Behavioral/Cognitive

Catecholamine-Mediated Increases in Gain Enhance the Precision of Cortical Representations

Christopher M. Warren,[1,2] Eran Eldar,[3] Ruud L. van den Brink,[1,2] Klodianna-Daphne Tona,[1,2] Nic J. van der Wee,[2,4] Eric J. Giltay,[4] Martijn S. van Noorden,[4] Jos A. Bosch,[5,6] Robert C. Wilson,[7] Jonathan D. Cohen,[1,2] and Sander Nieuwenhuis[8,9]

[1]Department of Psychology, Leiden University, 2333 AK Leiden, The Netherlands, [2]Leiden Institute for Brain and Cognition, Leiden University, Leiden, 2300 RC Leiden, The Netherlands, [3]Welcome Trust Centre for Neuroimaging, University College London, London WC1N 3BG, United Kingdom, [4]Department of Psychiatry, Leiden University Medical Center, 2333 ZA Leiden, The Netherlands, [5]Department of Clinical Psychology, University of Amsterdam, 1018 XA Amsterdam, The Netherlands, [6]Mannheim Institute of Public Health, Heidelberg University, 68167 Mannheim, Germany, [7]Department of Psychology, University of Arizona, Tucson, Arizona 85721, and [8]Department of Psychology and [9]Princeton Neuroscience Institute, Princeton University, Princeton, New Jersey 08540

Neurophysiological evidence suggests that neuromodulators, such as norepinephrine and dopamine, increase neural gain in target brain areas. Computational models and prominent theoretical frameworks indicate that this should enhance the precision of neural representations, but direct empirical evidence for this hypothesis is lacking. In two functional MRI studies, we examine the effect of baseline catecholamine levels (as indexed by pupil diameter and manipulated pharmacologically) on the precision of object representations in the human ventral temporal cortex using angular dispersion, a powerful, multivariate metric of representational similarity (precision). We first report the results of computational model simulations indicating that increasing catecholaminergic gain should reduce the angular dispersion, and thus increase the precision, of object representations from the same category, as well as reduce the angular dispersion of object representations from distinct categories when distinct-category representations overlap. In Study 1 ($N = 24$), we show that angular dispersion covaries with pupil diameter, an index of baseline catecholamine levels. In Study 2 ($N = 24$), we manipulate catecholamine levels and neural gain using the norepinephrine transporter blocker atomoxetine and demonstrate consistent, causal effects on angular dispersion and brain-wide functional connectivity. Despite the use of very different methods of examining the effect of baseline catecholamine levels, our results show a striking convergence and demonstrate that catecholamines increase the precision of neural representations.

Key words: catecholamine; dopamine; fMRI; norepinephrine; perception; psychopharmacology

Significance Statement

Norepinephrine and dopamine are among the most widely distributed and ubiquitous neuromodulators in the mammalian brain and have a profound and pervasive impact on cognition. Baseline catecholamine levels tend to increase with increasing task engagement in tasks involving perceptual decisions, yet there is currently no direct evidence of the specific impact of these increases in catecholamine levels on perceptual encoding. Our results fill this void by showing that catecholamines enhance the precision of encoding cortical object representations, and by suggesting that this effect is mediated by increases in neural gain, thus offering a mechanistic account of our key finding.

Introduction

In vivo and *in vitro* cell recordings indicate that the neuromodulators norepinephrine and dopamine increase the responsivity of neurons in target brain areas to both excitatory and inhibitory inputs (Berridge and Waterhouse, 2003; Winterer and Weinberger, 2004). An influential theory of catecholamine function proposes that this property of norepinephrine and dopamine allows them to perform the key computational function of gain regulation, thus tuning neural network dynamics to optimize processing in a given context (Servan-Schreiber et al., 1990; Aston-Jones and Cohen, 2005). Several theoretical analyses have

Received Sept. 16, 2015; revised March 17, 2016; accepted April 6, 2016.

Author contributions: C.M.W., J.D.C., and S.N. designed research; C.M.W., E.E., R.L.v.d.B., and K.-D.T. performed research; C.M.W., E.E., and J.A.B. analyzed data; C.M.W., E.E., R.L.v.d.B., K.-D.T., N.J.v.d.W., E.J.G., M.S.v.N., J.A.B., R.C.W., J.D.C., and S.N. wrote the paper.

This work was supported by a Consolidator Grant of the European Research Council. We thank P. Murphy and L. Nystrom for helpful discussion.

The authors declare no competing financial interests.

Correspondence should be addressed to Christopher M. Warren, Dept. of Psychology, Leiden University, Wassenaarseweg 52, 2333 AK Leiden, The Netherlands. E-mail: c.m.warren@fsw.leidenuniv.nl.

DOI:10.1523/JNEUROSCI.3475-15.2016

 0270-6474/16/365699-10$15.00/0

suggested that increases in catecholamine-mediated gain should enhance the quality of neural representations. For example, increasing the gain of all units in a neural network results in improved signal-detection performance of the network (Servan-Schreiber et al., 1990). Likewise, Bayesian theories of brain function indicate that increasing neural gain should narrow the posterior distribution representing the brain's perceptual interpretation of incoming information (Knill and Pouget, 2004; Ma et al., 2006; Friston, 2009). However, direct empirical evidence for an effect of catecholamine-mediated gain on the precision of neural representations in human subjects has never been reported.

We tested the relationship between pupil-linked and pharmacologically induced changes in catecholamine-mediated gain and precision using a powerful, multivariate metric of representational similarity (Schurger et al., 2010, 2015) to quantify the precision of cortical object representations in human neuroimaging data. In two studies we used functional MRI (fMRI) to measure brain activity while participants categorized monochromatic, isoluminant pictures of houses, faces, and scissors. For each stimulus, we determined the spatial pattern of neural activity in the ventral temporal cortex, where category-specific information is represented (Haxby et al., 2001). Treating these patterns of activation as vectors projecting into representational space enabled us to quantify the variability in neural representation across trials as the angular dispersion of the vectors (the degree to which they point in different directions). This measure (also referred to as angular deviation) is optimal for characterizing the spread of points in representational space (Fisher et al., 1987) and controls for global activation differences (Schurger et al., 2010, 2015). We focused on the effect of neural gain on within-category angular dispersion, measuring the precision of patterns of neural activity across different exemplars and instances of the same category. However, we anticipated that catecholamine-mediated changes in neural gain might also have an impact on the relationship between patterns of activity across distinct categories of information. In general, we hypothesized that increased gain would be associated with lower angular dispersion, and thus enhanced precision, of object representations.

To formalize our hypothesis and generate clear-cut predictions, we used a simple neural network model to determine how same-category and distinct-category angular dispersion should change with increasing gain. Below, we first report the model simulation results. Then we present empirical results from 48 participants across two studies, demonstrating that the precision of object representations in the ventral temporal cortex is directly related to catecholamine-mediated changes in gain as assessed by pupillometry (Aston-Jones and Cohen, 2005; Eldar et al., 2013; Murphy et al., 2014; Varazzani et al., 2015; Joshi et al., 2016), and as manipulated pharmacologically using the norepinephrine transporter blocker atomoxetine (Bymaster et al., 2002; Swanson et al., 2006; Koda et al., 2010; Kielbasa et al., 2015).

Pupil diameter and atomoxetine are typically associated with noradrenergic activity. However, there is a wealth of evidence that throughout the cortex noradrenergic terminals corelease dopamine, which acts on dopamine receptors that are also spread throughout the cortex (Devoto and Flore, 2006). Thus, pupil diameter may also index dopaminergic activity. Furthermore, in cortical areas, the norepinephrine transporter is responsible for the reuptake not only of norepinephrine but also of dopamine (Devoto and Flore, 2006). Thus, by blocking norepinephrine, norepinephrine transporter blockers must increase both norepinephrine and dopamine levels throughout the cortex, a prediction borne out by available data (Bymaster et al., 2002; Devoto et al., 2004).

Materials and Methods

Simulation. To simulate the effect of neural gain on the representation of stimuli that belong to different categories, we modeled a 100-unit neural network, set up to represent two categories of stimuli, and examined its response to individual exemplars from each category. Network units were fully interconnected, and the activation of each unit was computed as follows:

$$a_i = \tanh\left(gain \sum_j w_{i,j} a_j + \epsilon\right)$$

where $w_{i,j}$ is the connection weight between unit i and unit j, ε is Gaussian noise with SD 5, tanh describes the sigmoidal-shaped activation function relating the unit's net input to its output, and the gain parameter influences the slope of the activation function. Network weights were set so as to create two attractors m^1 and m^2, each representing a category of stimuli, as follows:

$$w_{i,j} = \frac{m_i^1 m_j^1 + m_i^2 m_j^2}{500000}$$

Setting the weights thus, in a manner that resembles Hebbian learning (Hebb, 1949), ensures that the energy function of the network is minimized by the attractor states, and thus the network naturally tends to return to these states (Hopfield, 1982, 1984). All self-connections ($w_{i,j}$) were set to 0. Both m^1 and m^2 were constructed by setting 50 units to 1 and 50 units to -1, and the similarity between m^1 and m^2 was manipulated along the continuum between completely identical to completely orthogonal.

Input represented individual exemplars of each category. Thus, 1000 input patterns were created for each category by taking either m^1 or m^2 and redrawing five randomly selected units. On each trial, unit activations were initialized to one of the input patterns and then all activations were updated asynchronously in a random order for five iterations. We then measured the angular dispersion of the end-state activations between pairs of trials in which inputs were of the same category (i.e., inputs were variants of the same attractor), as well between pairs of trials in which inputs were of distinct categories (i.e., inputs were variants of different attractors), as a function of gain and the level of overlap between categories (i.e., the similarity between the two attractors).

Participants. Twenty-seven healthy volunteers (19 females; mean ± SD: age, 22.1 ± 2.4 years; range, 18–28 years) were recruited for the pupil study. Three participants were excluded because half or more of the 14 blocks yielded insufficient artifact-free pupil data. Twenty-five healthy volunteers (15 females; mean ± SD: age, 22.0 ± 1.7 years; range, 19–26 years) were recruited for the atomoxetine study. One participant experienced nausea beginning ~1 h and 15 min after drug intake. The participant was able to complete the task, but the data were nonetheless discarded because of concerns that the nausea had influenced the results. Participants in the atomoxetine study were screened by a medical doctor for physical health, drug contraindications, standard contraindications for MRI research, and other exclusion criteria, including the following: current use of prescription medication, a history of psychiatric illness, cardiovascular disease, renal failure, hepatic insufficiency, glaucoma, head trauma, hypertension, and drug or alcohol abuse. Participants with learning disabilities, poor eyesight (severe myopia of −6 diopters or worse), who smoked >5 cigarettes a day, who were pregnant, or who were left-handed were also excluded. Participants who took part in the pupil study were paid €30. Participants who took part in the atomoxetine study were paid €135. All participants gave informed consent. Both studies were approved by the Leiden University medical ethics committee.

Procedure. Participants categorized pictures of houses, faces, and scissors, taken from the set used by Haxby and colleagues (2001). All stimuli were monochromatic, 200 × 200 pixels in size, and presented on a gray

background. Stimuli images, fixation cross, response screen, and feedback images were processed by a Matlab toolbox (Mathworks; Willenbockel et al., 2010) to make them isoluminant to each other. Stimuli were presented with the lights off, at an average rate of one per 8 s and consisted of five exemplars from each category, pulled randomly for each participant from a larger group of 12 exemplars each. In Study 2, five new exemplars of each category were randomly selected without replacement for the second session so that no exemplar was repeated between sessions. A trial began with a fixation cross presented for a duration that varied between 4000 and 6000 ms, followed by the picture for 1500 ms, and then a response screen for 1500 ms. If participants made an incorrect response or failed to respond, feedback appeared for 500 ms to indicate the error. When feedback was presented, fixation time was cut by 500 ms so as not to disrupt the rate of stimulus presentation. Participants were told to answer with their right hand using a button box strapped to their leg, and to respond only during the response screen, not before or after. The mapping of buttons to categories was shown on the response screen and was consistent throughout the entire experiment, but counterbalanced across participants. Each block consisted of 15 trials, with all stimuli presented once in random order. There were two runs of seven blocks each in the pupil study, and two runs of five blocks each in the atomoxetine study. There was no break between blocks, but between runs there was a short break, during which researchers briefly queried the participant on their well-being and task engagement.

The atomoxetine study followed a randomized, double-blind design. Atomoxetine and placebo were administered in separate test sessions, spaced 1 week apart, scheduled at approximately the same time of day. Atomoxetine was administered orally as a single, encapsulated, 40 mg pill, with a glass of water. A starting dose of 40 mg is used in clinical practice, and is a dose associated with limited side effects (Heil et al., 2002). Aside from its effect on catecholamine levels, atomoxetine can also act as an NMDA receptor blocker (Ludolph et al., 2010). The placebo consisted of 125 mg of lactose monohydrate with 1% magnesium stearate, encapsulated, and identical in appearance to the drug. The pill was administered ~1 h and 40 min before the participants were put into the scanner, and 1 h and 55 min before the first block of trials began.

Salivary cortisol is a reliable index of arousal/stress (Buchanan et al., 1999; Bosch et al., 2009). Catecholaminergic drug effects are known to interact with baseline levels of arousal (Coull, 2001; Sikström and Söderlund, 2007). With this in mind, we obtained saliva samples four times per session, beginning immediately before taking the pill ($t = 0$), at $t = 80$, at $t = 140$, and at $t = 220$ min after pill ingestion.

fMRI data acquisition and preprocessing. All fMRI scans were acquired using a Philips Achieva 3 tesla scanner at the Leiden University Medical Center in Leiden, The Netherlands. An EPI scan with 38 transverse slices covering the whole brain was employed in descending order, with a TR of 2.2 s, TE of 30 ms, a flip angle of 80°, a 220 mm field of view, and a voxel resolution of 2.75 × 2.75 × 2.75 mm (+10% slice gap). The first five volumes were discarded to remove magnetic disequilibrium effects. Structural images were acquired using a T1-weighted sequence with a TR of 9.8 ms, a TE of 4.60 ms, a flip angle of 8°, a 224 mm field of view, and a voxel resolution of 0.875 × 0.875 × 1.2 mm. All data were processed using Matlab (Mathworks) and SPM8 (Wellcome Trust Centre for Neuroimaging, University College London). EPI images were realigned to the first image of the session (no participant exhibited motion of >2 mm), and slice time was corrected to the center slice. Brain images were coregistered with the mean functional image, segmented into gray matter, white matter, and CSF, and then normalized to Montreal Neurological Institute (MNI) space. Functional images were then normalized using the same parameters as the structural images. No spatial smoothing was applied.

Heart and breathing rate. In Study 2, because of the known effect of atomoxetine on heart rate, we monitored heart rate and breathing rate using a pulse oximeter and a respiration belt provided as accessories with the MRI scanner. Physiological data were exported to Matlab (Mathworks) and down-sampled from 500 to 100 Hz. The pulse oximetry time series was smoothed using a finite impulse response bandpass filter (0.6–2.0 Hz), and the respiration time series was low-pass filtered using a 1 Hz cutoff. Peaks in each time series, which correspond to maximal expansions of the lungs in the respiration time series, and maximal local blood oxygenation in the pulse time series were detected using the peakdet function with a delta value of half the SD of the respective time series. Breathing and heart rates were both quantified as peaks per minute.

Because the results do not directly address our key questions, we report the results here. We ran ANOVAs of mean breathing rate and mean heart rate recorded over the course of the two scan sessions. The breathing rate (breaths per minute) for the atomoxetine group ($M = 32.1$) and placebo group ($M = 31.0$) was nearly the same ($F_{(1,22)} = 0.46$, $p = 0.50$). There was no interaction between treatment and treatment order ($F_{(1,22)} = 3.04$, $p = 0.10$). Heart rate (beats per minute) was somewhat higher in the atomoxetine group ($M = 66.1$) than in the placebo group ($M = 62.5$), but this effect was not robust ($F_{(1,22)} = 3.36$, $p = 0.08$). There was no interaction between treatment and treatment order ($p = 0.91$).

Saliva sample collection and processing. Whole saliva was collected using the spitting method, which involves allowing participants to let saliva collect passively on the floor of their mouth over a 3 min period, interrupted by spitting the accumulated saliva into a polypropylene tube every minute (Navazesh, 1993). Samples were frozen at −60°C until processing. Samples were thawed and homogenized using a Vortex mixer, and centrifuged at room temperature for 4 min at 4000 × *g*. The clear supernatant was assayed using a competitive enzyme-linked immunosorbent, according to the manufacturer's instructions (IBL International), with a sensitivity of 0.45 ng/ml, and an intra-assay variability of 2.2%. All samples of the same participant were run simultaneously.

Pupil diameter recording and preprocessing. Pupil diameter was recorded using an MRI-compatible Eyelink 1000 eye tracker (SR Research) at a sampling rate of 500 Hz. Samples with artifacts were identified using a slope criterion, whereby any sample-to-sample change of >150 pixels was eliminated, as were periods where tracking of the pupil was briefly lost (due to blinks or large eye movements). Pupil diameter was assessed in pixels, during the 2 s window of time preceding stimulus onset (giving ≥5 s for the phasic dilation from the previous trial to subside). Only trials for which at least half the assessment epoch contained no artifacts were included. A block was only included if at least half of the trials (≥8 of 15) yielded sufficient, artifact-free pupil data (Eldar et al., 2013). For included participants, on average 1.1 of the 14 blocks were excluded from analysis (8%; $SD = 13\%$; range: 0% excluded for 14 of 24 participants; 0–29% for 22 of 24 participants; 47% for 2 of 24 participants). For included blocks, an average of 1.6 trials per block were excluded (10.1%; $SD = 0.3\%$; range: 2–9% excluded for 6 of 24 participants; 9–12% for 12 of 24 participants; 12–17% for 6 of 24 participants).

Angular dispersion analysis. Patterns were taken from volumes at the peak of the stimulus-driven hemodynamic response (~6 after stimulus onset) and, using MNI masks included with the MarsBaR toolbox (Brett et al., 2002), restricted to the ventral temporal cortex (Haxby et al., 2001). Patterns representative of trial-baseline activity were taken from volumes at stimulus onset. The masks employed included the left and right inferior temporal cortex, the left and right middle temporal cortex, and the left and right fusiform gyrus. This resulted in vectors with 6695 voxels. The time series for each voxel within each run were *z*-scored separately. To have masks that were the exact same in all conditions, across all sessions and participants, no dimensionality reduction was done. The drawback of this decision is that including uninformative voxels resulted in mean angular dispersion values very close to perfect decorrelation. Nonetheless, the metric was able to distinguish between same-category and distinct-category comparisons, and was sensitive to gain in two studies. Furthermore, very similar effects were obtained using different methods of determining the mask (see Results).

Angular dispersion was calculated between fMRI volumes in the manner described by Schurger and colleagues (2010, 2015). This method allows comparisons of the quality of neural representations, independent of differences in the strength of activation. Angular dispersion reflects trial-to-trial variability in neural representation by treating each pattern as a vector projecting into a multidimensional state space defined by the number of voxels in the patterns: each voxel activation gives the location in the state space along that dimension. The length (norm) of the vector projecting to the point defined by all voxel coordinates corresponds to the strength of activation across all voxels, and the angle between two

activation vectors reflects the difference between the two patterns of representation. Angular dispersion is calculated as the length of the normalized vector sum, divided by the number of vectors summed. Following Schurger and colleagues (2015), we present angular dispersion as the inverse of this value so that lower values indicate lower angular dispersion, and greater precision. We calculated angular dispersion for both same-category and distinct-category comparisons, one pair of vectors at a time, and then took the mean of all possible same-category comparisons, and the mean of all possible distinct-category comparisons. Outlier values further than 2.5 SDs from the mean were not included in the calculation of each mean, leading to the exclusion of 1.54% of data in Study 1, and 1.49% of data in Study 2. Calculating angular dispersion one pair of vectors at a time has two benefits. First, the inverse of angular dispersion between two vectors yields values that can be interpreted in terms of the relative direction of the two vectors. That is, a value of 1 indicates two vectors pointing in the same direction, a value of infinity indicates two vectors pointing in opposite directions, and a value of 1.4142 (square root of 2) indicates two vectors perpendicular to each other. Second, angular dispersion increases as the number of vectors in the calculation increases, such that angular dispersion calculated for all same-category trials at once could not be validly compared with angular dispersion for the mean of all distinct-category comparisons, which must be calculated one pair of vectors at a time. Based on our clear a priori predictions, one-tailed *t* tests were employed in testing for effects of pupil size and drug on angular dispersion.

Functional connectivity analysis. The functional connectivity analysis was implemented as described by Eldar and colleagues (2013). A voxel-to-voxel correlation matrix was computed across the time series for all gray matter voxels (ranging from 32,810 to 40,190 voxels across participants). Functional connectivity was calculated as the mean of all absolute (i.e., unsigned) correlations in the matrix. To minimize spurious correlations, movement parameters from the realignment, as well as mean activation from the normalized gray, white, and CSF matter masks, were regressed out of the functional data, and low-frequency drifts were removed using a high-pass filter with a 0.0078 Hz cutoff. The normalized gray matter mask was used to restrict analysis of functional connectivity to gray matter. The calculations were done across the entire 28 min of scan time for each session, and the mean absolute correlation across the time series was compared between sessions. The clustering coefficient (the degree to which functional connections were clustered) was calculated using graph-theoretic analysis (Eguíluzet al., 2005) for the same time periods as described for the functional connectivity analysis. Voxels were defined as connected if the absolute correlation between them was in the top 1% of all unsigned correlations calculated in the correlation matrix. A triplet of connected voxels was defined as any three voxels connected in such a way that one voxel was connected to the other two ($\geq$2 connections). A closed triplet was defined as a triplet in which each voxel was connected to both the other voxels in the triplet (three connections). The clustering coefficient was defined as the number of closed triplets divided by the number of total triplets (Luce and Perry, 1949). The statistical test of a difference in clustering coefficients between drug and placebo in Study 2 was corrected for unequal variances.

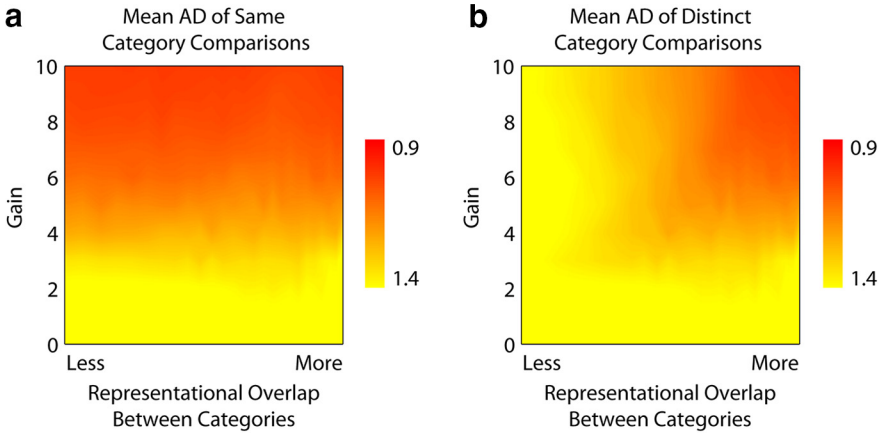

Figure 1. Simulation results. ***a***, For all levels of representational overlap, same-category angular dispersion decreases (i.e., becomes more precise) with increasing gain. ***b***, In contrast, decreases in distinct-category angular dispersion due to increasing gain depend upon the level of representational overlap. In both panels, hotter colors indicate decreased angular dispersion.

Results

Simulation: increased gain lowers angular dispersion for both same-category and distinct-category comparisons

Model simulations allowed us to formally test our prediction that increased gain should enhance same-category representational similarity (i.e., precision) and that angular dispersion captures this effect. In addition, the simulations allowed us to examine the effect of gain on distinct-category angular dispersion.

We modeled the effect of gain and of category overlap on angular dispersion using a standard neural network model designed to accept variable input and settle into a network state with a pattern of activation that could be compared with other end states using the angular dispersion metric. The network weights were set to simulate two category attractors with varying levels of pattern overlap (categories with many shared features vs categories with few shared features; see Materials and Methods). Simulations revealed that increasing gain lowered same-category angular dispersion, consistent with other computational models that predict that increasing gain should increase precision (Servan-Schreiber et al., 1990; Knill and Pouget, 2004; Fig. 1*a*). We also computed the average angular dispersion associated with all possible comparisons of the activation patterns of exemplars from distinct categories. Object category representations are distributed and overlapping within the ventral temporal cortex (Haxby et al., 2001). To the extent that there is overlap between neural representations of distinct categories, a general increase in the precision of representations should also manifest in increased representational similarity between distinct categories. Indeed, this was confirmed by our simulations: distinct-category angular dispersion decreased with increasing gain for all levels of representational overlap between the categories except when the two attractors were fully decorrelated (Fig. 1*b*). Thus, the simulations suggest that increases in catecholamine-mediated gain should result in lower angular dispersion for both same-category and distinct-category comparisons.

Study 1: pupil diameter covaried with the precision of cortical representations

In Study 1 (n = 24 participants) we measured baseline pupil size, which is correlated with tonic levels of norepinephrine activity (Aston-Jones and Cohen, 2005; Nieuwenhuis et al., 2011; Murphy et al., 2014; Varazzani et al., 2015; Joshi et al., 2016), and therefore has been hypothesized to track neural gain (Aston-Jones and Cohen, 2005; Gilzenrat et al., 2010; Nassar et al., 2012; Eldar et al., 2013; Cheadle et al., 2014; McGinley et al., 2015). We examined whether precision of neural representations differed between miniblocks of trials associated with relatively small and large baseline pupil sizes. As predicted by our model, angular dispersion was lower for

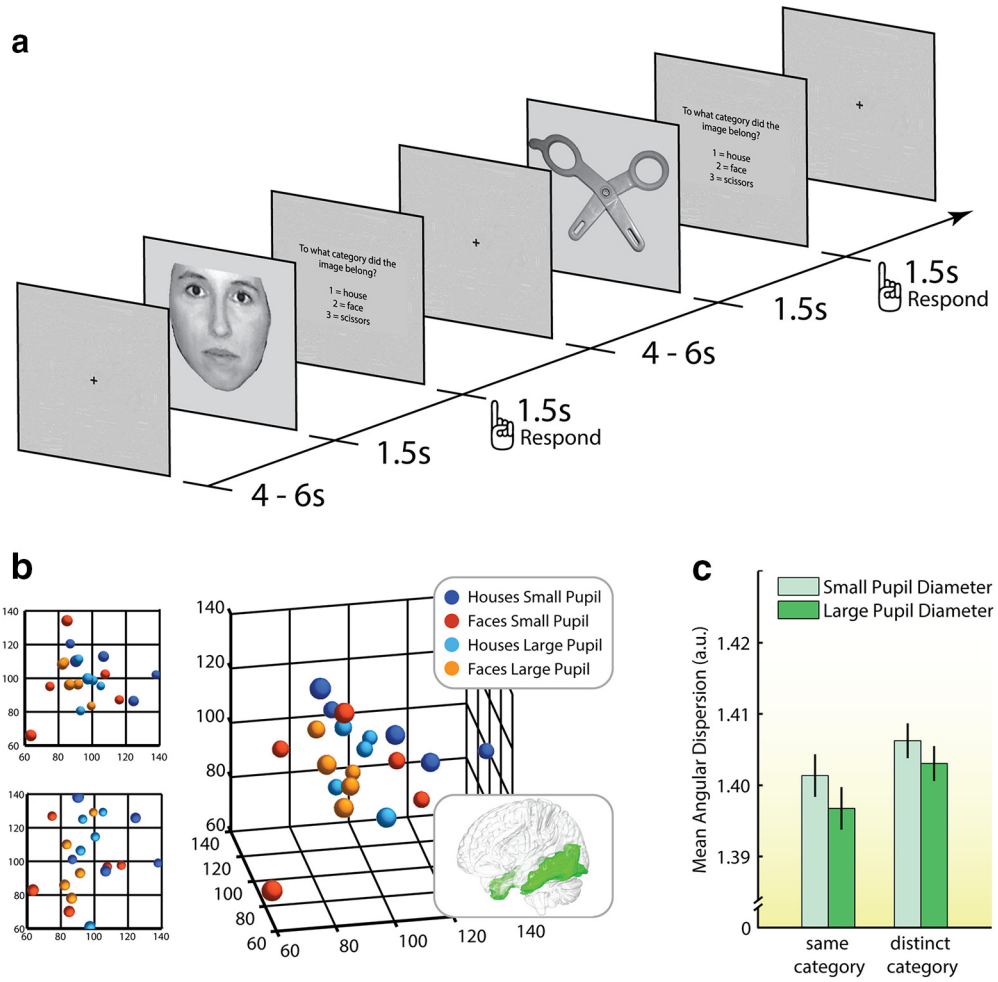

Figure 2. Trial procedure and empirical results in Study 1. ***a***, Participants categorized three categories of stimuli with a jittered interstimulus interval of 7–9 s. ***b***, Representative sample of data points (colored spheres) with coordinates defined by voxel-activation patterns, scaled down to three dimensions using nonmetric multidimensional scaling. Units on each axis are arbitrary: the three axes define a three-dimensional, representational space. Distance "into" the page is represented with smaller spheres. Angular dispersion characterizes the degree to which points are tightly clustered together (more precise) or loosely spread out (less precise) in the representational space. For both house and face stimuli, the points are closer together in the large-pupil blocks, and across categories large-pupil points are closer together than small-pupil points. These data points were chosen to best illustrate the overall pattern of results. The smaller plots on the left show two-dimensional views of the same data. ***c***, Larger baseline pupil size was associated with lower angular dispersion (higher precision) for both same-category and distinct-category comparisons. Error bars reflect SEM. The glass-brain inset shows the mask of ventral temporal cortex that was employed for all analyses.

large-pupil blocks, which is consistent with the hypothesis that higher gain was associated with increased precision of category representations ($F_{(1,23)} = 6.29$, $p = 0.020$; Fig. 2*b*,*c*). There was no evidence that this effect was different for same-category versus distinct-category comparisons ($F_{(1,23)} = 0.55$, $p = 0.47$; Fig. 2*c*). Indeed, paired *t* tests revealed that large-pupil blocks were characterized by significantly lower angular dispersion for both same-category ($t_{(23)} = 2.20$, $p = 0.019$) and distinct-category comparisons ($t_{(23)} = 2.14$, $p = 0.022$).

According to our simulations, this suggests that the stimulus categories had some degree of representational overlap in the temporal cortex, and that overlapping portions, or features shared between categories, were represented more precisely in the large-pupil blocks. The notion that the effects on same-category and distinct-category angular dispersion reflect the same influence of catecholamine-mediated gain is corroborated by our finding that participants with a larger decrease in same-category angular dispersion also showed a larger decrease in distinct-category angular dispersion ($r = 0.49$, $p = 0.015$).

Additional analyses showed that the effects of pupil-linked gain on angular dispersion were specific for the time period associated with stimulus processing (peak of the hemodynamic response); similar effects were not observed at stimulus onset (i.e., trial baseline), before the stimulus-elicited hemodynamic response had a chance to develop ($F_{(1,23)} = 0.15$, $p = 0.70$). Furthermore, condition-averaged angular dispersion was lower during the hemodynamic response peak than during trial baseline ($M_{peak} = 1.402$, $M_{baseline} = 1.413$, $t_{(23)} = 6.38$, $p < 0.001$), indicating that patterns of activity across ventral temporal cortex became more consistent during stimulus processing than at rest between trials (Schurger et al., 2015). Finally, same-category comparisons were associated with lower angular dispersion than distinct-category comparisons ($M_{same} = 1.399$, $M_{distinct} = 1.405$, $F_{(1,23)} = 24.98$, $p < 0.001$; Fig. 2*b*), as expected given that same-category exemplars should be coded as more similar to each other than exemplars of other categories, and consistent with seminal work on the subject (Haxby et al., 2001).

Study 1: exclusion of alternative interpretations; small-pupil and large-pupil blocks did not differ in gaze position, head movements, time on task, signal intensity, and task performance

It is possible that small-pupil blocks were associated with less precise category representations because participants looked away from the stimulus more often than in large-pupil blocks. To exclude this possibility, we calculated the variance in gaze position for the 1500 ms epoch of samples after stimulus onset, when gaze position/movement could affect the quality of visual input. Paired-samples *t* tests comparing the natural log of the variance in gaze position (distributions were skewed toward high values) along both the vertical and (separately) the horizontal axis yielded no significant differences between small-pupil and large-pupil blocks [horizontal: $M_{small} = 5.42$ ($SD = 1.03$), $M_{large} = 5.46$ ($SD = 0.88$), $t_{(23)} = 0.52$, $p = 0.61$; vertical: $M_{small} = 5.69$ ($SD = 1.16$), $M_{large} = 5.61$ ($SD = 0.98$), $t_{(23)} = 0.80$, $p = 0.43$]. An analysis of absolute deviation from fixation on each trial yielded similar results. Furthermore, small-pupil and large-pupil blocks did not differ in the amount of head movement [$M_{small} = 0.025$ ($SD = 0.001$), $M_{large} = 0.025$ ($SD = 0.001$), $F_{(1,23)} = 0.03$, $p = 0.85$].

We also considered the possibility that our comparison of small-pupil and large-pupil blocks was confounded with time on task. More specifically, small-pupil blocks may, on average, have occurred later in the experiment, when participants were perhaps less engaged in the task, resulting in impaired stimulus encoding and higher angular dispersion. To test for this potential confound, we quantified the timing of small-pupil and large-pupil blocks in terms of block number (1–14; i.e., the time since the start of the session), and compared the distributions across participants. This analysis did not support a time-on-task explanation: large-pupil blocks (mean block number, 7.13; $SD = 2.06$) did not differ from the small-pupil blocks (mean, 7.88; $SD = 2.08$; $t_{(23)} = 0.89$, $p = 0.38$).

The vector lengths associated with each stimulus-elicited pattern of hemodynamic activation reflect the intensity of brain activation: the responsivity of voxels to stimulus presentation relative to mean activation (Schurger et al., 2010, 2015). In the case of relatively short vector norms—when some vector coordinates are very close to zero—small perturbations due to noise can have large effects on the direction of the vector, increasing angular dispersion. Thus, it is important to test whether the small-pupil and large-pupil blocks differed in vector norm. A 2 × 3 ANOVA with pupil size and stimulus category (houses, faces, scissors) as repeated-measures factors revealed no effect of pupil size [$M_{small} = 78.4$ ($SD = 3.2$), $M_{large} = 78.8$ ($SD = 3.2$), $F_{(1,23)} = 1.58$, $p = 0.22$], and stimulus category ($F_{(2,46)} = 0.57$, $p = 0.57$), nor an interaction of pupil size with stimulus category ($F_{(2,46)} = 0.55$, $p = 0.58$). This indicates that the effect of pupil size on precision was not due to an increased susceptibility to noise in the small-pupil blocks.

Finally, small-pupil and large-pupil blocks did not differ in task performance—an anticipated result due to ceiling effects for accuracy and only slight time pressure to respond within 1500 ms: mean categorization accuracy was 98% ($SD = 3\%$) for small-pupil blocks, and also 98% ($SD = 4\%$) for large-pupil blocks ($t_{(23)} = 1.02$, $p = 0.32$). Likewise, mean response time did not differ with pupil size [$M_{small} = 607$ ms (SD = 122 ms), $M_{\text{large}} = 603$ ms (SD = 123 ms), $t_{(23)} = 0.38$, $p = 0.71$]. This indicates that the enhanced precision of category representations in the large-pupil blocks cannot be explained in terms of differences in task performance. The absence of systematic performance differences is important; if the task had been more difficult and had required perceptual learning, gain effects on perceptual learning (which would have manifested in performance differences) could have confounded gain effects on the precision of category representations.

Study 1: relationship between pupil diameter and network gain, as indicated by whole-brain functional connectivity

A global increase in gain increases the activity of neuron populations receiving excitatory input and decreases the activity of neuron populations receiving inhibitory input. This type of signal amplification yields a system-wide facilitation of signal transmission (Aston-Jones and Cohen, 2005). Recent computational work (Eldar et al., 2013) suggests that pupil-linked increases in neuromodulatory gain boost all temporal correlations between the activities of local neuron populations, regardless of whether these are positive or negative—that is, higher gain increases (absolute) functional connectivity throughout the network. Note that although both functional connectivity and angular dispersion are measures of similarity, these measures are, in an important way, orthogonal to each other: functional connectivity involves correlating time series (or time vectors) across voxels, whereas angular dispersion is akin to correlating voxel vectors across time points. Another effect of increased gain, confirmed by model simulations and empirical data, is that network connectivity becomes more tightly clustered (Eldar et al., 2013).

We computed these two network measures of neural gain to examine whether they would provide converging evidence for the assumption that large-pupil and small-pupil blocks differed in the level of gain. In line with Eldar and colleagues (2013), we found that both functional connectivity ($M_{large} = 0.132$ vs $M_{small} = 0.130$) and clustering coefficient ($M_{large} = 0.118$ vs $M_{small} = 0.110$) were greater in large-pupil blocks than in small-pupil blocks. A multivariate ANOVA (MANOVA) with functional connectivity and clustering coefficient as dependent variables and pupil size as a repeated-measures factor revealed a marginally significant effect of pupil size (Wilk's $\lambda = 0.79$, $F_{(2,23)} = 3.14$, $p = 0.06$). The mean variance across the time series for all voxels was not significantly different between atomoxetine and placebo ($t_{(23)} = 0.70$, $p = 0.49$), which rules out the possibility that the difference in functional connectivity was due to a difference in variance. Thus, although the difference between large-pupil and small-pupil blocks just failed to reach significance at the 0.05 level, the results of these functional-connectivity analyses are broadly consistent with the findings reported by Eldar and colleagues (2013), and with our hypothesis that the observed enhancement in precision of cortical representations was produced by an increase in catecholamine-mediated network gain.

Study 2: atomoxetine administration enhanced the precision of cortical representations in Session 1 but not in Session 2

In Study 2 ($n = 24$ participants) we sought to test the causal influence of catecholamine-mediated gain on representational precision. We did so by administering a single dose of 40 mg of atomoxetine, a norepinephrine transporter blocker that increases extracellular levels of norepinephrine (Bymaster et al., 2002; Swanson et al., 2006; Koda et al., 2010; Kielbasa et al., 2015) and (probably) dopamine (Bymaster et al., 2002; Devoto et al., 2004; Devoto and Flore, 2006), and examined the effect on precision in a double-blind, placebo-controlled, cross-over design. Aside from the treatment, the only difference between the two sessions was that five new exemplars of each category were randomly selected as stimuli in the categorization task.

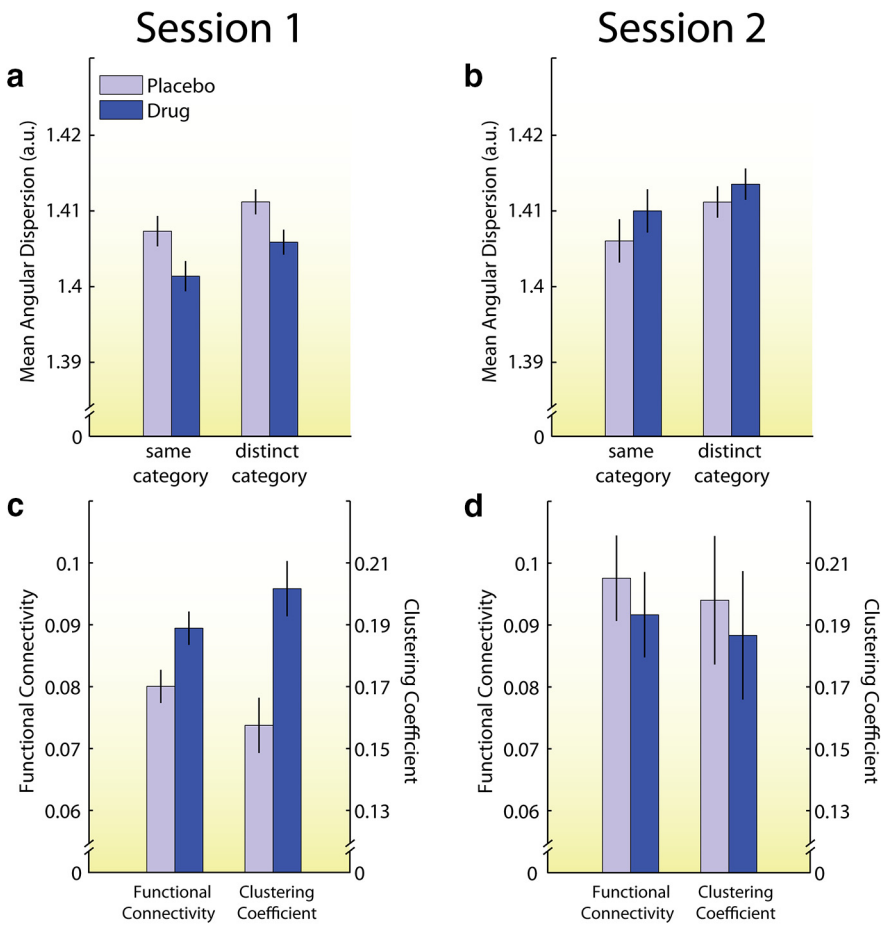

Figure 3. Results from both sessions of Study 2. Error bars reflect SEM. ***a–d***, Measures of precision (***a***, ***b***) and connectivity (***c***, ***d***) were affected by drug in Session 1 (***a***, ***c***) but not in Session 2 (***b***, ***d***). Error bars reflect SEM.

To test the prediction that an atomoxetine-induced increase in neural gain would produce an increase in precision, we analyzed the angular dispersion data from both sessions with a $2 \times 2 \times 2$ mixed ANOVA with treatment (drug vs placebo) and same-category versus distinct-category comparisons as repeated-measures factors, and treatment order (drug in first session vs drug in second session) as between-subjects factor. There was no effect of treatment ($M_{drug} = 1.407$, $M_{placebo} = 1.409$, $F_{(1,22)} = 0.68$, $p = 0.42$). However, there was a significant interaction of treatment with treatment order ($F_{(1,22)} = 6.22$, $p = 0.021$), indicating that atomoxetine had very different effects on angular dispersion in Session 1 and Session 2. This type of treatment-by-treatment order interaction is not uncommon in studies using noradrenergic agents (for review, see Coull et al., 2004), and is thought to reflect the complex interplay between baseline arousal and noradrenergic actions (Coull, 2001). In line with this view, salivary cortisol measures from our participants (reported below) suggested that stress and/or arousal levels differed between the two sessions (Buchanan et al., 1999; Bosch et al., 2009). In accordance with our treatment-by-treatment order interaction, we analyzed each session separately with treatment as a between-subjects factor.

In Session 1, as predicted by our model, atomoxetine was associated with lower angular dispersion, and hence higher precision, than placebo ($F_{(1,22)} = 4.36$, $p = 0.048$). There was no evidence that the effect of atomoxetine differed between same and distinct comparisons ($p = 0.67$). Indeed, atomoxetine significantly lowered angular dispersion for both same-category ($t_{(22)} = 1.96$, $p = 0.031$; Fig. 3*a*) and distinct-category comparisons ($t_{(22)} = 2.10$, $p = 0.024$; Fig. 3*a*). These effects were again specific for the time period associated with stimulus processing; neither effect was significant at trial baseline (both p's > 0.63).

Additional analyses revealed that, as in Study 1, condition-averaged angular dispersion was lower during the hemodynamic response peak than during trial baseline ($M_{peak} = 1.406$, $M_{baseline} = 1.414$, $t_{(23)} = 5.92$, $p < 0.001$), and same-category comparisons were associated with lower angular dispersion than distinct-category comparisons ($M_{same} = 1.404$, $M_{distinct} = 1.408$, $t_{(23)} = 7.98$, $p < 0.001$; Fig. 3*a*).

In Session 2, atomoxetine did not have a significant effect on angular dispersion collapsed across comparison type ($p > 0.30$), nor on either same-category or distinct-category comparisons tested separately (same category: $M_{drug} = 0.1.410$, $M_{plac} = 1.406$, $t_{(22)} = 1.16$, $p = 0.26$; distinct category: $M_{drug} = 1.413$, $M_{plac} = 1.411$, $t_{(22)} = 0.87$, $p = 0.39$; Fig. 3*b*). Other results were similar as in Session 1 and in Study 1: atomoxetine did not affect same-category and distinct-category angular dispersion at trial baseline (both p's > 0.57). Condition-averaged angular dispersion was lower during the hemodynamic response peak than during trial baseline ($M_{peak} = 1.410$, $M_{baseline} = 1.416$, $t_{(23)} = 5.69$, $p < 0.001$), and same-category comparisons were associated with lower angular dispersion than distinct-category comparisons ($M_{same} = 1.408$, $M_{distinct} = 1.412$, $t_{(23)} = 5.67$, $p < 0.001$; Fig. 3*b*).

Study 2: atomoxetine administration increased network gain in Session 1 but not in Session 2, as indicated by whole-brain functional connectivity

As noted above, pharmacological effects of catecholaminergic drugs on brain and behavior are often dose, context, and group specific (Berridge and Waterhouse, 2003; Cools and Robbins, 2004; Coull et al., 2004). Therefore, we asked whether the difference between sessions in angular dispersion effects might be associated with differential effects of atomoxetine on neural gain. Accordingly, we tested whether atomoxetine increased the strength and clustering of functional connectivity in both Sessions 1 and 2.

A MANOVA with functional connectivity and clustering coefficient as dependent variables, treatment as a repeated-measures factor, and treatment order as a between-subjects factor revealed a significant interaction between treatment and treatment order (Wilk's $\lambda = 0.71$, $F_{(2,21)} = 4.32$, $p = 0.027$). The main effect of treatment was not significant (Wilk's $\lambda = 0.93$, $F_{(2,21)} = 0.81$, $p = 0.46$). Follow-up MANOVAs, separately for each session, indicated that in Session 1 atomoxetine significantly increased functional connectivity ($M_{drug} = 0.089$ vs $M_{plac} = 0.080$) and the clustering coefficient ($M_{drug} = 0.202$ vs $M_{plac} = 0.158$), as

indicated by a main effect of treatment (Wilk's $\lambda = 0.74$, $F_{(2,21)} = 3.67$, $p = 0.043$; Fig. 3*c*). The mean variance across the time series for all voxels was not significantly different between atomoxetine and placebo ($t_{(22)} = 0.80$, $p = 0.43$), which rules out the possibility that the difference in functional connectivity was due to a difference in variance. These results suggest that the drug-related enhancement in precision in Session 1 was accompanied by an increase in catecholamine-mediated network gain.

In contrast to Session 1, in Session 2 atomoxetine did not increase functional connectivity and the clustering coefficient (Wilk's $\lambda = 0.98$, $F_{(2,21)} = 0.24$, $p = 0.79$; Fig. 3*d*). Indeed, if anything, atomoxetine slightly reduced both functional connectivity ($M_{drug} = 0.092$ vs $M_{plac} = 0.908$) and the clustering coefficient ($M_{drug} = 0.187$ vs $M_{plac} = 0.198$), although these reductions were not significant. As in Session 1, the mean variance across the time series for all voxels was not significantly affected by treatment ($t_{(22)} = 0.65$, $p = 0.53$). As in the work of Eldar and colleagues (2013), functional connectivity and clustering coefficient were correlated across participants: the mean Pearson *r*, averaged across the two sessions, was 0.78, an acceptable level of multicollinearity. These functional connectivity results suggest that atomoxetine increased neural gain in Session 1, but had no effect on gain in Session 2 (Eldar et al., 2013). This pattern of results aligns with computational models and theoretical frameworks that posit a direct relationship between neural gain and precision (Servan-Schreiber et al., 1990; Knill and Pouget, 2004; Ma et al., 2006; Friston, 2009), and thus provides an explanation for why atomoxetine decreased angular dispersion only in Session 1.

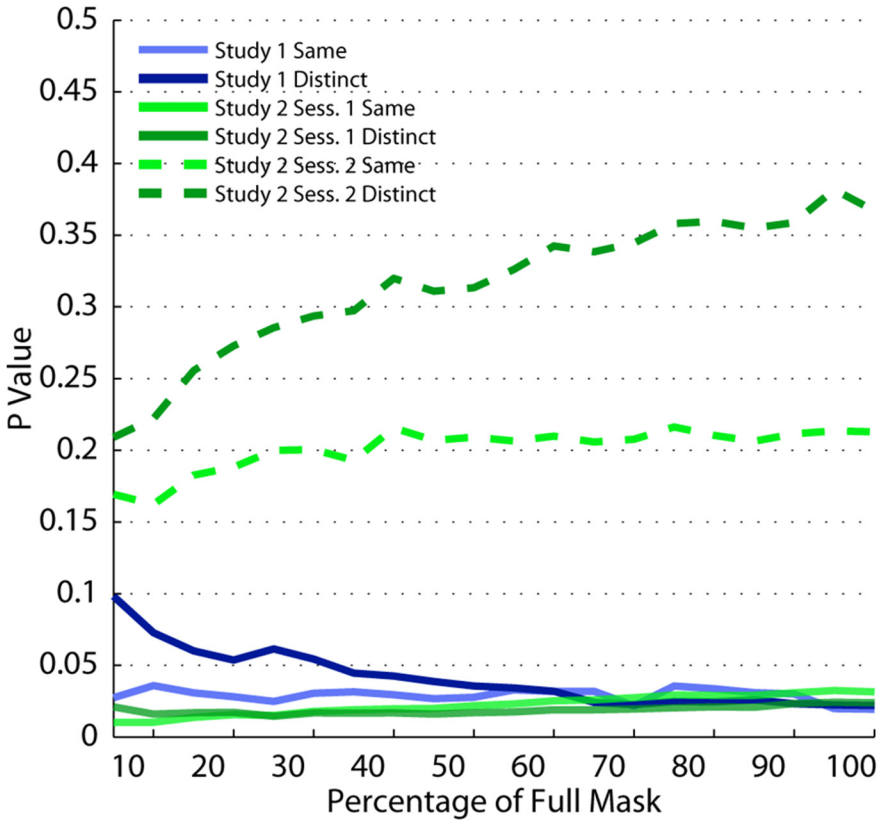

Figure 4. Robustness of angular dispersion effect to voxel selection. Selecting only voxels within the ventral temporal cortex that reliably discriminated between conditions (as per *F* tests) did not alter the observed effects. We plot the *p* value from each statistical test (paired *t* tests on effects of pupil size in Study 1, independent-samples *t* test on the effects of treatment in Study 2) as a function of the percentage of the most informative voxels from the full ventral temporal mask.

Study 2: treatment and session effects on salivary cortisol

To examine differences in arousal between the two sessions, we examined salivary cortisol (at four different time points; see Materials and Methods), a reliable and commonly used index of arousal/stress (Buchanan et al., 1999; Bosch et al., 2009). Cortisol values were analyzed using an ANOVA with treatment and time point (1–4) as repeated-measures factors and treatment order as between-subjects factor. In line with previous research (Chamberlain et al., 2007), this analysis showed a main effect of treatment ($F_{(1,21)} = 8.82$, $p = 0.007$) and an interaction between treatment and time point ($F_{(3,63)} = 6.35$, $p = 0.003$), indicating that atomoxetine increased cortisol values relative to placebo, especially at later time points. There was also a main effect of time ($F_{(3,63)} = 7.48$, $p < 0.001$): cortisol levels decreased over the course of each session, indicating that subjects habituated to the test context. Finally, we found a significant treatment-by-treatment order interaction ($F_{(1,21)} = 5.12$, $p = 0.034$), which reflected a main effect of session: cortisol levels were higher in Session 2 than in Session 1 at all time points, but especially so at baseline, although this comparison was only marginally significant ($t_{(23)} = 1.99$, $p = 0.059$).

These results suggest that participants' arousal levels were higher in Session 2, and that this effect was already present before treatment. As noted above, this type of session effect is known to interact with noradrenergic drugs. Although the notion of increased arousal levels in Session 2 seems counterintuitive, the debriefing process after Session 1 suggested a potential explanation. During this debriefing, most participants expressed a belief that they had not received the drug in the first session, a belief that may have caused them to be more anxious during the second session.

Study 2: exclusion of alternative interpretations; atomoxetine and placebo did not differ in head movements, signal intensity and task performance

To check for treatment effects on head movements in Study 2, we ran a 6×2 mixed ANOVA on the mean of the absolute derivative values with the head-movement parameters as the within-subjects factor and treatment (drug vs placebo) as a between-subjects factor. There was no effect of treatment [$M_{drug} = 0.022$ ($SD = 0.005$), $M_{plac} = 0.023$ ($SD = 0.006$), $F_{(1,22)} = 0.03$, $p = 0.86$] and no interaction of treatment with movement parameters ($p = 0.98$).

A $2 \times 3 \times 2$ mixed ANOVA with treatment (drug/placebo) and stimulus category (houses, faces, scissors) as repeated-measures factors and treatment order as a between-subjects factor revealed an effect of stimulus category, indicating that faces ($M = 73.2$, $SD = 6.6$) were associated with shorter vector norms than houses ($M = 77.8$, $SD = 1.4$) and scissors ($M = 77.8$, $SD = 1.7$; $F_{(2,44)} = 10.07$, $p < 0.001$). More importantly, there was no

effect of treatment [$M_{drug} = 76.2$ ($SD = 3.2$), $M_{plac} = 76.0$ ($SD = 1.7$), $F_{(1,22)} = 0.07$, $p = 0.79$], nor an interaction of stimulus category with treatment ($F_{(1,22)} = 0.83$, $p = 0.44$).

The drug and placebo groups did not differ in task performance. Mean categorization accuracy was 97% ($SD = 5\%$) on drug and 98% ($SD = 6\%$) on placebo ($t_{(22)} = 0.49$, $p = 0.63$). Mean response time was somewhat faster for the drug group (507 ms, $SD = 129$ ms) than for the placebo group (549 ms, $SD = 129$ ms), but the difference was not significant ($t_{(22)} = 0.81$, $p = 0.43$). Thus, there were no systematic performance differences that could confound the drug-related gain effects on precision.

Studies 1 and 2: effects of alternative voxel selections

In assessing the precision of category representations, we used a normalized mask of the ventral temporal cortex that included all voxels, rather than specifying a variable subset of voxels for each participant on the basis of the degree to which the voxels predicted category membership. However, our results were robust to changes in this approach. We performed the angular dispersion analysis on data taken from 20 different voxel selections, ranging from the 5% most informative voxels, to 10%, up to the full mask. Restricting the analysis to smaller portions of the most informative voxels resulted in much lower overall angular dispersion values (indicating higher precision). However, it did not change the direction or significance of our effects (Fig. 4): precision was enhanced for the high-gain condition in each study (large pupil, Study 1; atomoxetine condition in Session 1, Study 2) for all masks tested, except in the case of distinct-category comparisons in Study 1, where masks of <35% of the full mask yielded marginally significant p values.

Discussion

Neurophysiological recordings (Usher et al., 1999; Berridge and Waterhouse, 2003; Aston-Jones and Cohen, 2005) and an influential theory of catecholamine function (Servan-Schreiber et al., 1990; Aston-Jones and Cohen, 2005) suggest that endogenous fluctuations in neuromodulators, such as norepinephrine and dopamine, tune neural gain in target brain areas to optimize cognition to varying circumstances. Bayesian theories of brain function and computational models (Servan-Schreiber et al., 1990; Knill and Pouget, 2004; Ma et al., 2006; Friston, 2009) suggest that one mechanism by which gain optimizes neural processing is through enhanced precision of representation. Confirming these ideas, our empirical results provide the first direct evidence that the precision of widely distributed cortical representations is augmented by catecholamine-mediated gain. The two studies showed a striking convergence of the effects of pupil-linked and pharmacologically induced increases in catecholamine-mediated gain, and hence suggest an important mechanism by which catecholamines can influence signal-detection performance and other behaviors (Servan-Schreiber et al., 1990; Usher et al., 1999; Aston-Jones and Cohen, 2005; Friston et al., 2012).

Our neural network simulations generated two predictions about the effect of gain on angular dispersion. Confirming the first prediction, we demonstrated that the precision of representation of exemplars within a given category is improved by catecholamine-mediated increases in neural gain. This suggests an important role for catecholaminergic gain modulation in perceptual encoding. In line with our second prediction, we observed that higher gain also increased representational overlap between patterns from distinct object categories. This somewhat counterintuitive finding suggests that distinct category representations in the ventral temporal cortex are overlapping (Haxby et al., 2001), and that these overlapping portions, which are assumed to code for features shared between categories, are represented more precisely under high gain. This can be thought of as enhanced precision within a higher-order category to which all stimuli in the experiment belong. Further work is needed to establish whether these findings generalize to representations in other brain areas than the ventral lateral temporal cortex, where catecholamines may have different effects (Hirata et al., 2006; Castro-Alamancos and Gulati, 2014).

Our aim was to identify the effects of catecholamine-mediated gain on a multivariate metric of representational similarity. Although fMRI in human subjects lends itself well to the study of multivariate, distributed object representations, this approach also forced us to rely on noninvasive measures of neural gain. In Study 1, we built on a growing literature linking pupil size to neural gain (Aston-Jones and Cohen, 2005; Gilzenrat et al., 2010; Nassar et al., 2012; Eldar et al., 2013; Cheadle et al., 2014; McGinley et al., 2015). However, the exact relationship between pupil size and neural gain requires further research. In Study 2, we used functional connectivity measures that show promise as noninvasive measures of gain, but previous validation relied on pupillometry and computational modeling (Eldar et al., 2013), other indirect methods for studying gain. It is also worth noting that gain modulation is unlikely to be the only mechanism by which catecholamines influence perception (Hurley et al., 2004). Thus, it remains a possibility that gain modulation was not the (only) mechanism by which precision was enhanced in the current studies. Nevertheless, our results follow directly from the predictions about the impact of increasing gain made by our own and previous neural network models (Servan-Schreiber et al., 1990; Usher et al., 1999), as well as from Bayesian characterizations of gain and precision (Knill and Pouget, 2004; Ma et al., 2006; Friston, 2009).

We tried to characterize the effects of neural gain, a computational property of a neuron's input–output function, on a whole-brain fMRI metric. This forms a promising approach by which low-level principles of neuromodulation may be linked via computational modeling to system-level brain function. The adaptive value of catecholamine-mediated gain modulation has been clearly specified in the domain of decision making and action (Aston-Jones and Cohen, 2005). An important goal for future research is to examine in detail the adaptive nature of catecholamine-mediated modulation of perceptual encoding.

References

Aston-Jones G, Cohen JD (2005) An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. Annu Rev Neurosci 28:403–450. CrossRef Medline

Berridge CW, Waterhouse BD (2003) The locus coeruleus–noradrenergic system: modulation of behavioral state and state-dependent cognitive processes. Brain Res Rev 42:33–84. CrossRef Medline

Bosch JA, de Geus EJ, Carroll D, Goedhart AD, Anane LA, van Zanten JJ, Helmerhorst EJ, Edwards KM (2009) A general enhancement of autonomic and cortisol responses during social evaluative threat. Psychosom Med 71:877–885. CrossRef Medline

Brett M, Anton JL, Valabregue R, Poline JB (2002) Region of interest analysis using an SPM toolbox. Paper presented at the 8th International Conference on Functional Mapping of the Human Brain. Sendai, Japan, June.

Buchanan TW, al'Absi M, Lovallo WR (1999) Cortisol fluctuates with increases and decreases in negative affect. Psychoneuroendocrinology 24:227–241. CrossRef Medline

Bymaster FP, Katner JS, Nelson DL, Hemrick-Luecke SK, Threlkeld PG, Heiligenstein JH, Morin SM, Gehlert DR, Perry KW (2002) Atomoxetine increases extracellular levels of norepinephrine and dopamine in prefrontal cortex of rat: a potential mechanism for efficacy in attention

deficit/hyperactivity disorder. Neuropsychopharmacology 27:699–711. CrossRef
Castro-Alamancos MA, Gulati T (2014) Neuromodulators produce distinct activated states in neocortex. J Neurosci 34:12353–12367. CrossRef Medline
Chamberlain SR, Muller U, Cleary S, Robbins TW, Sahakian BJ (2007) Atomoxetine increases salivary cortisol in healthy volunteers. J Psychopharmacol 21:545–549.
Cheadle S, Wyart V, Tsetsos K, Myers N, de Gardelle V, Herce Castañón S, Summerfield C (2014) Adaptive gain control during human perceptual choice. Neuron 81:1429–1441. CrossRef Medline
Cools R, Robbins TW (2004) Chemistry of the adaptive mind. Philos Trans A Math Phys Eng Sci 362:2871–2888. CrossRef Medline
Coull JT (2001) Modulation of attention by noradrenergic a 2-agents varies according to arousal level. Drug News Perspect 14:5–11. CrossRef Medline
Coull JT, Jones ME, Egan TD, Frith CD, Maze M (2004) Attentional effects of noradrenaline vary with arousal level: selective activation of thalamic pulvinar in humans. Neuroimage 22:315–322. CrossRef Medline
Devoto P, Flore G (2006) On the origin of cortical dopamine: is it a co-transmitter in noradrenergic neurons? Curr Neuropharmacol 4:115–125. CrossRef Medline
Devoto P, Flore G, Pira L, Longu G, Gessa GL (2004) Alpha2-adrenoceptor mediated co-release of dopamine and noradrenaline from noradrenergic neurons in the cerebral cortex. J Neurochem 88:1003–1009. CrossRef Medline
Eguíluz VM, Chialvo DR, Cecchi GA, Baliki M, Apkarian AV (2005) Scale-free brain functional networks. Phys Rev Lett 94:018102. CrossRef Medline
Eldar E, Cohen JD, Niv Y (2013) The effects of neural gain on attention and learning. Nat Neurosci 16:1146–1153. CrossRef Medline
Fisher NI, Lewis T, Embleton BJ (1987) Statistical analysis of spherical data. Cambridge, UK: Cambridge UP.
Friston K (2009) The free-energy principle: a rough guide to the brain? Trends Cogn Sci 13:293–301. CrossRef Medline
Friston KJ, Shiner T, FitzGerald T, Galea JM, Adams R, Brown H, Dolan RJ, Moran R, Stephan KE, Bestmann S (2012) Dopamine, affordance and active inference. PLoS Comput Biol 8:e1002327. CrossRef Medline
Gilzenrat MS, Nieuwenhuis S, Jepma M, Cohen JD (2010) Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function. Cogn Affect Behav Neurosci 10:252–269. CrossRef Medline
Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P (2001) Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science 293:2425–2430. CrossRef Medline
Hebb DO (1949) The organization of behavior: A neuropsychological approach. New York: Wiley.
Heil SH, Holmes HW, Bickel WK, Higgins ST, Badger GJ, Laws HF, Faries DE (2002) Comparison of the subjective, physiological, and psychomotor effects of atomoxetine and methylphenidate in light drug users. Drug Alcohol Depend 67:149–156. CrossRef Medline
Hirata A, Aguilar J, Castro-Alamancos MA (2006) Noradrenergic activation amplifies bottom-up and top-down signal-to-noise ratios in sensory thalamus. J Neurosci 26:4426–4436. CrossRef Medline
Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. Proc Natl Acad Sci U S A 79:2554–2558. CrossRef Medline
Hopfield JJ (1984) Neurons with graded response have collective computational properties like those of two-state neurons. Proc Natl Acad Sci U S A 81:3088–3092. CrossRef Medline
Hurley LM, Devilbiss DM, Waterhouse BD (2004) A matter of focus: monoaminergic modulation of stimulus coding in mammalian sensory networks. Curr Opin Neurobiol 14:488–495. CrossRef Medline
Joshi S, Li Y, Kalwani RM, Gold JI (2016) Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex. Neuron 89:221–234. CrossRef Medline
Kielbasa W, Pan A, Pereira A (2015) A pharmacokinetic/pharmacodynamic investigation: assessment of edivoxetine and atomoxetine on systemic and central 3,4-dihydroxyphenylglycol, a biochemical marker for norepinephrine transporter inhibition. Eur Neuropsychopharmacol 25:377–385. CrossRef Medline
Knill DC, Pouget A (2004) The Bayesian brain: the role of uncertainty in neural coding and computation. Trends Neurosci 27:712–719. CrossRef Medline
Koda K, Ago Y, Cong Y, Kita Y, Takuma K, Matsuda T (2010) Effects of acute and chronic administration of atomoxetine and methylphenidate on extracellular levels of noradrenaline, dopamine and serotonin in the prefrontal cortex and striatum of mice. J Neurochem 114:259–270. CrossRef Medline
Luce RD, Perry AD (1949) A method of matrix analysis of group structure. Psychometrika 14:95–116. CrossRef Medline
Ludolph AG, Udvardi PT, Schaz U, Henes C, Adolph O, Weigt HU, Fegert JM, Boeckers TM, Föhr KJ (2010) Atomoxetine acts as an NMDA receptor blocker in clinically relevant concentrations. Br J Pharmacol 160:283–291. CrossRef Medline
Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. Nat Neurosci 9:1432–1438. CrossRef Medline
McGinley MJ, David SV, McCormick DA (2015) Cortical membrane potential signature of optimal states for sensory signal detection. Neuron 87:179–192. CrossRef Medline
Murphy PR, O'Connell RG, O'Sullivan M, Robertson IH, Balsters JH (2014) Pupil diameter covaries with BOLD activity in human locus coeruleus. Hum Brain Mapp 35:4140–4154. CrossRef Medline
Nassar MR, Rumsey KM, Wilson RC, Parikh K, Heasly B, Gold JI (2012) Rational regulation of learning dynamics by pupil-linked arousal systems. Nat Neurosci 15:1040–1046. CrossRef Medline
Navazesh M (1993) Methods for collecting saliva. Ann NY Acad Sci 694:72–77. CrossRef Medline
Nieuwenhuis S, De Geus EJ, Aston-Jones G (2011) The anatomical and functional relationship between the P3 and autonomic components of the orienting response. Psychophysiology 48:162–175. CrossRef Medline
Schurger A, Pereira F, Treisman A, Cohen JD (2010) Reproducibility distinguishes conscious from nonconscious neural representations. Science 327:97–99. CrossRef Medline
Schurger A, Sarigiannidis I, Naccache L, Sitt JD, Dehaene S (2015) Cortical activity is more stable when sensory stimuli are consciously perceived. Proc Natl Acad Sci U S A 112: E2083–E2092. CrossRef Medline
Servan-Schreiber D, Printz H, Cohen JD (1990) A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior. Science 249:892–895. CrossRef Medline
Sikström S, Söderlund G (2007) Stimulus-dependent dopamine release in attention-deficit/hyperactivity disorder. Psychol Rev 114:1047–1075. CrossRef Medline
Swanson CJ, Perry KW, Koch-Krueger S, Katner J, Svensson KA, Bymaster FP (2006) Effect of the attention deficit/hyperactivity disorder drug atomoxetine on extracellular concentrations of norepinephrine and dopamine in several brain regions of the rat. Neuropharmacology 50:755–760. CrossRef Medline
Usher M, Cohen JD, Servan-Schreiber D, Rajkowski J, Aston-Jones G (1999) The role of locus coeruleus in the regulation of cognitive performance. Science 283:549–554. CrossRef Medline
Varazzani C, San-Galli A, Gilardeau S, Bouret S (2015) Noradrenaline and dopamine neurons in the reward/effort trade-off: a direct electrophysiological comparison in behaving monkeys. J Neurosci 35:7866–7877. CrossRef Medline
Willenbockel V, Sadr J, Fiset D, Horne GO, Gosselin F, Tanaka JW (2010) Controlling low-level image properties: the SHINE toolbox. Behav Res Methods 42:671–684. CrossRef Medline
Winterer G, Weinberger DR (2004) Genes, dopamine and cortical signal-to-noise ratio in schizophrenia. Trends Neurosci 27:683–690. CrossRef Medline