., 2015), the generative adversarial network learns to model the statistics of natural images without merely memorizing the training set (Dosovitskiy and Brox, 2016); Fig. S1, thus representing a vast and general image space, constrained only by natural image statistics. We reasoned that this would be an efficient space in which to perform the genetic algorithm, because the brain also learns from real-world images, so its preferred images are also likely to follow natural image statistics. Moreover, convolutional neural networks emulate aspects of computation along the primate ventral visual stream (Yamins et al., 2014), and this particular generative network has been used to synthesize images that strongly activate units in several convolutional neural networks, including ones not trained on ImageNet (Nguyen et al., 2016). The network takes 4096-dimensional vectors (image codes) as input and deterministically transforms them into 256×256 RGB images (STAR METHODS, Figure 1). In combination, a genetic algorithm used responses of neurons recorded in alert macaques to optimize image codes input to this network. Specifically, each experiment started from an initial population of 40 images created from random achromatic textures (Portilla and Simoncelli, 2000) (Figure 1B). We recorded responses of IT neurons (spike counts 70–200 ms after stimulus onset minus background) while monkeys engaged in a passive fixation task. Images subtended 3°×3° and covered the unit's receptive field (Figure 1C). Neuronal responses to each synthetic image were used to score the image codes. In each generation, images were generated from the top 10 image codes from the previous generation, unchanged, plus 30 new image codes generated by mutation and recombination of all the codes from the preceding generation selected based on firing rate (Figure 1D). This process was repeated for up to 250 generations over 1–3 hours; session duration depended on the monkey's willingness to maintain fixation. To monitor changes in firing rate due to adaptation and to compare synthetic-image responses to natural-image responses, we interleaved reference images that included faces, body parts, places and simple line drawings.

We term the overall approach XDREAM (E**X**tending **D**eepDream with **R**eal-time **E**volution for **A**ctivity **M**aximization in real neurons; Figure 1D). We conducted evolution experiments on IT neurons in six monkeys: two with chronic microelectrode arrays in posterior IT (PIT, monkeys Ri and Gu), two with chronic arrays in central IT (CIT, monkeys Jo and Y1), one monkey, Ge, with chronic arrays in both CIT and PIT, and one with a recording chamber over CIT (monkey B3). Lastly we validated the approach in a seventh monkey with a chronic array in primary visual cortex (V1, monkey Vi).

RESULTS

Evolution of preferred stimuli for units in CaffeNet

We first validated XDREAM on units in an artificial neural network, as models of biological neurons. Our method generated super stimuli for units across layers in CaffeNet, a variant of AlexNet (Figure 2) (Krizhevsky et al., 2012). The evolved images were frequently better stimuli than all of >1.4 million images, including the training set for CaffeNet, for all 4 layers that we tested (Figure 2C). For units in the first and last layers, the method produced stimuli that matched the ground-truth best stimuli in the first layer and category labels in the last layer. XDREAM is also able to recover the preferred stimuli of units constructed to have a single preferred image (Fig. S1). Importantly, well-known methods for feature visualization in artificial neural networks, such as DeepDream, rely on knowledge of network weights (Erhan et al., 2009; Mordvintsev et al., 2015; Nguyen et al., 2016), whereas our approach does not, making it uniquely applicable to neuronal recordings.

Evolution of preferred stimuli by one biological neuron

We first show an example of an evolution experiment for one PIT single unit (Ri-10) in chronic-array monkey Ri. The synthetic images changed with each generation as the genetic algorithm optimized the images according to the neuron's responses (Figure 3; Video S1). At the beginning of the experiment, this unit responded more strongly to the reference images than to the synthetic images, but over generations, the synthetic images evolved to become more effective stimuli (Figure 4A). To quantify the change in responses over time, we fit an exponential function to the cell's mean firing rate per generation, separately for the synthetic and for the reference images (solid thick lines in Figure 4A). This neuron showed an increase of 51.5±5.0 (95% CI) spikes/s/generation in response to the synthetic images and a decrease of −15.5±3.5 spikes/s/generation to the reference images — thus the synthetic images became gradually more effective, despite the neuron's slight reduction in firing rate to the reference images, presumably due to adaptation.

We conducted independent evolution experiments with the same single unit on different days, and all final-generation synthetic images featured a brown object against a uniform background, topped by a smaller round pink/brown region containing several small dark spots; the object was centered toward the left half of the image, consistent with the recording site being in the right hemisphere (Figure 4B). The synthetic images generated on different days were similar by eye, but not identical, potentially due to invariance of the neuron, response variability, and/or stochastic paths explored by the algorithm in the neuron's response landscape. Regardless, given that this unit was located in PIT, just anterior to the

tip of the inferior occipital sulcus and thus early in the visual hierarchy, it was impressive that it repeatedly directed the evolution of images that contained such complex motifs and that evoked such high firing rates. Two days following the evolution experiment in Figure 3 this unit was screened with 2,550 natural images, including animals, bodies, food, faces and line drawings, plus the top synthetic images from each generation. Among the natural images this neuron responded best to monkey torsos and monkey faces. Of the 10 natural images in this set giving the largest responses, five were of the head and torso of a monkey (Figure 4C). The worst natural images were inanimate or rectilinear objects (Figure 4D, E).

Evolution of preferred stimuli in other neurons

We conducted 46 independent evolution experiments on single- and multi-unit sites in IT in six different monkeys. During almost all the evolutions, the synthetic images evolved gradually to become increasingly effective stimuli. To quantify the change in stimulus effectiveness over each experiment we fit an exponential function to the mean firing rate per generation, separately for synthetic and reference images (as in Figure 3A). Synthetic-image firing rate change over the course of each experiment was on average between 25 to 84 spikes/s for the different animals (Figure 5A); synthetic image changes were significantly different from zero in 45 out of 46 individual experiments (95% CI of amplitude estimate not including zero per bootstrap test). In contrast, responses to reference images were stable or decreased slightly across generations (reference firing rate change average for different animals ranged from −11 to 9 spk/s; this change was significant in 15 out of 46 individual experiments, Figure 5A, Table S1). Thus IT neurons could consistently guide the evolution of highly effective images, despite minor adaptation. Moreover, these evolved images were often more powerful stimuli for these neurons than were the best natural images tested, despite the fact that the synthetic images were far from naturalistic. When comparing the cells' maximum responses to natural vs. evolved images in every experiment, cells showed significant differences in 25 of 46 experiments ($P < 0.03$, permutation test after false discovery correction), and, in all but one case, the synthetic images evoked the greater response (Figure 5B; see Table S4 for further quantification of natural and evolved image responses). Figure 5C shows a histogram of response magnitudes for PIT cell Ri-10 to the top synthetic image in each of the 210 generations and responses to each of the 2550 natural images (data for both synthetic and natural images collected 2 days later). Early generations are indicated by lighter gray and later by darker, so it is apparent that later generation synthetic images gave larger responses. We also illustrate one of the few experiments where natural images evoked stronger responses than did synthetic images in Figure 5D (monkey Ge, site 7), which compares the site's responses to synthetic images against 2,550 natural images. This site responded slightly better (by an average of four spikes/s, or 3.7% of its maximum rate) to images of an animal-care person who visits the animals daily, wearing our institution-specific protective mask and gown. Even in this case, one clear benefit of XDREAM is that, by coming up with effective stimuli in a manner independent from the hand-picked image set, it reveals specific features of the natural image that drove the neuron's selectivity, stripped of incidental information.

IT neurons guided the evolution of images that varied from experiment to experiment, but retained consistent features for any given recording site (see Fig. S4 for measure of

similarity across and between sites), features that bore some similarity to each neuron's preferences in natural image sets. Figure 6 shows the final-generation evolved images from two independent evolution experiments for IT sites in five monkeys, along with each site's top 10 natural images. In each case a reproducible figure emerged in the part of the synthetic image corresponding to the contralateral visual field. In three monkeys (Ri, Gu, and Ge) response profiles to natural images indicated that the arrays were located in face-preferring regions, whereas in monkey Y1, the array was in a place-preferring region, and in monkey Jo the array was in an object selective region. Face-selective unit Ri-23 evolved something face-like in the left (contralateral) half of the image; this is most apparent if one covers up the right (ipsilateral) half of the synthetic image. The images evolved by unit Ge-7 bore some resemblance to the unit's top natural image, a familiar person wearing protective clothing (see Fig. S3 for additional independent evolutions from this site). Unit Ge-15 consistently evolved a black dot just to the left of fixation on a tan background (see Fig. S3 for additional independent evolutions from this site). This unit may be similar to posterior-face-patch neurons described previously that responded optimally to a single eye in the contralateral field (Issa and DiCarlo, 2012) (see Fig. S3 for additional independent evolutions from this site). Monkey-face-selective unit Ge-17 evolved a tan area with two large black dots aligned horizontally, and a light area below (see Fig. S3 for additional independent evolutions from this site). Unit Jo-6 responded to various body parts, and evolved something not inconsistent with a mammalian body; interestingly, a whole 'body' in one evolution and a larger partial 'body' in the other. Unit Jo-5 evolved a small black square, and unit Jo-4 something black with orange below. Unit Jo-21 consistently evolved a small dark shape in the contralateral half of the image. Scene-selective unit Y1–14 evolved rectilinear shapes in the left (contralateral) field. Additional independent evolutions for some of these and other units are shown in Fig. S3.

Predicting neuronal responses to a novel image from its similarity to the evolved stimuli

If these evolved images are telling us something important about the tuning properties of IT neurons, then we should be able to use them to predict neurons' responses to novel images. The deep generator network had been trained to synthesize images from their encoding in layer fc6 of AlexNet (4096 units;(Krizhevsky et al., 2012)), so we used the fc6 space to find natural images similar to the evolved images. In particular, we could ask whether a neuron's response to a novel image was predicted by the distance in fc6 space between the novel image and the neuron's evolved synthetic image. To do this we calculated the activation vectors of the evolved synthetic images in AlexNet fc6 and searched for images with similar fc6 activation vectors. We used 2 databases for this: the first comprised ~60,000 images collected in our laboratory over several years, and the second set comprised 100,300 images from the ILSVRC2012 dataset (Russakovsky et al., 2015); we included 100 randomly sampled images from each of its 1,000 categories, and two additional ImageNet categories of *faces* [ID *n09618957*], *macaques* [ID *n02487547*], as well as 100 images of local animal care personnel with and without personal protective garb.

First, we focus on the evolution experiment for PIT single unit Ri-17. This cell evolved a discrete shape near the top left of the image frame, comprising a darkly outlined pink convex shape with two dark circles and a dark vertical line between them (Fig. S5A). When tested

with the 2,550 natural images, this neuron responded best to images of monkeys, dogs, and humans (Fig. S5B). We propagated this evolved image through AlexNet along with the 100,300 ImageNet examples, ranked all the fc6 vectors by their Pearson correlation to the evolved image vector, and identified the closest, middle, and farthest 100 matches. The synthetic image showed an average vector correlation of 0.38 to the closest images, 0.06 to the middle and −0.14 to the farthest images. The 9 nearest ImageNet images were cats, dogs and monkeys (Fig. S5C). To visualize the common shape motifs of this image cluster, we identified the individual fc6 units most strongly activated by the synthetic image and used activation maximization (*deepDreamImage.m*) to generate examples of preferred shapes for those fc6 units. All the units preferred round tan/pink regions with small dark spots (Fig. S5D). To rule out that these matches could be due to an overrepresentation of animals in ImageNet, we also looked at the least correlated matches, which were indeed not animals, but were pictures of places, rectilinear textures, or objects with long straight contours (Fig. S5E).

We applied this image-search approach to all evolution experiments by identifying the top 100 matches to every synthetic image in fc6 space (the Pearson correlation coefficients of these images ranged from 0.30 to 0.61, median 0.36) and visualized the WordNet (PrincetonUniversity, 2010) labels of the matching images via word clouds. In monkey Ri, whose array showed natural-image preferences for faces, the categories that best matched the synthetic images were “macaques”, “toy terrier”, and “Windsor tie” (the latter containing images of faces and bodies) (Fig. S5F); in contrast, in monkey Y1, where most of the neurons in the array had shown natural-image preferences for places, the categories that best matched were “espresso maker,” “rock beauty” (a type of fish), and “whiskey jug”; by inspection these images all contained extended contours (Fig. S5G). We confirmed this matching trend by quantifying the WordNet hierarchy labels associated with every matched natural image (Table S2).

To find out whether similarity in fc6 space between a neuron’s evolved synthetic image and a novel natural image predicted that neuron’s response to that novel image, we first performed 3–4 independent evolution experiments using the same (single- or multi-) unit in each of three animals. After each evolution, we took the top synthetic image from the final generation and identified the top 10 nearest images in fc6 space, 10 images from the middle of the distance distribution and the farthest (most anticorrelated) 10 images (9 of each shown in Figure 7A). Then, during the same recording session, we presented these images to the same IT neurons and measured the responses to each group (*near, middle, far*) as well as to all 40 evolved images of the last generation. Figure 7B shows that synthetic images gave the highest responses, the nearest natural images the next highest responses, and the middle and farthest images the lowest. To quantify this observation, we fit linear regression functions between the ordinal distance from the synthetic image (near, middle, and far) and the unit’s mean responses, and found median negative slopes ranging from −5.7 to −21.1 spikes/s across monkeys (Table S3). Thus distance from the evolved synthetic image in fc6 space predicted responses to novel natural images. This does not indicate that this space is the best model for IT response properties; instead, this shows that it is possible to use the neurons’ evolved images to predict responses to other images. Importantly, responses to the synthetic images were the highest of all.

Invariance to evolved vs natural images

IT neurons retain selectivity despite changes in position, size and rotation (Ito et al., 1995; Kobatake and Tanaka, 1994), although it has been reported that more selective neurons are less transformation-invariant (Zoccolan et al., 2007). The later observation is also consistent with the alternative interpretation that the more optimal a stimulus is for the neuron, the less invariant the neuron will be, and this is consistent with what we found. To compare the invariance of IT neurons to synthetic and natural images, we presented 3 natural and 3 evolved synthetic images at different positions, sizes and fronto-parallel rotations in two animals (monkeys Ri and Gu). The natural images were chosen from the nearest, middle, and farthest matches from ImageNet. The synthetic images were chosen from the final generation. Every image was presented at three positions relative to the fovea: (−2.0°, −2.0°), (−2.0°, 2.0°) and (0.0°, 0.0 °); three sizes (width of 1°, 2°, 4°) and 4 rotations (0°, 22°, 45° and 80°, counterclockwise from horizontal) (Fig. S6A). Invariance was defined as the similarity (correlation coefficient) in the neuron's rank order of preferences for images under different transformation conditions (the more similar the rank order, the higher the invariance). The rank order was better maintained across transformations for the natural images than for the synthetic images (Fig. S6 B&C). Thus the degree of invariance for these neurons changed depending on the stimulus set, and the neurons were the least invariant for the more optimal synthetic images. This result suggests that the degree of invariance measured for particular neurons may not be a fixed feature of that neuron, but rather may depend on the effectiveness of the stimulus used to test the invariance.

Evolution of preferred stimuli for populations of neurons

Single- and multi-units in IT successfully guided the evolution of synthetic images that were stronger stimuli for the neuron guiding the evolution than large numbers of natural images. To see if our technique could be used to characterize more coarsely sampled neuronal activity than a single site, we asked whether we could evolve strong stimuli for all 32 sites on an array. Each of the chronically implanted arrays had up to 32 visually responsive sites, spaced 400 μm apart. We conducted a series of evolution experiments in 3 monkeys (Ri, Gu, and Y1) guided by the average population response across the array. Evolution experiments for all 3 monkeys showed increasing responses to synthetic images over generations compared to reference images: The median population response changes to synthetic images for monkeys Ri, Gu and Y1 were 9.4 spikes/s/generation (2.8 – 19.6, 25th-75th percentile), 30.7 (17.6 – 48.3, 25th-75th %tile) and 27.8 (18.8 – 39.6, 25th-75th %tile). In these population-guided evolutions, 61%, 93% and 99.5% of individual sites showed increases in firing rate (statistical significance defined by fitting an exponential function to 250 resampled firing rate/generation curves per site; an increase was significant if the 95% CI of the amplitude bootstrap distribution did not include zero). Therefore larger populations of IT neurons could successfully create images that were on average strong stimuli for that population. When the populations were correlated in their natural-image preferences, the synthetic images were consistent with those evolved by individual single sites in the array: for example, in monkey Ri, the population-evolved images contained shape motifs commonly found in ImageNet pictures labeled "macaques," "wire-haired fox terrier," and "Walker hound." This suggests that the technique can be used with coarser sampling

techniques than single-unit recordings, such as local field potentials, electrocorticography electrodes, or even functional magnetic resonance imaging.

Testing XDREAM using the ground truth of primary visual cortex

We recorded from one single unit and three multiunit sites (six evolution experiments total) in monkey Vi, which had a chronic microelectrode array in V1. The stimuli were centered on each receptive field (measuring ~0.79° square root of the area), but wer e kept at the same size as in the IT experiments (3°×3°-wide). In addition to the synthetic images, we interleaved reference images of gratings (3°×3° area) of different orientations (0°, 45°, 90° and 1 35°) and spatial frequencies (~0.5, 1 and 2 cycles/°) at 100% contrast. In all experiments, neurons showed an increase in firing rate to the synthetic images (median 84.0 spikes/s/generation, 77.4 – 91.2, 25th-75th percentile; Table S1). Thus on average, V1 sites, like IT, responded well to late-generation synthetic images (Table S4). To measure the distribution of orientations of the region of the synthetic images that fell within each V1 receptive field (~0.8°×0.8°), we performed a discrete Fourier transform analysis on the central 0.8°×0.8° of the synthetic images and correlated the resulting spectrogram to the spectrograms expected from 16 gratings with orientations ranging from 0° to 135°. Across experiments, the mean correlation between the orientation content profile of the patch and the orientation tuning measured from the gratings was 0.59±0.09 (mean ± SEM), compared to 0.01±0.26 for a shuffled distribution (P-values ≤ 0.006 in 5/6 experiments, permutation test, $N_{iterations}$ = 999). Thus V1 neurons guided the evolution of images dominated by their independently measured preferred orientation.

DISCUSSION

We introduce XDREAM, a new algorithm for studying the response properties of visual neurons using a vast generative image space. Our approach is an extension of previous work that uses adaptive sampling to reveal neuronal tuning in the ventral stream (Carlson et al., 2011; Hung et al., 2012; Kalfas et al., 2017; Vaziri et al., 2014; Vaziri and Connor, 2016; Yamane et al., 2008). In our approach, adaptive sampling enables the exploration of a stimulus space potentially large enough to span the selectivity of ventral pathway neurons. Starting with random textures and evolving images based on neuronal responses, the algorithm created images that elicited large responses in V1, in different parts of IT, in single-units, in multi-units, and in average population responses. Remarkably, XDREAM evolved stimuli that evoked higher responses than the best natural images found by an extensive exploration of large image sets.

A powerful approach for modeling neuronal responses has been to use stimuli sampled in a fully-defined parametric space. For example, one innovative study (Yamane et al., 2008), which inspired our approach, defined response properties in IT according to 3D curvature and orientation. In a more recent study, Chang and Tsao (2017) used face images parameterized along appearance and shape axes to describe and predict neuronal responses in face patches. Parametric stimulus spaces lead to quantitative neuronal models that are easier to describe than models built on the learned, vast, and latent space of a deep network. But these approaches are complementary: standard parametric models operate in a

circumscribed shape space, so these models may not capture the entire response variability of the neuron. Generative neural networks are powerful enough to approximate a wide range of shape configurations, even shapes diagnostic of monkey faces or bodies. If the response evoked by the best stimulus in a parametric space is less than that evoked by a generative-network stimulus, it would indicate that the parametric model is overly restricted. This is important because a neuron could be tuned to different aspects of a stimulus: a face-preferring neuron might show tuning for high curvature, but curvature tuning alone is insufficient to explain the neuron's selectivity. A generative-network-based approach can serve as an independent, less constrained, test of parametric models. Indeed, in some instances, the evolved stimuli contained features of animal faces, bodies, and even animal-care staff known to the monkeys, consistent with theories of tuning to object categories and faces, whereas in other instances, the evolved stimuli were not identifiable objects or object parts, suggesting aspects of neuronal response properties that current theories of visual cortex have missed. Thus our approach can also serve as an initial step in discovering novel shape configurations for subsequent parametric quantification. In sum, our approach is well-positioned to corroborate, complement, contrast with, and extend the valuable lessons learned from these previous studies.

What are these powerful stimuli 'dreamed of' by the neurons? To start to answer this, we looked for images that were nearest to the evolved images in AlexNet layer fc6 space, because the generative network had been trained on that space and because higher layers of this network can predict neuronal responses (Yamins et al., 2014). Similarity to the evolved images in fc6 space predicted neuronal selectivity. For neurons defined using natural images as monkey- and face-preferring the closest images to the evolved ones were macaques and dogs or contained faces and torsos of monkeys or other mammals. In contrast, for place-preferring neurons, the nearest images instead included a variety of objects with extended straight contours, such as espresso makers and moving vans.

Although the evolved synthetic images were not life-like, sometimes not even identifiable objects, they nevertheless tell us something quite novel about what information might be encoded by IT neurons. PIT neurons have been reported to be selective for low-level features like radial gratings (Pigarev et al., 2002), color (Zeki, 1977), or single eyes (Issa and DiCarlo, 2012), and face cells in CIT have been shown to be tuned to distinct face-space axes (Chang and Tsao, 2017; Freiwald et al., 2009). Our experiments revealed that both PIT and CIT neurons evolved complex synthetic images, usually consisting of intact objects with multiple, spatially organized, features and colors, not just disjointed object parts. This is consistent with the finding that even cells tuned to particular feature parameters, like eye size, do not respond well to that parameter in isolation, but rather only in context (Freiwald et al., 2009). It has been proposed that neurons are tuned to abstract parameters, such as axes (Chang and Tsao, 2017; Freiwald et al., 2009; Leopold et al., 2006) that distinguish between things in the environment, rather than being tuned to the things themselves. That is, neuronal responses may carry information about an object in the environment, not a veridical representation of it. The unrealistic nature of our evolved images, plus the fact that these images were more effective than most or all the images in an extensive natural-image database, suggest that these images may lie somewhere on tuning axes that extend beyond anything the animal would normally encounter. Lastly, the neurons that evolved images

resembling lab staff wearing face masks indicate a major role for experience in the development of neuronal response properties (Arcaro et al., 2017).

It is unclear whether our approach has uncovered globally optimum stimuli for these cells, but it is equally unclear whether there should be a single global optimum stimulus for a neuron. Different evolutions for the same units yielded synthetic images that shared some features but differed in others (e.g. Figure 4B), as would be expected from a cell that shows invariance to nuisance transformations like position, color, or scale. Even in early visual areas complex cells respond equally well to the same stimulus in multiple locations within a cell's receptive field (Hubel and Livingstone, 1985; Hubel and Wiesel, 1968), and this complexification/OR gate-like operation is likely to occur at multiple levels in the hierarchy (Hubel and Livingstone, 1985; Riesenhuber and Poggio, 1999), so it follows that multiple different feature combinations could strongly activate a particular neuron in IT. However, the facts that our approach can recover known properties of neurons; that it can discover unknown response properties; that the synthetic images can be better than any in a large set of natural stimuli; and even better than highly-optimized stimuli based on convolutional neural network models of neurons—currently the best models of visual neurons (Supplementary Fig. S7 Supplementary Table S6) —all indicate that we have discovered stimuli that give us a best glimpse at what IT neurons are tuned to. These results complement classical methods for defining neuronal selectivities and demonstrate the potential for uncovering internal representations in any modality that can be captured by generative neural networks.

STAR METHODS

CONTACT FOR REAGENT AND RESOURCE SHARING

Requests for resources should be directed to and will be fulfilled by the Lead Contact, Margaret Livingstone (mlivingstone@hms.harvard.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

All procedures were approved by the Harvard Medical School Institutional Animal Care and Use Committee, and conformed to NIH guidelines provided in the Guide for the Care and Use of Laboratory Animals.

Behavior—Six adult male macaca mulatta (9 −13 kg; 5–14 years old) and one adult male macaca nemestrina (9kg, 10 years old) were socially housed in standard quad cages on 12/12hr light/dark cycles.

Recording arrays—Monkeys Ri, Gu, Ge, Jo, and Y1 were implanted with custom floating microelectrode arrays manufactured by MicroProbes for Life Sciences (Gaithersburg, MD); each had 32 platinum/iridium electrodes per ceramic base, electrode lengths of 2–5 mm, impedances between 0.7– 1.0 MΩ. The recording sites in the chronic arrays were stable across days; in particular, for all the experiments where we compare evolutions from the same site recorded on different days, the correlation in category preference was greater than 0.89. Monkey Vi was implanted with a 96channel Utah array

(Blackrock microsystems, Salt Lake City, Utah). Monkey B3 had an acute recording chamber, and neuronal activity was recorded using a 32 channel NeuroNexus Vector array (Ann Arbor, Michigan) that was inserted each recording day.

Surgical procedures—All animals were implanted with custom-made titanium or plastic headposts before fixation training. After several weeks of fixation training, the animals underwent a second surgery for array or chamber implantation. PIT insertion sites were just anterior to the inferior occipital sulcus; CIT sites were on the lower lip of the STS 6–8 mm anterior to the interaural line. All surgeries were done under full surgical anesthesia using sterile technique.

METHOD DETAILS

Behavioral task—The monkeys were trained to perform a fixation task. They fixated on a 0.2°diameter fixation spot in the middle of the screen. Eye position was monitored using an ISCAN system (Woburn, MA). Animals were rewarded with a drop of water or juice for maintaining fixation within 1.0° of the fixation spot for 2–7 image presentations; the interval was gradually decreased over the experimental session as the monkey's motivation decreased.

Physiological recording—Neural signals were amplified and extracellular action potentials were isolated using the box method of an on-line spike sorting system (Plexon, Dallas, TX). Spikes were sampled at 40 kHz.

Deep Generative Neural Network—The pre-trained generative network (Dosovitskiy and Brox, 2016) was downloaded from the authors' website (lmb.informatik.uni-freiburg.de/resources/software.php) and used without further training with the Caffe library (Jia et al., 2014) in Python. To synthesize an image from an input image code, we forward propagated the code through the generative network, clamped the output image pixel values to the valid range between 0 and 1, and visualized them as an 8-bit color image. Some images synthesized by the network contained a patch with a stereotypical shape that occurred in the center right of the image (e.g., second and third image in Figure 4B, and 'late synthetic' in Figure 5 C,D). This was identified as an artifact of the network commonly known as "mode collapse" and it appeared in the same position in a variety of contexts, including a subset of simulated evolutions. This artifact was easily identifiable and it did not affect our interpretations. In the future, more modern GNN training methods should avoid this problem (personal correspondence with Alexey Dosovitskiy).

Initial generation—The initial generation of image codes for all evolution experiments reported here was 40 achromatic textures constructed from a set of Portilla and Simoncelli (2000) textures, derived from randomly sampled photographs of natural objects on a gray background. We started from all-zero codes and optimized for pixelwise loss between the synthesized images and the target images using backpropagation through the network for 125 iterations, with a learning rate linearly decreasing from 8 to 1×10^{-10}. The resulting image codes produced blurred versions of the target images, which was expected from the

pixelwise loss function and accepted because the initial images were intended to be quasi-random textures.

Genetic algorithm—The algorithm began with an initial population of 40 image codes ('individuals'), each consisting of a 4096-dimensional vector ('genes') and associated with a synthesized image. Images were presented to the subject, and the corresponding spiking response was used to calculate the 'fitness' of the image codes by transforming the firing rate into a Z-score within the generation, scaling it by a selectiveness factor of 0.5, and passing it through a softmax function to become a probability. The 10 highest-fitness individuals were passed on to the next generation without recombination or mutation. Another 30 children image codes were produced from recombinations between two parent image codes from the current generation, with the probability for each image code to be a parent being its fitness. The two parents contributed unevenly (75%:25%) to any one child. Individual children genes had a 0.25 probability of being mutated, with mutations drawn from a 0-centered gaussian with standard deviation 0.75. Hyperparameter values were not extensively optimized. All source code is available upon request.

Evolutions in CaffeNet units—We selected 100 random units each in 4 layers in CaffeNet as targets for evolution. For convolutional layers, only the center unit in each channel was used. Each unit was evolved for 500 generations, 10,000 image presentations total. The best image in the last generation was used in the analysis, although most of the total activation increase was achieved by 200 generations. As a control, we recorded activations of the units to all 1,431,167 images in the ILSVRC2012 dataset, including the training set of CaffeNet. To visualize the ground truth best in CaffeNet layer conv1, 11 × 11 × 3 images were produced from the 11 × 11 × 3 filter weights according to $image(x, y, c) = 0.5 + sign(weight(x, y, c))/2$, because this is a linear transformation. In other words, positive weights corresponded to a pixel value of 1 and negative weights 0, and the ground truth optimum only contained black, white, red, green, blue, magenta, cyan, and yellow pixels. To further visualize the magnitude of the weights, each pixel was made transparent to a grey checkerboard in inverse proportion to its contribution to the overall activation, so that the greyer a pixel, the less it affected the overall activation.

Visual stimuli—We used MonkeyLogic2 (https://www.nimh.nih.gov/labs-at-nimh/research-areas/clinics-and-labs/ln/shn/monkeylogic/index.shtml) as the experimental control software. Images were presented on an LCD monitor 53 cm in front of the monkey at a rate of 100 ms on, 100–200 ms off. Most of the images were from (Konkle et al., 2010); human and monkey images were from our lab, and the rest of the images were from public domain databases.

QUANTIFICATION AND STATISTICAL ANALYSIS

Spike rate changes during evolutions—We defined the neuronal response as the spike rate measured in the 70–200 ms time window after image onset and subtracted the rate in the 0–60 ms window. In the evolution experiments, there were 40 synthetic and 40 reference images per generation, each presented once. To track firing rate change per generation, we averaged responses to all 40 synthetic and 40 reference images separately. To estimate the

inter-generational changes in response we fit the mean response/generation separately for synthetic and reference images with a decaying exponential function $-a \times \exp\left(-\frac{x}{\tau}\right) + c$, which models firing rate as starting at the first block with rate and asymptotically approaching the rate ($a + c$) with decay constant τ. We restricted the amplitude change to be within physiologically plausible values (magnitude no more than the absolute maximum rate difference between any two generations in that day). To assess statistical significance, we generated new mean rate per generation curves by resampling responses (with replacement) from each of the 40 synthetic and 40 reference image presentations within one generation, then fit the exponential function each time (N = 500 repetitions). The 95% confidence intervals reported are the 0.025/500 and 0.975/500 (12th and 488th) value of the bootstrapped distribution. We used responses from all generations except for one acute experiment in monkey B3, where a single-unit isolation change nullified the initial 15 generations and two chronic experiments in monkey Ge, where spike thresholding adjustments nullified the initial 17 generations in one experiment, and the final 35 generations in the other.

Responses to evolved vs. reference images—For every evolution experiment, we measured how the neurons' maximum firing rate to the evolving images compared to the neurons' maximum firing rate to the reference images. To estimate maximum firing rate and its standard error, we re-sampled (with replacement) all single-trial responses to synthetic or reference images, computed the observed maximum per sample, repeated 500 times and then reported the mean and standard deviation across samples. To determine if the max firing rates to synthetic images were different from those to reference images, we used a randomization test: We created a null distribution of differences between two maximum responses by randomly sampling twice from the pooled distribution of synthetic and reference images and measuring the difference in the maximum response between the samples. We repeated this process 499 times, and then measured how often the absolute differences for the mixed-distribution were larger than the observed absolute maximum difference. This two-tailed test indicated if either the reference or synthetic images evoked a significantly larger maximum response.

Assigning natural image labels to the evolved images—Every evolved image was propagated through AlexNet and its fc6-layer activation was compared to those of 100,300 natural images sampled from ImageNet. Every photograph in ImageNet is labeled by categories defined by WordNet, a hierarchically organized lexical database (PrincetonUniversity, 2010). After ranking every natural image by its proximity to the evolved image, measured by Pearson correlation coefficient, we used a tree search algorithm to crawl through each labeled image's label hierarchy for the specific search terms--"macaque," "monkey", "face," and "appliance" ("place" was not used because place images often contained people). We measured the frequency of labels associated with all evolved images for every subject. To estimate confidence intervals for every observed frequency, we re-sampled the top matches to each evolved image 200 times (with replacement) and repeated the analysis. To test if the frequency of photographs labeled "monkeys" and "appliance" were statistically different between subjects Ri and Y1, we used a permutation test. The null hypothesis was that these frequency values arose from the same distribution, so

we shuffled labels from the Ri and Y1 populations, sampling twice with replacement, and measured the difference, 500 times. We then compared the observed difference in frequency values with the null distribution.

Calculation of stimulus rank order—Background was firing rate 0 to 60 ms after stimulus onset. Spikes evoked in response to each stimulus were averaged over all stimulus presentations from 70 to 200 ms after stimulus onset, minus background. Rank order was determined from the evoked firing rate for each stimulus.

DATA AND SOFTWARE AVAILABILITY

Source code is available at https://github.com/willwx/XDream.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

This work was supported by NIH grants R01EY16187, R01EY25670, R01EY011379, P30EY012196, R01EY026025, NSF STC award CCF-1231216 to CBMM at MIT. We thank Richard Born, Shimon Ullman and Christof Koch for comments.

REFERENCES

Abbasi-Asl R, Chen Y, Bloniarz A, Oliver M, Willmore BDB, Gallant JL, and Yu B (2018). The DeepTune framework for modeling and characterizing neurons in visual cortex area V4. https://www.biorxiv.org/content/10.1101/465534v1

Arcaro MJ, Schade PF, Vincent JL, Ponce CR, and Livingstone MS (2017). Seeing faces is necessary for face-domain formation. Nat Neurosci 20, 1404–1412. [PubMed: 28869581]

Barlow HB, Blakemore C, and Pettigrew JD (1967). The neural mechanism of binocular depth discrimination. J Physiol (Lond) 193, 327–342. [PubMed: 6065881]

Bashivan P, Kar K, and DiCarlo JJ (2018). Neural Population Control via Deep Image Synthesis https://www.biorxiv.org/content/10.1101/461525v1.

Carlson ET, Rasquinha RJ, Zhang K, and Connor CE (2011). A sparse object coding scheme in area V4. Curr Biol 21, 288–293. [PubMed: 21315595]

Chang L, and Tsao DY (2017). The Code for Facial Identity in the Primate Brain. Cell 169, 1013–1028 e1014. [PubMed: 28575666]

Desimone R, Albright TD, Gross CG, and Bruce C (1984). Stimulus-selective properties of inferior temporal neurons in the macaque. J Neurosci 4, 2051–2062. [PubMed: 6470767]

Dosovitskiy A, and Brox T (2016). Generating Images with Perceptual Similarity Metrics based on Deep Networks. Advances in Neural Information Processing (NIPS).

Erhan D, Bengio Y, Courville A, and Vincent P (2009). Visualizing higher-layer features of a deep network. University of Montreal.

Felleman DJ, and Van Essen DC (1987). Receptive field properties of neurons in area V3 of macaque monkey extrastriate cortex. J Neurophysiol 57, 889–920. [PubMed: 3585463]

Freiwald WA, Tsao DY, and Livingstone MS (2009). A face feature space in the macaque temporal lobe. Nat Neurosci 12, 1187–1196. [PubMed: 19668199]

Gaska JP, Jacobson LD, Chen HW, and Pollen DA (1994). Space-time spectra of complex cell filters in the macaque monkey: a comparison of results obtained with pseudowhite noise and grating stimuli. Visual Neuroscience 11, 805–821. [PubMed: 7918230]

Gross CG, Rocha-Miranda CE, and Bender DB (1972). Visual properties of neurons in inferotemporal cortex of the Macaque. J Neurophysiol 35, 96–111. [PubMed: 4621506]

Hubel D, and Wiesel T (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. The Journal of Physiology 160, 106. [PubMed: 14449617]

Hubel DH (1959). Single unit activity in striate cortex of unrestrained cats. J Physiol (Lond) 147, 226–238. [PubMed: 14403678]

Hubel DH, and Livingstone MS (1985). Complex-unoriented cells in a subregion of primate area 18. Nature 315, 325–327. [PubMed: 2987703]

Hubel DH, and Wiesel TN (1968). Receptive fields and functional architecture of monkey striate cortex. J Physiol (Lond) 195, 215–243. [PubMed: 4966457]

Hung CC, Carlson ET, and Connor CE (2012). Medial axis shape coding in macaque inferotemporal cortex. Neuron 74, 1099–1113. [PubMed: 22726839]

Issa EB, and DiCarlo JJ (2012). Precedence of the eye region in neural processing of faces. J Neurosci 32, 16666–16682. [PubMed: 23175821]

Ito M, Tamura H, Fujita I, and Tanaka K (1995). Size and position invariance of neuronal responses in monkey inferotemporal cortex. J Neurophysiol 73, 218–226. [PubMed: 7714567]

Jia Y, Shelhamer E, Donohue J, Karayev S, Long J, Girshick R, Guadarrama S, and Darrell T (2014). Caffe: Convolutional Architecture for Fast Feature Embedding. Proceedings of the 22nd ACM international conference on Multimedia, 675–678.

Jones JP, and Palmer LA (1987). The two-dimensional spatial structure of simple receptive fields in cat striate cortex. J Neurophysiol 58, 1187–1211. [PubMed: 3437330]

Kalfas I, Kumar S, and Vogels R (2017). Shape Selectivity of Middle Superior Temporal Sulcus Body Patch Neurons. eNeuro 4.

Kobatake E, and Tanaka K (1994). Neuronal selectivities to complex object features in the ventral visual pathway of the macaque cerebral cortex. J Neurophysiol 71, 856–867. [PubMed: 8201425]

Konkle T, Brady TF, Alvarez GA, and Oliva A (2010). Conceptual distinctiveness supports detailed visual long-term memory for real-world objects. J Exp Psychol Gen 139, 558–578. [PubMed: 20677899]

Krizhevsky A, Sutskever I, and Hinton GE (2012). ImageNet classification with deep convolutional neural networks. NIPS, 1106–1114.

Leopold DA, Bondar IV, and Giese MA (2006). Norm-based face encoding by single neurons in the monkey inferotemporal cortex. Nature 442, 572–575. [PubMed: 16862123]

Michael CR (1978). Color vision mechanisms in monkey striate cortex: simple cells with dual opponent-color receptive fields. J Neurophysiol 41, 1233–1249. [PubMed: 100586]

Mordvintsev A, Olah C, and Tyka M (2015). Inceptionism: Going deeper into neural networks. Google Research Blog.

Nguyen A, Dosovitskiy A, Yosinski J, Brox T, and Clune J (2016). Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. 30th Conference on Neural Information Processing Systems (NIPS).

Pasupathy A, and Connor CE (1999). Responses to contour features in macaque area V4. J Neurophysiol 82, 2490–2502. [PubMed: 10561421]

Pasupathy A, and Connor CE (2002). Population coding of shape in area V4. Nat Neurosci 5, 1332–1338. [PubMed: 12426571]

Pigarev IN, Nothdurft HC, and Kastner S (2002). Neurons with radial receptive fields in monkey area V4A: evidence of a subdivision of prelunate gyrus based on neuronal response properties. Exp Brain Res 145, 199–206. [PubMed: 12110960]

Portilla J, and Simoncelli EP (2000). A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients. International Journal of Computer Vision 40, 49–71.

PrincetonUniversity (2010). About WordNet. WordNet. Princeton University.

Riesenhuber M, and Poggio T (1999). Hierarchical models of object recognition in cortex. Nat Neurosci 2, 1019–1025. [PubMed: 10526343]

Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, Huang Z, Karpathy A, Khosla A, Bernstein M, et al. (2015). ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision 115, 211–252.

Vaziri S, Carlson ET, Wang Z, and Connor CE (2014). A channel for 3D environmental shape in anterior inferotemporal cortex. Neuron 84, 55–62. [PubMed: 25242216]

Vaziri S, and Connor CE (2016). Representation of Gravity-Aligned Scene Structure in Ventral Pathway Visual Cortex. Curr Biol 26, 766–774. [PubMed: 26923785]

Walker EY, Sinz FH, Froudarakis E, Fahey PG, Muhammad T, Ecker AS, Cobos E, Reimer J, Pitkow X, and Tolias AS (2018). Inception in visual cortex: in vivo-silico loops reveal most exciting images. https://www.biorxiv.org/content/10.1101/506956v1.

Yamane Y, Carlson ET, Bowman KC, Wang Z, and Connor CE (2008). A neural code for threedimensional object shape in macaque inferotemporal cortex. Nat Neurosci 11, 1352–1360. [PubMed: 18836443]

Yamins DL, Hong H, Cadieu CF, Solomon EA, Seibert D, and DiCarlo JJ (2014). Performanceoptimized hierarchical models predict neural responses in higher visual cortex. Proc Natl Acad Sci U S A 111, 8619–8624. [PubMed: 24812127]

Zeki SM (1973). Colour coding in rhesus monkey prestriate cortex. Brain Res 53, 422–427. [PubMed: 4196224]

Zeki SM (1974). Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. J Physiol 236, 549–573. [PubMed: 4207129]

Zeki SM (1977). Colour coding in the superior temporal sulcus of rhesus monkey visual cortex. Proc R Soc Lond B Biol Sci 197, 195–223. [PubMed: 17866]

Zoccolan D, Kouh M, Poggio T, and DiCarlo JJ (2007). Trade-off between object selectivity and tolerance in monkey inferotemporal cortex. J Neurosci 27, 12292–12307. [PubMed: 17989294]

Highlights

- A generative deep neural network evolved images guided by neuronal firing
- Evolved images maximized neuronal firing in alert macaque visual cortex.
- Evolved images activated neurons more than large numbers of natural images.
- Similarity to evolved images predicts response of neurons to novel images.

Neurons guided the evolution of their own best stimuli with a generative deep neural network.

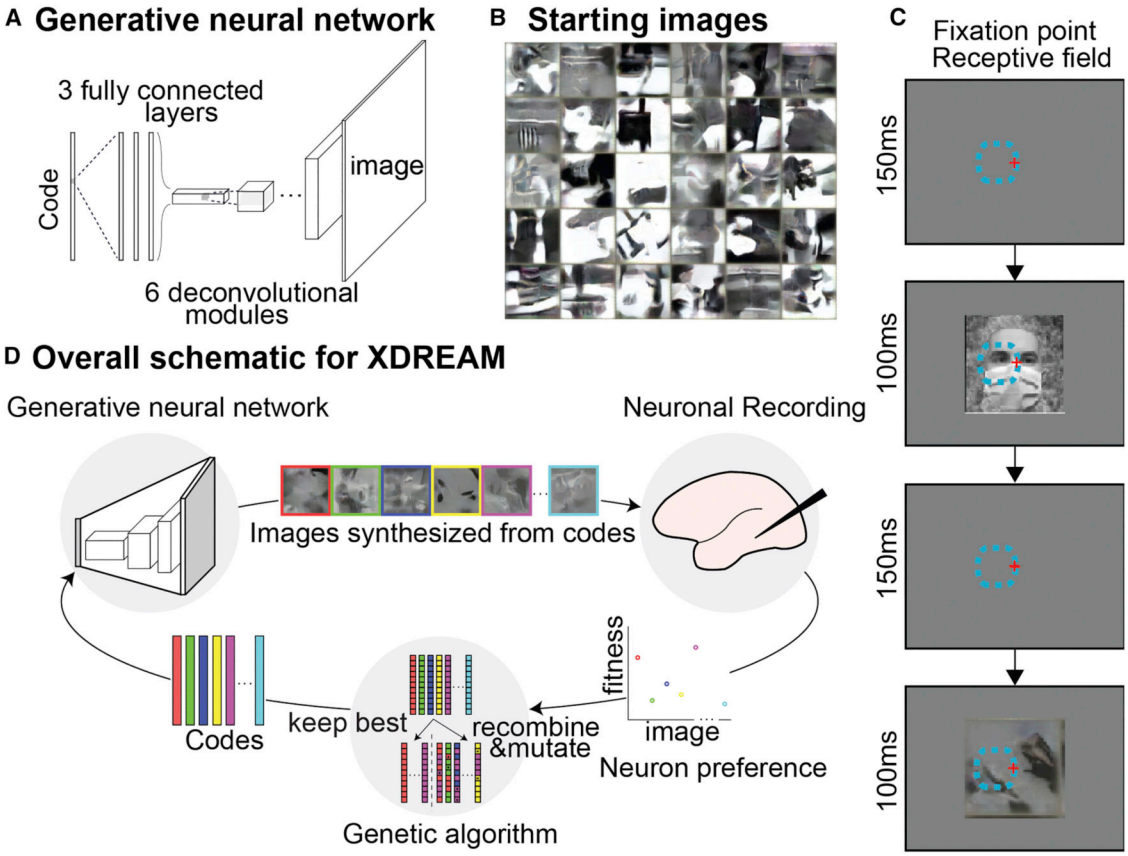

Figure 1. Synthesis of preferred stimuli via neuron-guided evolution.

(**A**) Generative adversarial network. Architecture of the pre-trained deep generative network (Dosovitskiy and Brox, 2016). The network comprised three fully connected layers and six deconvolutional modules. (**B**) The initial synthetic images were random achromatic Portilla and Simoncelli (2000) textures; 30 examples are shown here. (**C**) Behavioral task. Animals fixated within a 2° diamet er window while images were presented for 100 ms followed by a 100 to 200 ms blank period. Red cross: fixation point; dashed line, population RF. (**D**) Experimental flow. Image codes were forwarded through the deep generative adversarial network to synthesize images presented to the monkey. Neuronal responses were used to rank image codes, which then underwent selection, recombination and mutation to generate new image codes (for details see Star Methods).

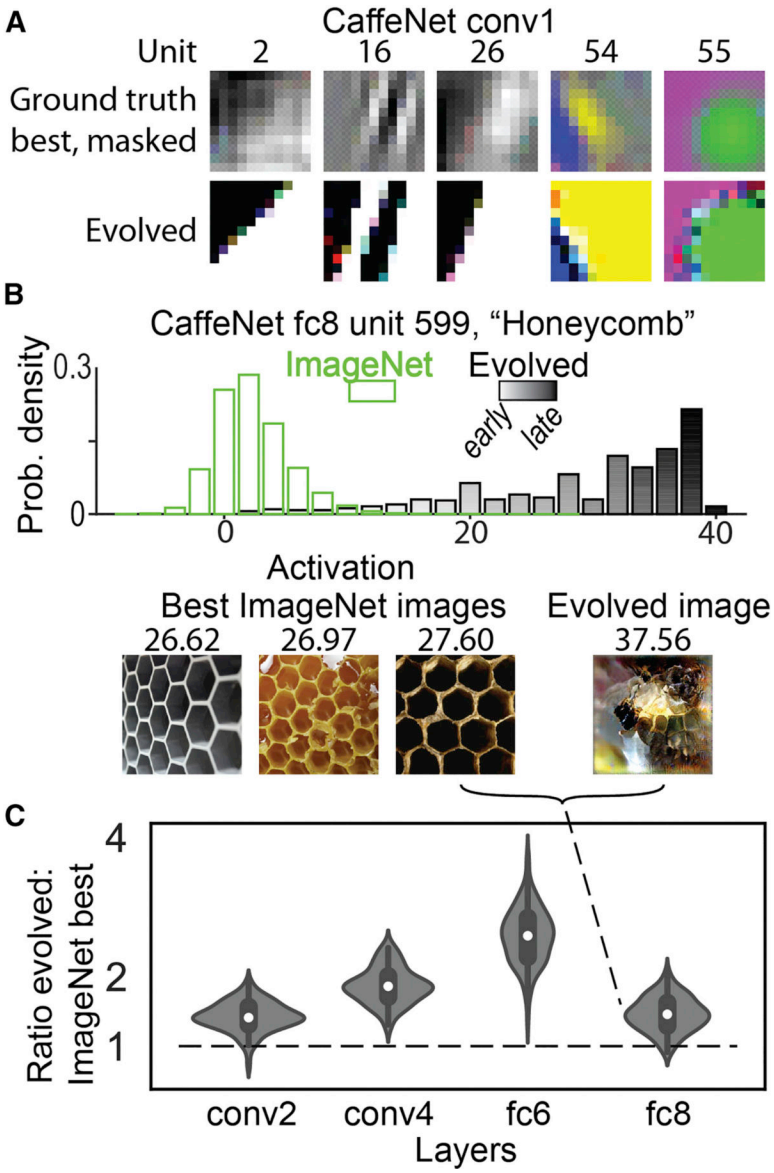

Figure 2. The XDREAM algorithm produced super stimuli for units in CaffeNet.

(A) Evolved images resembled the ground truth best images in the first layer of CaffeNet. In the ground truth best, transparency indicates the relative contribution of each pixel to the unit's activation. Only the center 11×11 pixels of the evolved images are shown, matching the filter size of the units. **(B, C)** Most evolved images activated artificial units more strongly than all of > 1.4 million images in the ILSVRC2012 dataset (Russakovsky et al., 2015). **(B)** Top, distribution of activations to ImageNet images and evolved images for one unit in the classification layer fc8, corresponding to the "Honeycomb" label. Grayscale gradient indicates generation of evolved images. Bottom, best 3 ImageNet images, and one evolved image, labeled with their respective activations. In this case, the evolved image activated the unit ~1.4× as strongly as did the best ImageNet image. **(C)** Distribution of (evolved / best in ImageNet) ratios across 4 layers in CaffeNet, 100 random units each layer.

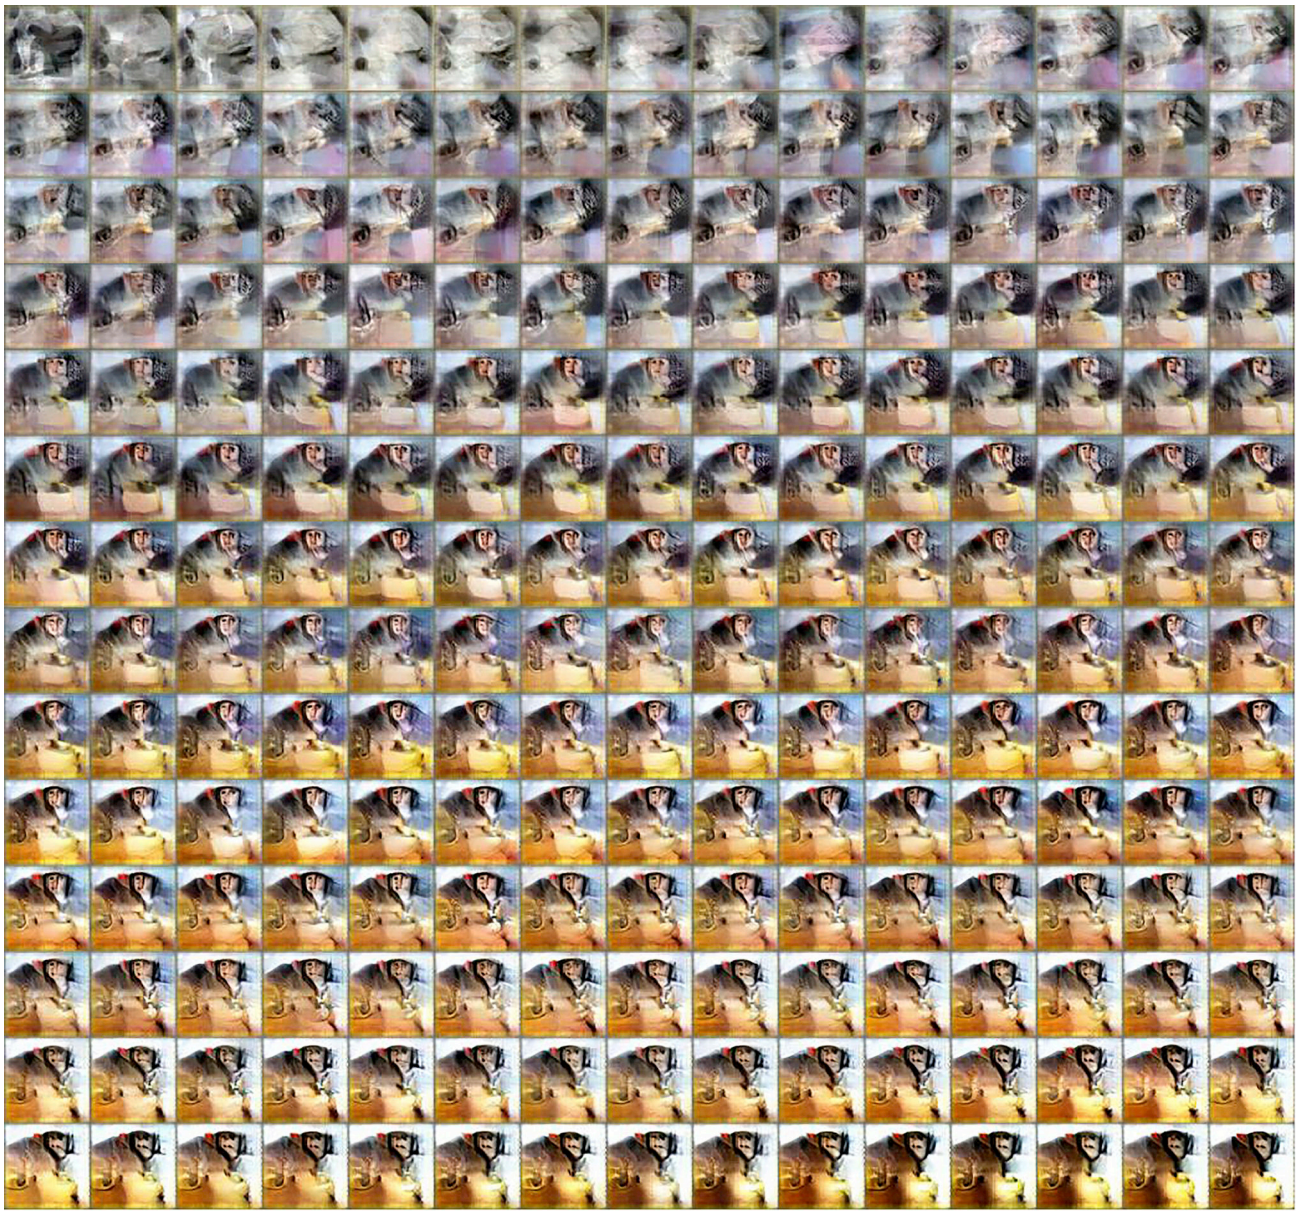

Figure 3. Evolution of synthetic images by a single monkey-selective neuron, Ri-10.

Each image is the average of the top 5 synthetic images for each generation (ordered L to R, top to bottom). The response of this neuron in each of these generations is shown in Figure 4A.

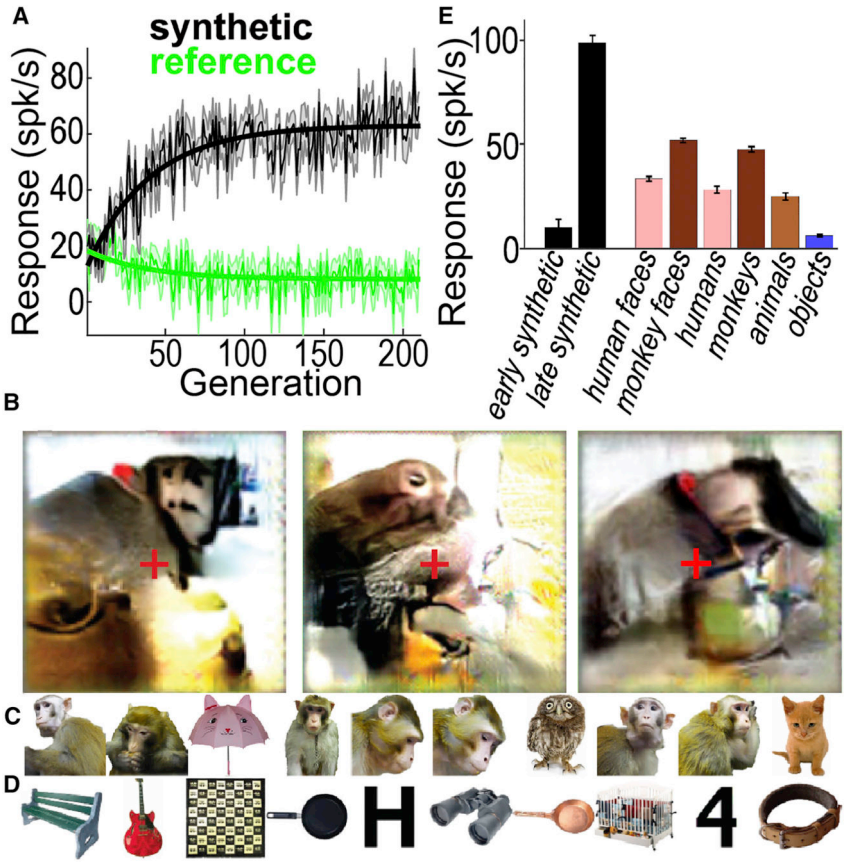

Figure 4. Evolution of synthetic images by maximizing responses of single neuron Ri-10 (same unit as Figure 3).

(**A**) Mean response to synthetic (black) and reference (green) images for every generation (spikes/s ± sem). Solid straight lines show an exponential fit to the response over the experiment. (**B**) Last-generation images evolved during three independent evolution experiments; the leftmost image corresponds to the evolution in (**A**); the other two evolutions were carried out on the same single unit on different days. Red crosses indicate fixation. The left half of each image corresponds to the contralateral visual field for this recording site. Each image shown here is the average of the top 5 images from the final generation. (**C**-**E**) Selectivity of this neuron to 2,550 natural images. (**C**) Top 10 images from this image set for this neuron. (**D**) Worst 10 images from this image set for this neuron. The entire rank ordered natural image set is shown in Fig. S2. (**E**) Selectivity of this neuron to different image categories (mean±sem). The entire image set comprised 2,550 natural images plus the best synthetic image from each of the 210 generations; 10–12 repetitions each; early synthetic was the first 10 generations and late the last 10. See Fig. S3 for additional independent evolutions from this site.

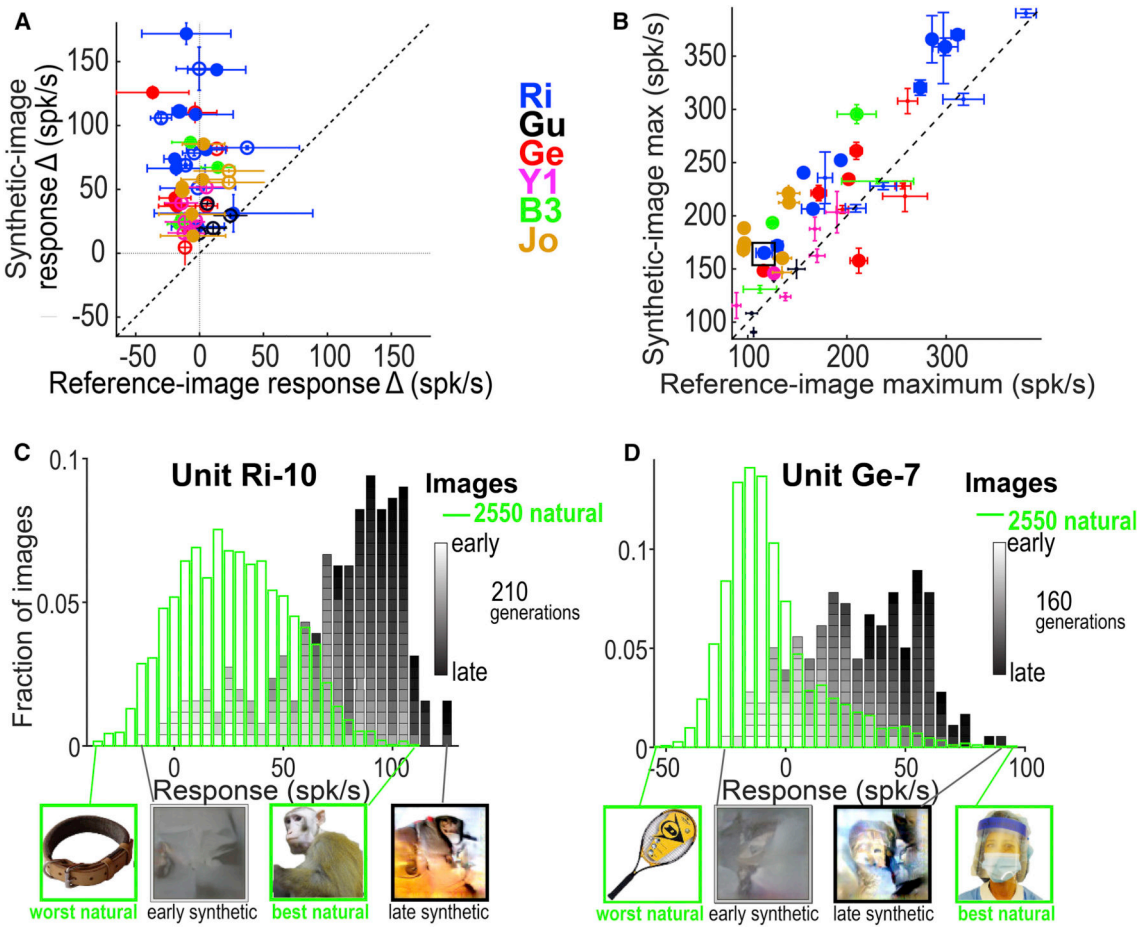

Figure 5. Evolutions for other IT cells.

(A) Change in response to synthetic vs reference images over generations. Each point shows the mean change in firing rate to reference versus synthetic images in each experiment (change estimated by the amplitude coefficient of an exponential function fitted to the neuron's mean response per generation; error bars = ±sem, per bootstrap, 500 iterations of data re-sampling). Solid circles indicate single units; open circles multi-units. **(B)** Scatter plot of maximum responses across all images for synthetic vs reference images (measured across all generations, max±SE per bootstrap). Colors indicate animal. The size of the circle indicates statistical significance (large circle: $P < 0.03$ after false discovery correction). Black square indicates the experiment in Figure 3. **(C)** Histogram of response magnitudes to natural (green) and synthetic (gray-to-black) images for unit Ri-10 (same unit as Figures 3&4). Below the histogram are shown the best and worst natural and synthetic images. **(D)** Same for unit Ge-7. The evolution for this neuron can be seen in Video S1.

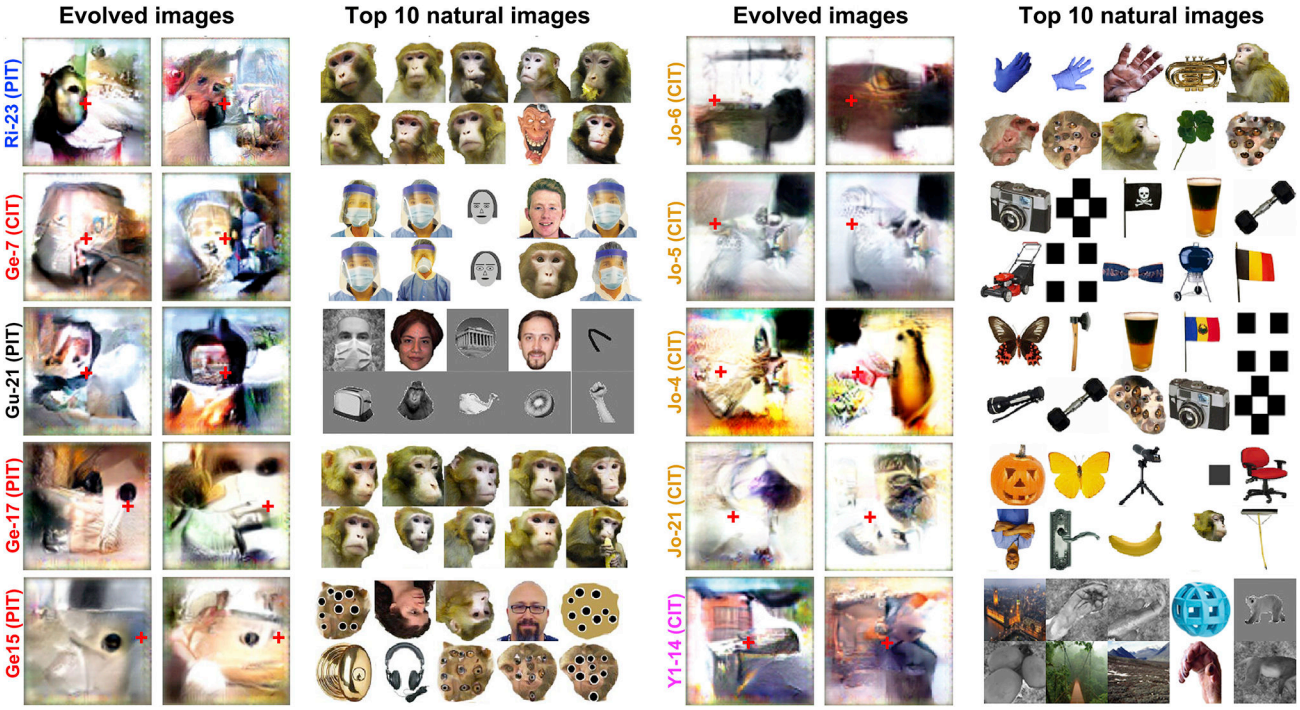

Figure 6. Evolution of synthetic images in other IT neurons.

Each pair of large images show the last-generation synthetic images from two independent evolution experiments for a single chronic recording site in 5 different animals. To the right of the synthetic images are shown the top 10 images for each neuron from a natural image set. Red crosses indicate fixation. The arrays were in the left hemisphere of monkey Jo, and in the right hemisphere of all the other animals. The natural images shown interleaved during each evolution were from either a 108 reference image set containing faces, bodies, places, and line segments (used for cells Gu-21 and Y1–14) or the set of 2550 natural images rank ordered in Fig. S2 for unit Ri-10 (used for all the other units in this figure). We also conducted a separate session testing Ge-7 and Ri-10 responses to natural and synthetic images two days after one of their evolution experiments.

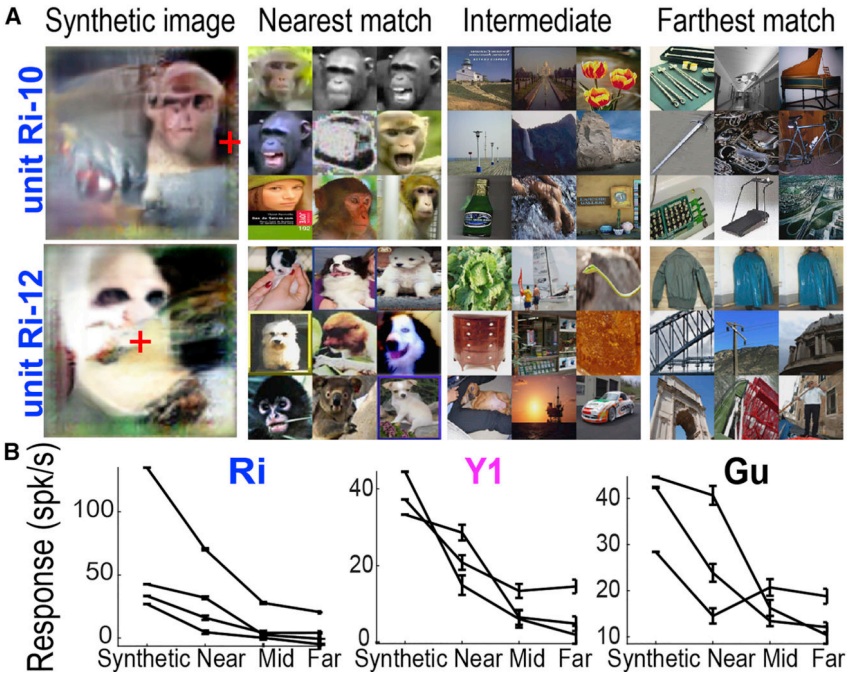

Figure 7. Using evolved images to predict responses to novel images.

(**A**) Final-generation synthetic images from units Ri-10 and Ri-12 and the closest, intermediate, and farthest 9 images from the image set for each. For unit Ri-10 we used the 60,000 image database, and for unit Ri-12 the 100,300 image database. (**B**) Responses from each unit to the last-generation evolved images compared to the nearest, intermediate, and farthest images from the evolved images in the fc6 space of AlexNet (mean±sem).