

Do you see the forest or the tree? Neural gain and breadth versus focus in perceptual processing

Eran Eldar^{1,2,*}, Yael Niv^{1,3}, Jonathan D. Cohen^{1,3}

¹ Princeton Neuroscience Institute, Princeton University

² currently at: Wellcome Trust Centre for Neuroimaging, University College London

³ Psychology Department, Princeton University

* Correspondence should be addressed to: Wellcome Trust Centre for Neuroimaging
University College London
London WC1N 3BG, UK
e.eldar@ucl.ac.uk

Running head: Neural gain and focus vs. breadth

Abstract

When perceiving rich sensory information, some may integrate its various aspects, while others may selectively focus on its most salient aspects. We propose that neural gain modulates the tradeoff between breadth and selectivity, such that high gain focuses perception on those aspects of the information that have the strongest, most immediate influence, whereas low gain allows broader integration of different aspects. We illustrate our hypothesis using a neural network model of ambiguous letter perception. We then show experimentally that, in line with the model, pupil-diameter indices of high gain are associated with letter perception that is more selectively focused on the letter's shape, or if primed, its semantic content. Finally, we use a recognition-memory experiment to show that the relationship between gain and selective processing also applies when the influence of different stimulus features is voluntarily modulated by task demands.

Keywords: Neural gain, perception, attention, memory, neural network, pupillometry

Introduction

The ability to focus on particular cues while ignoring others is necessary for us to perform many daily activities (Doverspike & Arthur, 1992; Green & Bavelier, 2003), and is especially useful when a particular cue signals an available reward or imminent danger that require our immediate attention (e.g., when seeing a bear in the woods). Such high-stakes situations typically increase physiological arousal which has long been thought to narrow attentional focus (Easterbrook, 1959). However, narrow focus on one or few cues can compromise performance in situations that require integration of a broad range of cues (Baddeley, 1972)—even basic functions such as recognizing a face depend on simultaneous integration of multiple cues (Richler et al., 2011).

Here, we propose that the balance between focus and breadth in perceptual processing is controlled by brain-wide levels of neural gain (Servan-Schreiber et al., 1990; Aston-Jones & Cohen, 2005; Eldar et al., 2013). Our hypothesis follows from the idea that gain enhances both excitation and inhibition, and thus increases the contrast between weak and strong neural inputs (Fig. 1). As a result, perceptual processing may become dominated by the strongest inputs—those that reflect the most salient signals—at the expense of weaker sources of information that are effectively ignored. In contrast, with low gain, weak and strong inputs produce more comparable levels of activity, and therefore, perception may reflect a broader range of sources of information.

Converging evidence suggests that neural gain is modulated throughout the brain by the locus-coeruleus norepinephrine (LC-NE) system (Servan-Schreiber et al., 1990; Aston-Jones & Cohen, 2005; Eldar et al., 2013; Gilzenrat et al., 2010; Jepma & Nieuwenhuis, 2011; Waterhouse et al., 1980; Waterhouse et al., 1984; Waterhouse & Woodward, 1980; Einhäuser et al., 2008; Murphy et al., 2011). Pertinent to the present study, pupil diameter indices of high LC-NE activity (Aston-Jones & Cohen, 2005; Joshi et al., 2016) have been shown to be associated with signatures of high gain in functional magnetic resonance imaging (fMRI) including, in particular, a higher contrast between weak and strong activations (Eldar et al., 2013). In line with our hypothesis concerning gain and focused attention, these same pupillary indices were also associated with more locally-focused neural dynamics, and with learning behavior that was more selectively focused on particular aspects of experimental stimuli (Eldar et al., 2013).

Building on this previous work, here we investigate the effects of gain on the balance between focus and breadth in perceptual processing. We begin by simulating the effects of gain on perception in a neural network model, to demonstrate that with high gain processing is more selectively dominated by the most salient stimulus features, whereas with low gain other features are taken into account as well. We then test for this effect experimentally, by manipulating feature saliency via subliminal priming and examining how the effect of this manipulation on perception varies with gain, indexed using pupillometry. In a second experiment, we extend our hypothesis to the domain of memory, and to circumstances in which differences in feature saliency arise from explicit task demands.

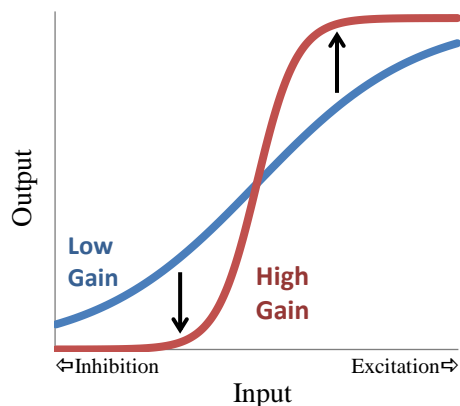


Fig. 1. Input-output function of a model processing unit (neuron, or possibly population of neurons) with low and high neural gain. Variations in neural gain can be captured in computational models by changing the gain of a standard non-linear activation function (e.g.,

$$\text{output} = \frac{1}{1 + e^{-\text{gain} \cdot \text{input}}}.$$

Simulation & Experiment 1: letter perception

Method

To test the degree to which perception reflects sources of information that differ in salience, we used ambiguous stimuli—characters whose shape most resembles one letter, but the letters each character is presented with favor its perception as a different letter (e.g., the middle character in *CAT* resembles the letter H, but resemblance of the whole string to the word *CAT* favors perception of the ambiguous character as the letter A; Fig. 2). Perception of such a stimulus involves competition between the letter's shape and its potential to form a familiar word with the adjoining letters. Since processing of words relies on and is thus secondary to processing of character shapes (McClelland and Rumelhart, 1981), we assumed that information about character shape is more immediately salient than is word context information. As a result, we predicted that participants with high levels of neural gain will perceive the letter that accords more with the character's shape, whereas participants with low gain will integrate the shape and word information more equally. To further test the relationship between gain and salience, we attempted to manipulate the relative salience of the character shape and word context by semantically priming half of the stimuli so as to increase the salience of the word context. We predicted that increasing word salience would reverse the relationship between gain and letter perception, such that with high gain, perception of primed stimuli will more strongly reflect the letter that accords with the word context.

THE
CAT

Fig. 2. Perception of the ambiguous middle character reflects its shape as well as the letters that surround it. Resemblance of the trigram stimuli to known words favors perception of an H in the top stimulus and A in the bottom stimulus.

It is not possible to measure gain directly in human participants, nor the norepinephrine activity thought to regulate gain. However, pupil diameter has been shown to closely correlate

with LC-NE activity in non-human primates (Aston-Jones & Cohen, 2005; Joshi et al., 2016) and with behaviors hypothesized to be associated with LC-NE activity in humans (Gilzenrat et al., 2010, Jepma & Nieuwenhuis, 2011; Einhäuser et al., 2008; Murphy et al., 2011). Furthermore, we recently showed that pupil dilation response, which is thought to be inversely related to baseline LC-NE activity (Aston-Jones & Cohen, 2005), correlates inversely with hallmarks of brain-wide fluctuations of gain in functional magnetic resonance imaging (fMRI; Eldar et al., 2013). The stimulus-locked phasic pupil dilation response, which is anticorrelated with the baseline pupil diameter, is particularly useful for between-subject comparisons, since it can be normalized to the baseline diameter, and thereby dissociated from factors that confound between-subject baseline comparisons. We therefore utilized the mean pupil dilation response to task stimuli throughout the experiment as an *inverse* measure of gain—below, high pupillary responses will be assumed to reflect low gain and vice versa.

Participants. 86 participants (mean age 21.7, age range 18-61, 69 female) performed the main experiment. A sample size of 80 participants was chosen *a priori* based on previous studies of semantic priming effects (Lucas, 2000), and data collection continued until the desired sample size was reached (6 participants who had fewer than 20 trials in which at least half of the baseline pupil diameter and pupil response measurements were free of artifacts had to be excluded from the sample). Participants were from the Princeton University area, and gave written informed consent before taking part in the study, which was approved by the university's institutional review board. Participants received either monetary compensation (\$10) or course credit for participation.

Experimental task. Participants were presented with 88 3-letter strings, 52 of which included an ambiguous character, one interpretation of which formed a word with the other 2 letters. To manipulate the salience of the potential word, half of the letter strings were preceded by subliminal presentation (33 ms) of a semantically related word. The other half were preceded by subliminal presentation of a similarly sized non-word (each letter string was semantically primed for half of the participants). We used semantic rather than repetition priming since the latter would involve priming of both the visual shape of one of the letters and the semantic meaning of the potential word.

Following the priming stimulus, the 3-letter target stimulus was presented for 225 ms, flanked by %%% on both sides so as to mask the priming stimulus, which could consist of more than 3 letters. The 3-letter string then disappeared from the screen and an arrow pointed to where the target letter had previously appeared. Participants had 5 seconds to choose, out of a list of 4 letters, which letter the target letter resembled the most. The list always included the two letters that the ambiguous character resembled and, in addition, two other letters that did not appear in the letter string, allowing us to validate that participants were not choosing letters randomly. Choices of one of the two letters that did not appear in the letter string were infrequent (less than 5% of trials) and were not included in the analyses below. Inter-trial interval was varied randomly (uniformly) between 6 s and 10 s – long enough to allow the pupil dilation response to resolve after each trial (Hoeks & Levelt, 1993).

Participants were explicitly instructed to try to choose the letter that most resembled the target character, and disregard whether the letters formed words. Although we cannot rule out that participants sometimes deliberately chose letters that form words, we have no reason to expect that such behavior would correlate with indices of neural gain. To account for possible response biases that may have resulted from conscious awareness of the priming manipulation, participants were asked during debriefing whether they saw any words appearing immediately before any of the letter strings. 10 participants reported that they saw such words. The results presented include the data from these participants, but analyses performed with and without these data produced similar results.

Stimuli. We designed 52 ambiguous characters using the Processing programming environment (Reas & Fry, 2007), each created by morphing one letter halfway into a different letter. Each ambiguous character was then embedded in a 3-letter string that could either form or not form a word depending on which of the two possible letters was perceived. To counteract the contextual effect of the word on perception of the ambiguous character, ambiguous characters were morphed so that their shape was slightly closer to the letter that did not form a word. Ambiguous characters were positioned in either the beginning or the end of the letter string, whereas participants were directed to fixate at the center. This ensured that the distance between the ambiguous letters and the focus of gaze remained constant throughout the experiment, while allowing variability in the location of the ambiguous letter. The words that letter strings could form were all medium-to-high frequency (above 10 per million; Kučera & Francis, 1967) picked using the MRC Psycholinguistic Database (Coltheart, 1981).

To prime the words that ambiguous characters could form, we used semantically related words, three to seven letters long. To avoid shape-related priming effects, prime words included neither of the two letters that the ambiguous letter resembled, nor other visually confounding letters (e.g., due to visual resemblance F could favor perception of E).

To ensure that participants were paying attention to all three letters of each string and not just to the ambiguous letter, we designed 36 additional 3-letter strings, in which one letter was somewhat morphed, but participants were asked to identify one of the non-morphed letters.

To maximize the ambiguity of the ambiguous characters, we conducted several iterations of a preliminary experiment, the results of which were used to adjust the stimuli so as to equate the probability of the ambiguous character being perceived as the word-forming and non-word-forming letters. On each iteration, four to six participants performed the task described above. Then, every ambiguous letter that was perceived as one particular letter at least 80% of the time was morphed slightly toward the other letter. This process was iterated 6 times, for a total of 30 participants (mean age 20.4, age range 18-23, 25 female). Participants in this preliminary experiment were also from the Princeton University area, gave informed consent, and were compensated with \$10 or course credit.

To minimize luminance-related changes in pupil diameter, all stimuli were adjusted to be isoluminant with the background using the flicker-fusion procedure (Lambert et al., 2003) on the display system used in the experiment.

Pupillometry. An ASL Series 5000 remote optics eye tracker (Applied Science Laboratories, MA) was used to measure participants' left pupil diameter while they were performing the task. At the beginning of the experiment, a baseline measurement of pupil diameter at rest was taken for a period of 45 s. Pupil-diameter data were processed in MATLAB to detect and remove blinks and other artifacts. For each trial, baseline pupil diameter was computed as the average diameter over a period of 1 s prior to the beginning of the trial (at the end of the inter-trial interval, at which point pupil activity from the trial itself should have subsided). Pupil-dilation response was computed as the difference between the peak diameter recorded during the 4 s that followed the beginning of the trial and the preceding baseline diameter. All pupil dilation responses were normalized by the pre-experiment baseline pupil diameter. Horizontal displacement of gaze during stimulus presentation was quantified for all participants but one (for whom gaze data were not recorded due to a technical problem). Since gaze displacement might affect pupil diameter measurements, we used a control covariate indicating gaze displacement to validate that all reported correlations with pupil diameter could not be explained by differences in gaze displacement.

Neural network model of the task. To formalize our hypothesis, we first simulated perception of the stimulus 'CAT' using an established neural network model of letter and word perception (McClelland and Rumelhart, 1981). The network consisted of three layers: a 'visual' input layer, a letter layer and a word layer (Fig. 3a). Since C and T are unambiguous, their respective letter-layer units received maximal input (input = 1). In contrast, since the middle letter is ambiguous, the H and A letter-layer units received sub-maximal input (input < 1). To reflect the fact that the shape of the ambiguous letter was closer to H, we simulated stronger input to the H unit as compared to the A unit (see below for precise values). Since the task required participants to decide on a single percept for the ambiguous letter, in our simulation the A and H units competed through mutual inhibition, such that only one prevailed on any given trial. Finally, the unit representing the word CAT was connected with excitatory connections to the letters C, A and T with which it is consistent.

To simulate the limited exposure time used in the experimental task, input to the letter units was applied for 225 iterations. At each time step t , the activity a_i^t of every network unit i built up gradually according to a weighted sum of its inputs:

$$a_i^t = 0.9a_i^{t-1} + 0.1f\left(b_i + \sum_j w_{ij}a_j\right) + n \quad (1)$$

where b_i refers to the bias to unit i (initially set to -0.5 for all units), w_{ij} refers to the connection weight from unit j to unit i (set to +1 for excitatory connections, and -1 for inhibitory connections), $f(x)$ is the sigmoid activation function:

$$f(x) = \frac{1}{1 + e^{-gain \cdot x}} \quad (2)$$

and n is a normally distributed random noise variable.

The parameter *gain* in the sigmoid function was used to simulate the level of neural gain in the network, which was the same for all units¹. Finally, we simulated semantic priming of the word information by adding excitatory input to the CAT word unit for 33 iterations immediately prior to the stimulus input.

Since the network's task was to reach a decision between perception of the middle character as A or H, we simulated two mutually-inhibitory decision units (activity initialized to 0, bias = 0), each of which had a bidirectional excitatory connection with its corresponding letter unit. Following presentation of the stimulus, the network switched to a 'decision mode,' in which the biases of the letter-layer A and H units were increased from their resting state of -0.5 to 0, simulating the allocation of attention to the letter-decision task (Cohen et al., 1990). Activity continued to be updated using Equation 1 until one of the decision units reached activity level of 0.9 or 1000 iterations were completed, at which point the probability of choosing the word-forming letter (A) was computed as the activity of the A decision unit divided by the sum of the activity of both decision units. In addition, to simulate the relationship between letter choice, reaction time, and a noisy pupillary index of gain, we computed reaction time as the number of iterations the network needed to reach a decision (max 1000 iterations), and pupil response as the true level of gain used in the simulation plus randomly distributed noise (with standard deviation between 0 and 10). The strength of the inputs to the H (0.52) and A (0.25) letter-layer units and the level of noise (standard deviation = 0.035) were adjusted so as to make the network equally likely to decide in favor of H or A under conditions of low gain ($gain \leq 4$).

Statistical analysis. Analyses were carried out using MATLAB. All predictions concerning individual differences were tested using correlation and regression analyses across the whole group of participants. Median splits were only used for complementary analyses and for visualization of results. Reported correlation values are Pearson correlation coefficients. Significance of across-participant correlations was computed using the Student's *t*-distribution. Averaging of correlation coefficients was preceded by Fisher *r*-to-*z* transformation and followed by Fisher's *z*-to-*r* transformation, so as to mitigate the problem of the non-additivity of correlation coefficients. Group-level significance of within-participant correlations was computed using a one-tailed one-sample Student's *t*-test on the vector of correlation coefficients following Fisher *r*-to-*z* transformation. To account for potential outliers, correlations and interactions were additionally tested using robust regression analysis with default MATLAB options (bisquare weighting, tuning constant 4.685; Holland & Welsch, 1977; Rousseeuw & Leroy, 2005). All statistical tests were two tailed except for within-participant tests that were used to validate between-participant results, as mentioned below.

¹ Having the level of gain affect all network units similarly is consistent with the widespread distribution of LC-NE projections throughout the brain (Aston-Jones & Cohen, 2005).

Results

Simulation. We used the neural network shown in Fig. 3a to simulate perception of the stimulus CAT with different levels of neural gain. With low gain, the shape of the ambiguous character initially drove the network to perceive the letter H, but as the surrounding letters activated a representation of the word CAT, perception of the letter A increased. As a result, the network was equally likely to perceive the ambiguous character as A or H (left part of Fig. 3b,c). In contrast, with high gain, the effect of the ambiguous character's shape was enhanced, and thus, the network settled on the letter that does not complete the word before the word representation had a chance to influence the outcome (Fig. 3b, red line; Fig. 3c, left two plots). Thus, a higher level of gain, despite being applied similarly to all network units, focused processing on the ambiguous character's shape.

Our hypothesis suggests that the focusing effect of high gain acts in favor of the character's shape because the shape information has a stronger and more immediate impact (see Supplementary Fig. 2 for additional simulations investigating the distinction between strength and immediacy). To test this explanation, in a second set of simulations we pre-activated (i.e., primed) the word representation in order to strengthen the word's immediate impact relative to the character shape information. In this case, high gain had the opposite effect: when the word was primed, higher gain became associated with a higher frequency of word-congruent letter perception (Fig. 3b, blue line; Fig. 3c, right plots).

Experiment. To test for similar effects of neural gain on perception in humans, we showed participants letter strings such as CAT, and asked them to indicate which letter the ambiguous character resembled most. As predicted, and consistent with our simulations, participants with a lower mean pupil response (indicating higher sustained neural gain) were more likely to perceive the ambiguous character as the letter that does not complete a word ($r = 0.30$, $t_{78} = 2.78$, $p < 0.01$; robust regression: $t_{78} = 3.3$, $p < 0.005$; Fig. 4a), indicating that high gain was associated with perceptions that more strongly reflected the ambiguous character's shape. Moreover, the relationship between pupil response and letter perception changed towards the opposite direction when the words were subliminally primed using semantically related words (e.g., the stimulus CAT was preceded by subliminal presentation of the word DOG; $r = -0.12$ with priming vs. $r = 0.30$ without priming; correlation of pupil response and difference between priming and no-priming condition: $r = -0.28$, $p = 0.01$; robust regression, interaction between pupil response and priming: $t_{156} = 2.7$, $p < 0.01$; Fig. 4b). While semantic priming generally increased word-congruent letter choices (mean increase $4.6\% \pm 1.8\%$, $t_{79} = 2.6$, $p = 0.01$), it did so only in participants whose pupil responses indicated high gain (i.e., mean pupillary response below median; Fig. 4c). In these participants, perception primarily reflected character shape when the word was not primed, whereas when the word was primed, perception primarily reflected the word context (mean increase $7.7\% \pm 2.8\%$, $t_{39} = 2.80$, $p < 0.01$). In contrast, participants whose pupil responses indicated low neural gain (i.e., whose pupil responses were higher than the median) were relatively unaffected by the saliency manipulation, exhibiting almost equal sensitivity to letter shape and word in both conditions (mean increase $1.5\% \pm$

2.3%, $t_{39} = 0.69$, $p = 0.50$). Together, these findings suggest that participants whose pupil diameter indicated high gain predominantly processed the most salient aspect of the stimuli, irrespective of its source.

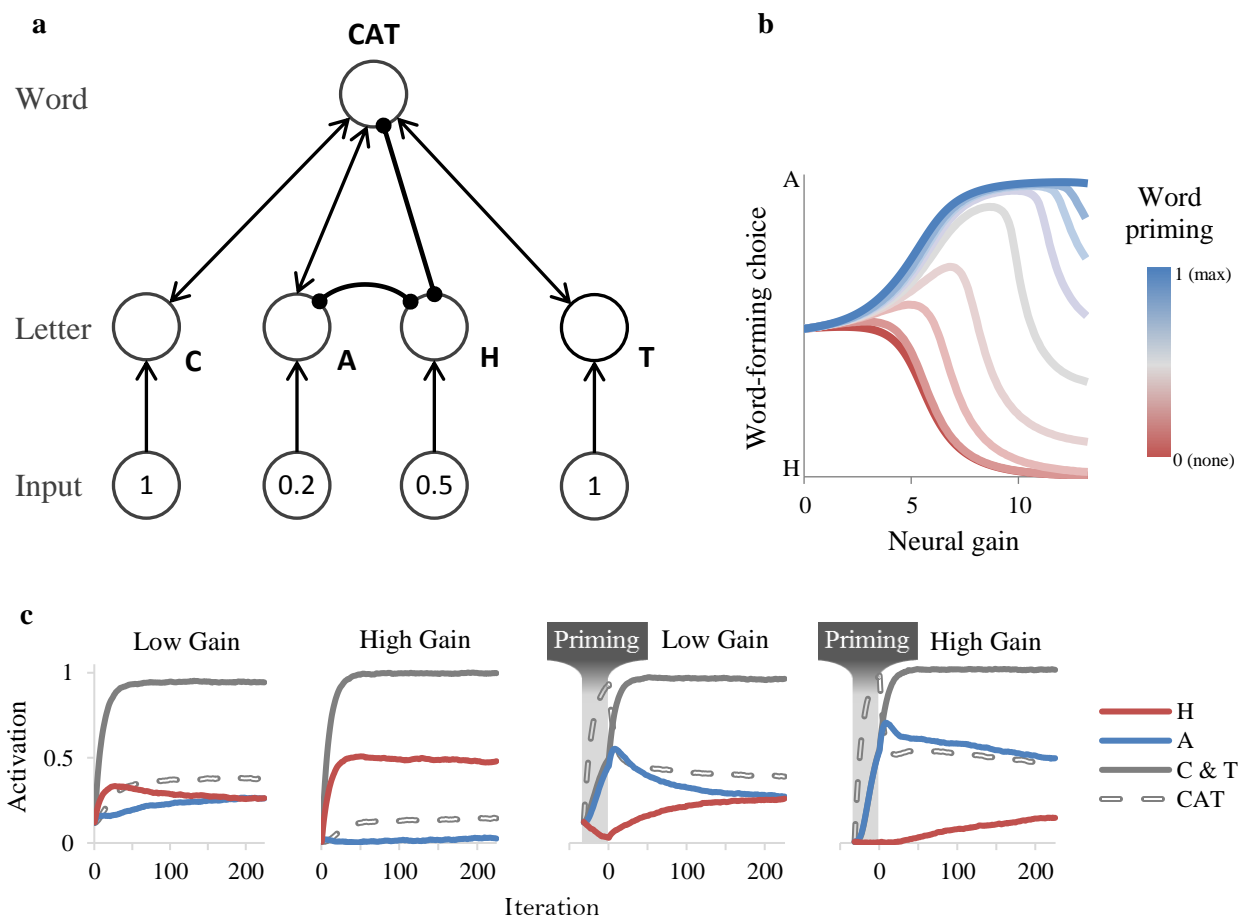


Fig. 3. A neural network model of the effect of neural gain on ambiguous letter perception. (a) The model's simulation of perception of a stimulus similar to the bottom one in Fig. 2. Arrows: excitatory connections, balls: inhibitory connections. (b) Simulated letter perception as a function of neural gain and the degree of priming of the CAT word-layer unit. 1000 simulations were conducted with each setting of gain and priming. Levels of gain higher than 3 are sufficient for the network to choose letters that correspond to the input (rather than choose randomly; see Supplementary Fig. 1). A similar network without inhibitory top-down connections produced similar results. (c) Trajectories of activation of the letter (H, A, C & T) and word (CAT) units, with low ($g = 4$) and high ($g = 10$) gain, without priming and with maximal priming. Iteration 0 indicates onset of the CAT stimulus input. Shaded area indicates duration of word priming.

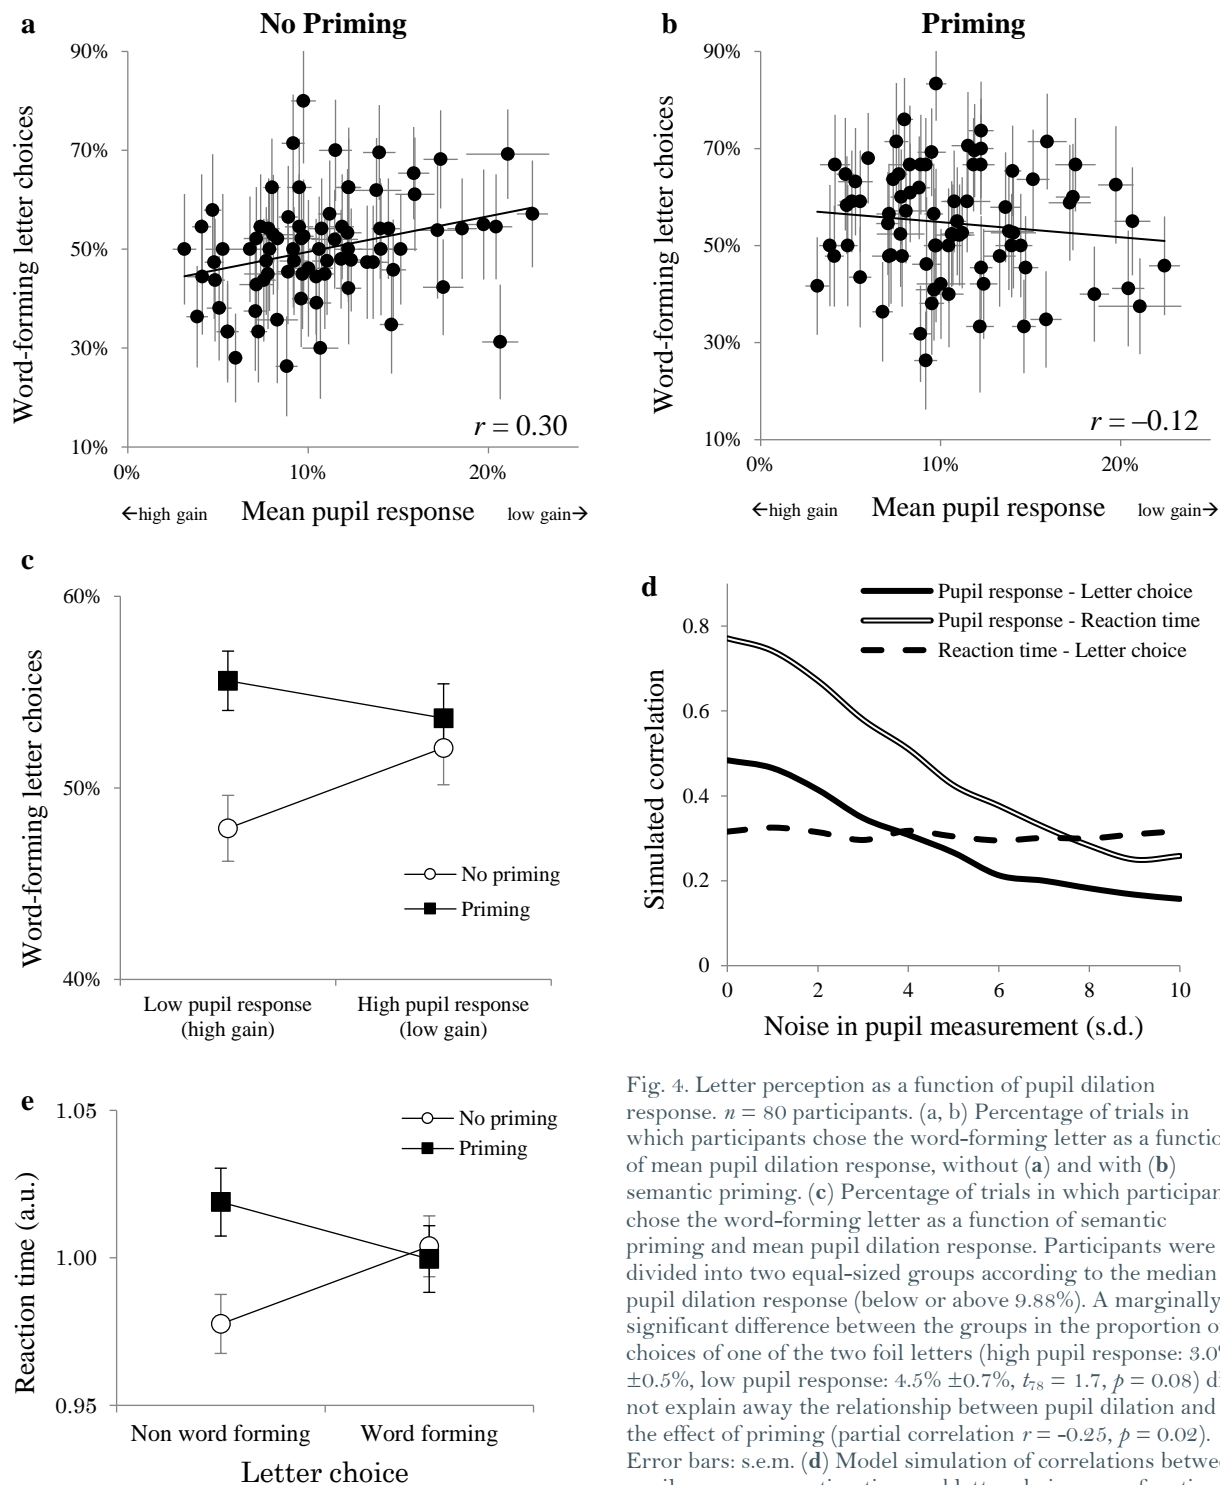


Fig. 4. Letter perception as a function of pupil dilation response. $n = 80$ participants. (a, b) Percentage of trials in which participants chose the word-forming letter as a function of mean pupil dilation response, without (a) and with (b) semantic priming. (c) Percentage of trials in which participants chose the word-forming letter as a function of semantic priming and mean pupil dilation response. Participants were divided into two equal-sized groups according to the median pupil dilation response (below or above 9.88%). A marginally significant difference between the groups in the proportion of choices of one of the two foil letters (high pupil response: $3.0\% \pm 0.5\%$, low pupil response: $4.5\% \pm 0.7\%$, $t_{78} = 1.7$, $p = 0.08$) did not explain away the relationship between pupil dilation and the effect of priming (partial correlation $r = -0.25$, $p = 0.02$). Error bars: s.e.m. (d) Model simulation of correlations between pupil response, reaction time and letter choices, as a function of noise in pupil measurement. The simulations predict that for

high levels of noise, the relationships between pupil response and reaction time, and between reaction time and letter choice, would be easier to detect than the direct relationship between pupil response and letter choice. 1000 simulation were conducted with each level of noise and neural gain, with and without priming. (b) Mean reaction time of human participants as a function of letter choice and semantic priming. Within participants, reaction times were faster for shape-related letter perceptions (“non-word”) in the no-priming condition, as compared to the priming condition (mean reaction time difference -0.05 ± 0.02 , group-level $t_{79} = 1.9$, $p < 0.05$). This result mirrors the between-participant interaction shown in (c).

While we observed the predicted relationship between pupil response and letter perception across participants, we did not find a similar relationship between perception and trial-by-trial variations in pupillary response within participants. One reason for this may be that within individual participants neural gain did not vary sufficiently over the course of the experiment for such a relationship to be detectable. Consistent with this possibility, the difference in mean pupil response between the first and second halves of the experiment was significantly lower within participant (mean 2.3%) than when measured between the first and second halves of different participants (mean 5.1%; $t_{79} = -8.6$, $p < 10^{-13}$). In addition, the high level of noise associated with pupillometric measurements makes it difficult to detect trial-by-trial within-participant effects.

To circumvent this problem, we used reaction time as an alternative index of neural gain. High gain leads to faster reaction times since all signals are amplified and thus the network settles on a decision more quickly. This is evident in the model (see Fig. 4d) as well as in the experiment, where the trial-to-trial correlation between pupil response and reaction time was significant in both the no-priming (mean $r = 0.08$, $t_{79} = 2.7$, $p < 0.01$) and priming (mean $r = 0.07$, $t_{79} = 2.3$, $p < 0.05$) conditions. Our simulations suggested that if pupillometric noise exceeds a certain level, the effects of neural gain on perception may be more robustly evident when using reaction time instead of pupil response as an indirect index of gain (Fig. 4d). Consistent with this, we found that reaction times were faster for shape-related letter choices in the no-priming condition, whereas when the word was primed, word-forming letter choices were faster (difference between conditions -0.05 ± 0.02 , $t_{79} = 1.9$, $p < 0.05$; Fig. 4e), indicating that high gain was associated in both cases with perceptions that more strongly reflected the more salient feature. This within-participant interaction between priming and reaction time mirrors the between-participant interaction between priming and pupil response (shown in Fig. 4c), suggesting that variations in gain between and within individuals had a similar impact on perception.

Experiment 2: recognition memory

In Experiment 1, we used subliminal priming to manipulate salience in order to investigate the interaction between salience and neural gain in the domain of perception. Here, we test whether these interactions extend to the domain of memory, and to salience that arises from voluntary allocation of attention in accord with task demands. To do this, we directed participants' attention toward the visual shapes of words by asking them to rate how easy it is to read each word, and then used a memory test to assess the degree to which memory for the words primarily relied on the words' visual shape (i.e., font). Based on our hypothesis that high gain focuses processing on the most salient information, we predicted that high gain would be associated with word memory that is more specific to the particular fonts the words originally assumed, when participants are directed to focus on word shape.

Method

Participants. 45 participants (mean age 19.8, age range 18-22, 28 females) performed the recognition memory experiment, and received course credit for participation. The sample size was chosen based on similar studies (Morris et al., 1977; Graf & Ryan, 1990), and data collection continued until the desired effective sample size was reached. Participants were Princeton University students who gave written informed consent before taking part in the study, which was approved by the university's institutional review board.

Experimental task. Participants were presented with 72 words in one of two highly dissimilar fonts, each for a period of 2 s. Half of the words were coupled with a task that focused participants' attention on word shape. Specifically, participants were asked to rate how readable the word was on a scale of 1 (very hard to read) to 4 (very easy to read). The other half of the words, which served as a control, were coupled with a semantic task that required processing both a word's shape (so as to read it) and its meaning. Specifically, participants were asked to report, for each word, whether it refers to an object that is man-made (for example, buildings) or not (for example, trees).

Words were presented in 4 blocks of 18 words each, and each block involved one of the two tasks. Task order was counterbalanced both within and between participants. In order to mitigate primacy and recency effects, each block started and ended with 4 words that were not included in the later recognition memory test. Words that participants indicated they were not able to read (on average 1.0 ± 0.18 words per block) were excluded from further analysis. Words were separated by a random inter-trial interval of 7 s to 9 s (uniform distribution).

Following an average period of 19.0 ± 0.18 minutes, during which participants performed an unrelated decision-making task, participants were tested on a word recognition memory test in which half of the words were foils, a quarter of the words had previously appeared in the same font (in either the readability or semantic task), and a quarter of the words had previously appeared in a different font. Participants were asked to indicate, for each word in the test, whether it had appeared in the first part of the experiment (regardless of font). Recognition memory performance was quantified by participants' hit rate.

Stimuli. 176 words, each 5 to 7 letters long, of medium-to-high frequency (above 10 per million; Kučera & Francis, 1967) were randomly assigned for each participant to different blocks or to be used as foils. Words were presented in an isoluminant color in capital letters in one of two fonts, Old English Text MT or Matura MT Script, chosen since they are relatively difficult to read and dissimilar from each other.

Pupillometry. A desk-mounted SMI RED 120Hz eye-tracker (SensoMotoric Instruments Inc., MA) was used to measure participants' left and right pupil diameters at a rate of 60 samples per second while they performed the experiment with their head fixed on a chinrest. Pupil diameter data were processed using the same methods as in Experiment 1. Mean pupil dilation response was computed separately for the readability and the semantic tasks.

Statistical analysis. Analyses were carried using the same methods as in Experiment 1.

Results

The hypothesis that high gain focuses processing on the most salient information predicts that for words from the readability task, in which participants focused on words' visual shape, recognition memory would be more strongly degraded by font change for participants with high gain. Consistent with this prediction, while font did not affect word recognition for participants with a large pupillary response during the readability task (mean pupil response above median, indicating low gain: same font hit rate minus different font hit rate $-3.1\% \pm 3.6\%$, $t_{21} = -0.9$, $p = 0.39$), it did have a significant effect on memory in participants with low pupillary response indicative of high gain (hit rate difference: $+12.5\% \pm 3.3\%$, $t_{21} = 3.7$, $p < 0.005$). This result suggests that high gain was associated with more selective processing of word shape. Consistent with this, there was a significant correlation between pupil response and hit rate difference across all participants: $n = 45$, $r = -0.43$, $p < 0.005$; robust regression: $t_{43} = 3.4$, $p < 0.005$ (Fig. 5a,c). Moreover, this effect was not evident for words for which participants performed the control semantic task (correlation between pupil response and hit rate difference: $n = 45$, $r = -0.05$, $p = 0.76$; difference between readability and semantic tasks: $z = 1.92$, $p = 0.05$; robust regression, interaction between pupil response and task: $t_{86} = 3.2$, $p < 0.005$; Fig. 5b,c). Finally, pupil response did not significantly correlate with general recognition performance levels ($n = 45$, $r = -0.01$, $p = 0.97$), indicating that pupillary indices of gain primarily reflected an interaction with the distribution of attention, not overall task engagement. Together, these results indicate that high gain amplified the specificity of memory to those stimulus features to which participants' attention was directed by the experimental task.

Discussion

We showed the degree to which perception and memory were selectively focused on the most salient sources of information was correlated with variations in pupillary indices of neural gain. Our priming results demonstrate that the focusing effect of gain applies to the most salient source of information irrespective of its source (i.e., visual or semantic). Our recognition memory results further suggest that gain interacts with saliency regardless of whether saliency is determined by automatic processes, as in the case of priming, or by voluntary attention in response to task demands.

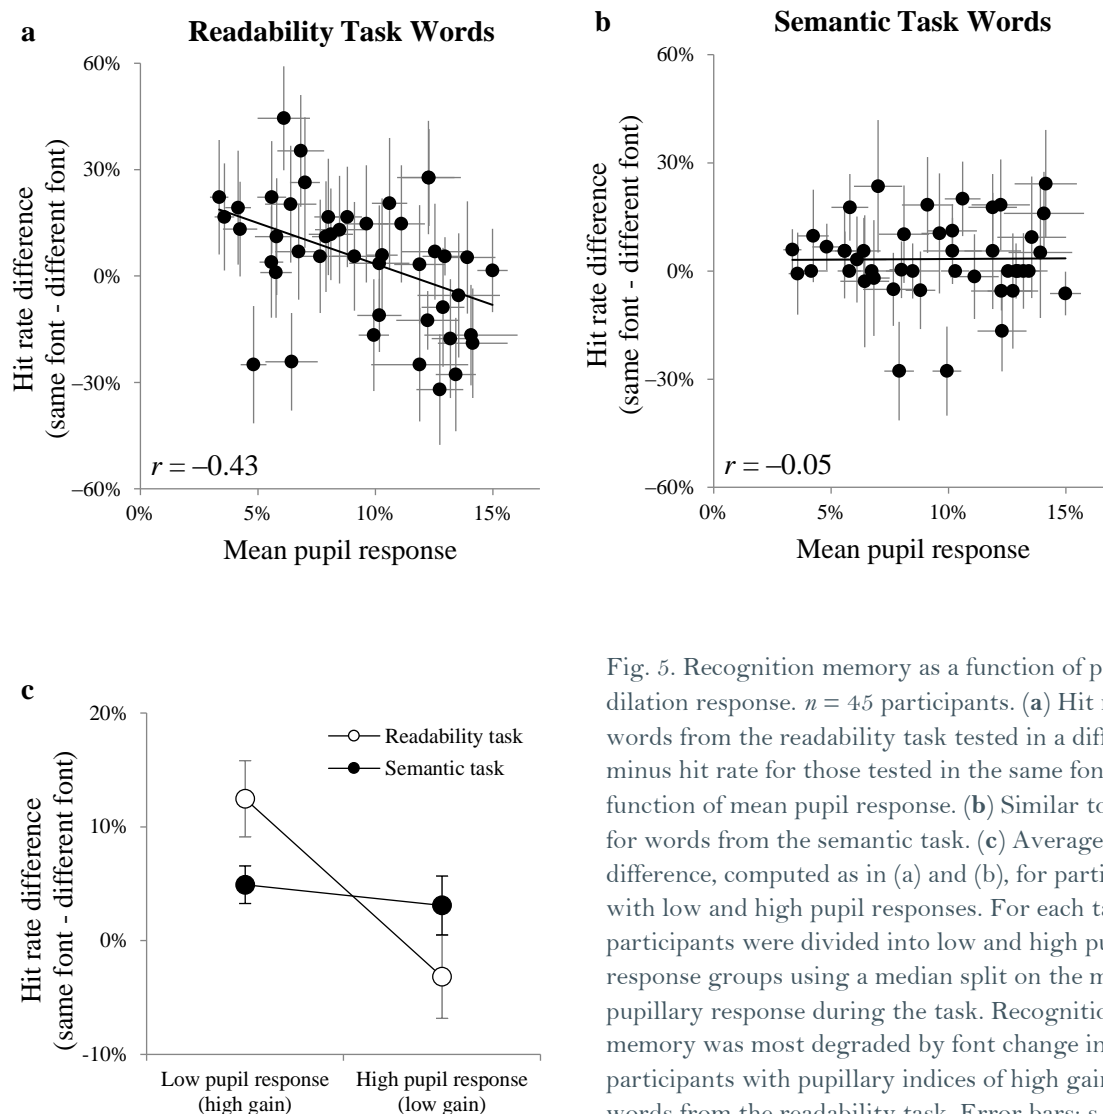


Fig. 5. Recognition memory as a function of pupil dilation response. $n = 45$ participants. (a) Hit rate for words from the readability task tested in a different font minus hit rate for those tested in the same font, as a function of mean pupil response. (b) Similar to (a), but for words from the semantic task. (c) Average hit rate difference, computed as in (a) and (b), for participants with low and high pupil responses. For each task, participants were divided into low and high pupil response groups using a median split on the mean pupillary response during the task. Recognition memory was most degraded by font change in participants with pupillary indices of high gain for words from the readability task. Error bars: s.e.m.

In the experiments, we jointly manipulated the strength and immediacy of particular sources of information. However, our simulations suggest that high gain should similarly favor stronger inputs and inputs that arrive earlier. We note also that the focusing effect of gain in our simulations was mainly driven by amplification of lateral inhibition, which has been suggested to underlie winner-takes-all dynamics (Coultrip et al., 1992). However, lateral inhibition cannot explain the faster reaction times associated with pupillary indices of higher gain, which suggest amplification of excitatory signals as well (amplified excitation could be feedforward or recurrent, as in Usher & Davelaar, 2002). Future work could clarify the contributions of these different factors to selectivity in information processing.

Our conclusions require several qualifications. The relationship between pupil diameter and behavior in our experiments was also evident in changes across time for individual participants, but only indirectly, through the relationship of both measures with reaction time. In addition,

although there is considerable evidence to support our pupillary measure of gain, the precise relationship between pupil dilation responses, central norepinephrine activity, and brain-wide fluctuations of gain has yet to be fully established. In particular, we cannot rule out that effects that are associated with high pupil dilations are driven by transient LC-NE activity, though we note that anticorrelations between pupil dilation and baseline pupil diameter are not thought to reflect variations in transient LC-NE activity (Joshi et al., 2016). Additional research is needed to dissociate the transient and sustained components of pupillary or LC-NE activity.

In sum, our findings suggest that neural gain modulates the tradeoff between focus and breadth in information processing: high gain causes us to focus on the most salient features of stimuli, whereas low gain favors more even-handed processing of all available features. Appropriate control over this tradeoff is essential in our daily activities, and its failure could underlie several neuropsychiatric conditions. In particular, impaired ability to focus is a hallmark of learning disabilities (Tarver et al., 1976; Richards et al., 1990), whereas the opposite end of spectrum, an excessively narrow focus, is thought to be a fundamental feature of autism spectrum disorders (Happé, 1996). Our findings implicate gain as a potential underlying mechanism for this important dimension of individual differences, and provide a practical way of measuring it.

Acknowledgements

This project was supported by the Howard Hughes Medical Institute (EE), and the John Templeton Foundation (JDC, YN).

Author Contributions

EE and JDC designed the study. EE ran the study and analyzed the data. All authors contributed to discussion and interpretation of the findings and writing the manuscript.

References

- Arthur Jr, W., & Doverspike, D. (1992). Locus of control and auditory selective attention as predictors of driving accident involvement: A comparative longitudinal investigation. *Journal of Safety Research*, 23, 73-80.
- Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annual Review of Neuroscience*, 28, 403-450.
- Baddeley, A. D. (1972). Selective attention and performance in dangerous environments. *British journal of psychology*, 63, 537-546.
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: a parallel distributed processing account of the Stroop effect. *Psychological review*, 97, 332.
- Coltheart, M. (1981). The MRC psycholinguistic database. *The Quarterly Journal of Experimental Psychology*, 33, 497-505.

Coultrip, R., Granger, R., & Lynch, G. (1992). A cortical model of winner-take-all competition via lateral inhibition. *Neural networks* 5, 47-54.

Easterbrook, J. A. (1959). The effect of emotion on cue utilization and the organization of behavior. *Psychological review*, 66, 183.

Einhäuser, W., Stout, J., Koch, C., & Carter, O. (2008). Pupil dilation reflects perceptual selection and predicts subsequent stability in perceptual rivalry. *Proceedings of the National Academy of Sciences*, 105, 1704-1709.

Eldar, E., Cohen, J. D., & Niv, Y. (2013). The effects of neural gain on attention and learning. *Nature neuroscience*, 16, 1146-1153.

Gilzenrat, M. S., Nieuwenhuis, S., Jepma, M., & Cohen, J. D. (2010). Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function. *Cognitive, Affective, & Behavioral Neuroscience*, 10, 252-269.

Graf, P., & Ryan, L. (1990). Transfer-appropriate processing for implicit and explicit memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 978.

Green, C. S., & Bavelier, D. (2003). Action video game modifies visual selective attention. *Nature*, 423, 534-537.

Happé, F. G. (1996). Studying weak central coherence at low levels: children with autism do not succumb to visual illusions. A research note. *Journal of Child Psychology and Psychiatry*, 37, 873-877.

Hoeks, B., & Levelt, W. J. (1993). Pupillary dilation as a measure of attention: A quantitative system analysis. *Behavior Research Methods, Instruments, & Computers*, 25, 16-26.

Holland, P. W., & Welsch, R. E. (1977). Robust regression using iteratively reweighted least-squares. *Communication in Statistics – Theory and Methods*, 6, 813-827.

Jepma, M., & Nieuwenhuis, S. Pupil diameter predicts changes in the exploration-exploitation trade-off: evidence for the adaptive gain theory. *Journal of cognitive neuroscience*, 23, 1587-1596 (2011).

Joshi, S., Li, Y., Kalwani, R. M., & Gold, J. I. (2016). Relationships between Pupil Diameter and Neuronal Activity in the Locus Coeruleus, Colliculi, and Cingulate Cortex. *Neuron*, 89, 221-234.

Kučera, H., & Francis, W. N. (1967). *Computational analysis of present-day American English*. Providence, RI: Dartmouth Publishing Group.

Lambert, A., Wells, I., & Kean, M. (2003). Do isoluminant color changes capture attention? *Perception & psychophysics*, 65, 495-507.

Lucas, M. (2000). Semantic priming without association: A meta-analytic review. *Psychonomic Bulletin & Review*, 7, 618-630.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological review*, 88, 375.

- Morris, C. D., Bransford, J. D., & Franks, J. J. (1977). Levels of processing versus transfer appropriate processing. *Journal of verbal learning and verbal behavior*, *16*(5), 519-533.
- Murphy, P. R., Robertson, I. H., Balsters, J. H., & O'Connell, R. G. (2011). Pupillometry and P3 index the locus coeruleus–noradrenergic arousal function in humans. *Psychophysiology*, *48*, 1532-1543.
- Reas, C., & Fry, B. (2007). *Processing: a programming handbook for visual designers and artists* (Vol. 6812). Cambridge, MA: MIT Press.
- Richards, G. P., Samuels, S. J., Turnure, J. E., & Ysseldyke, J. E. (1990). Sustained and selective attention in children with learning disabilities. *Journal of learning disabilities*, *23*, 129-136.
- Richler, J. J., Cheung, O. S., & Gauthier, I. (2011). Holistic processing predicts face recognition. *Psychological Science*, *22*, 464-71.
- Rousseeuw, P. J., & Leroy, A. M. (2005). *Robust regression and outlier detection*. Hoboken, New Jersey: John Wiley & Sons.
- Servan-Schreiber, D., Printz, H., & Cohen, J. D. (1990). A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior. *Science*, *249*, 892-895.
- Tarver, S. G., Hallahan, D. P., Kauffman, J. M., & Ball, D. W. (1976). Verbal rehearsal and selective attention in children with learning disabilities: A developmental lag. *Journal of Experimental Child Psychology*, *22*, 375-385.
- Usher, M., & Davelaar, E. J. (2002). Neuromodulation of decision and response selection. *Neural Networks*, *15*, 635-645.
- Waterhouse, B. D., Moises, H. C., & Woodward, D. J. (1980). Noradrenergic modulation of somatosensory cortical neuronal responses to iontophoretically applied putative neurotransmitters. *Experimental neurology*, *69*, 30-49.
- Waterhouse, B. D., Moises, H. C., Yeh, H. H., Geller, H. M., & Woodward, D. J. (1984). Comparison of norepinephrine- and benzodiazepine-induced augmentation of Purkinje cell responses to gamma-aminobutyric acid (GABA). *Journal of Pharmacology and Experimental Therapeutics*, *228*, 257-267.
- Waterhouse, B. D., & Woodward, D. J. (1980). Interaction of norepinephrine with cerebrocortical activity evoked by stimulation of somatosensory afferent pathways in the rat. *Experimental neurology*, *67*, 11-34.