



Published in final edited form as:

Nat Hum Behav. 2017 ; 1: . doi:10.1038/s41562-017-0107.

Arousal-related adjustments of perceptual biases optimize perception in dynamic environments

Kamesh Krishnamurthy^{*1}, Matthew R. Nassar^{*2}, Shilpa Sarode³, and Joshua I Gold¹

¹Department of Neuroscience, University of Pennsylvania 19104-6074

²Department of Cognitive, Linguistic, and Psychological Sciences, Brown University 02912

³Department of Electrical Engineering, University of Pennsylvania 19104-6074

Abstract

Prior expectations can be used to improve perceptual judgments about ambiguous stimuli. However, little is known about if and how these improvements are maintained in dynamic environments in which the quality of appropriate priors changes from one stimulus to the next. Using a sound-localization task, we show that changes in stimulus predictability lead to arousal-mediated adjustments in the magnitude of prior-driven biases that optimize perceptual judgments about each stimulus. These adjustments depend on task-dependent changes in the relevance and reliability of prior expectations, which subjects update using both normative and idiosyncratic principles. The resulting variations in biases across task conditions and individuals are reflected in modulations of pupil diameter, such that larger stimulus-evoked pupil responses correspond to smaller biases. These results suggest a critical role for the arousal system in adjusting the strength of perceptual biases with respect to inferred environmental dynamics to optimize perceptual judgements.

Introduction

Perception is shaped by prior expectations (“priors”) on the statistical structure of the sensory world¹⁻⁶. When the environmental statistics are stationary and well known, priors on those statistics can bias the perception of relevant sensory stimuli^{7,8}. For example, the prevalence of relatively slow-versus fast-moving objects in the world can lead to biases in the perception of object speed⁹. However, many environmental statistics that are relevant to

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use: http://www.nature.com/authors/editorial_policies/license.html#terms

Corresponding Author: Joshua I. Gold, Department of Neuroscience, University of Pennsylvania, Philadelphia, PA 19104-6074, Phone: 215-746-0028, jjgold@mail.med.upenn.edu.

*contributed equally and listed alphabetically

Author Contributions

All authors designed the experiment and analyses and wrote the manuscript, KK collected the data, and KK and MRN analysed the data.

Competing Interests

None of the authors have any competing interests to report, financial or otherwise.

Data availability

The data that support the findings of this study are available from the authors upon request. Matlab code for all models and analysis is available at: <https://github.com/TheGoldLab/Projects/KrishnamurthyNassarEtAl2017>.

perception can be highly non-stationary. For example, the locations of sources of sensory input are constantly changing relative to a given observer. The goal of this study was to examine how priors on such dynamic features of the environment are updated and used to shape perception.

To achieve this goal, we developed an auditory-localization task that required human subjects to both predict and report the perceived location of a simulated sound source as the predictability of the location varied over time (Fig. 1a–c). The statistical structure of the task is similar to ones we used previously to show that people can make effective predictions in dynamic environments by adaptively modulating the influence of new information on existing beliefs^{10,11}. However, here we focus on the questions of if and how such dynamically modulated predictions affect their influence on the perception of ambiguous stimuli. In principle, these predictions could govern perceptual biases through a form of optimal (Bayesian) inference that takes into account dynamic changes in the priors^{10,12,13}. Specifically, as long as the statistical structure of the sampled locations in our task remains stable, new sounds can be used to develop increasingly reliable priors about the locations of subsequent sounds. These increasingly reliable priors should, in turn, have an increasingly strong and beneficial influence on the perception of those sounds, reducing localization errors (Fig. 1d,e). However, the statistics of the sampled locations can undergo abrupt change-points that render previously held priors irrelevant to new sounds. These seemingly reliable but irrelevant priors should not influence the perception of sound-source location, which under these conditions should be limited entirely by sensory uncertainty (Fig. 1f).

We also measured pupil diameter, an index of arousal that can reflect the activation of the locus coeruleus (LC)-norepinephrine (NE) system and has been implicated in rapidly updating inference processes in response to unexpected events or errors^{14–19}. Pupil diameter tracks the extent to which predictions are updated in response to new information in dynamic and perceptually unambiguous cognitive tasks¹¹. Here we tested the hypothesis that such changes in arousal play an important role in shaping perception. In particular, we examined whether the arousal system controls the extent to which perceptual judgments about ambiguous sensory stimuli are biased toward prior expectations in accordance with the relevance and reliability of those expectations.

Our results yield new insights into the relationship between perception and arousal. We show that the subjects' priors had a variable influence on their perceptual reports. This variability was predicted by changes in the relevance and reliability of those priors, across task conditions and individual subjects. These effects were encoded in both baseline and stimulus-evoked changes in pupil diameter, such that larger diameters corresponded to less influence of priors on the perception of that stimulus. Taken together, these findings support a fundamental role for pupil-linked arousal systems, including the LC-NE system, in adaptively adjusting the influence of priors on perception in accordance with environmental dynamics.

Results

Twenty-nine subjects performed both the dynamic localization task (Fig. 1) and a control task that required perceptual reports of simulated sound-source locations that lacked predictable, sequential structure. Overall, the subjects tended to perform both tasks in an effective manner, providing predictions on the dynamic task and perceptual reports on both tasks that corresponded strongly to the simulated sound-source locations (Fig. 2). On the control task, the Pearson's correlation between simulated and reported location had median [interquartile range, or IQR] values of 0.926 [0.895–0.944] across subjects (Fig. 2a,d). On the dynamic task, there were similarly high correlations for both the predictions and perceptual reports (predictions on non-change-point trials: $r=0.907$ [0.895–0.921], Fig. 2b,e; perceptual reports on all trials: $r=0.948$ [0.941–0.964], Fig. 2c,f). However, the subjects also tended to make errors that varied considerably from trial to trial on both tasks (Fig. 2g–i). Subsequent analyses focus on how the subjects minimized their errors on the dynamic task by exploiting the fluctuating predictability of sound-source locations on that task.

Dynamic, task-dependent modulation of perceptual biases

The subjects used both sensory and prior information to guide their perceptual reports on the dynamic task. We measured performance in terms of the variability of the distribution of trial-by-trial errors (quantified as the standard deviation, or STD, and denoted as σ). This variability was lower for perceptual reports on the dynamic task than for either: 1) predictions from that task (σ_{prior} ; Fig. 2h), or 2) perceptual reports on the control task that lacked sequential predictability and thus reflected more purely sensory processing ($\sigma_{sensory}$; Fig. 2g). Moreover, for individual subjects, these different measures of variability were related to each other, such that perceptual errors from the dynamic task were well approximated using the optimal, reliability-weighted combination of prior and sensory information ($\sigma_{optimal}^{-2} = \sigma_{prior}^{-2} + \sigma_{sensory}^{-2}$; Fig. 2i). This result implies that, on average, the subjects tended to not only use these two sources of information, but also combine them according to their relative reliabilities to optimize perceptual performance on the dynamic task.

This integration of prior and sensory information took into account the changes in the relevance and reliability of the priors that occurred throughout the dynamic task. These changes are illustrated in Fig. 3a, which shows prediction-error STDs averaged across subjects as a function of the number of sounds after a clearly noticeable change-point, or SAC (see legend for details), separately for the two noise conditions. Figure 3b shows linear contrasts that captured the salient, dynamic aspects of these changes for each subject (see inset in Fig. 3e illustrating the three contrasts: CP, describing the effects of a noticeable change-point; Exp, describing the effects of the number of sounds experienced following a noticeable change-point; and Noise, describing the high or low noise condition). Specifically, on change-point trials, predictions were irrelevant and hence most variable with respect to the subsequent sound-source location (signed-rank test for H_0 : the median of the distribution of per-subject CP contrasts, which compared change-points to other trials=0, $p < 10^{-5}$). After change-points, predictions became steadily more reliable as the number of sound sources experienced from the new distribution increased in both noise conditions

($p < 10^{-4}$ for Exp_{low} and Exp_{high} contrasts, which identified linear trends across SAC 2–6 for each of the two noise conditions). The predictions were also more reliable overall in the low-versus high-noise condition (Noise contrast, $p < 10^{-5}$). These dynamic trends were consistent with predictions from a normative model of predictive inference that had full knowledge of the generative statistics¹⁰. The model, which produced simulated predictions that were analyzed in the same way as the data, had task-dependent effects that were in the same directions and of roughly the same magnitude as the data, although the subjects tended to produce more variable predictions than the model (Fig. 3a,b diamonds).

These task-dependent changes in the subjects' predictions were associated with similar changes in the variability of their perceptual reports (Fig. 3c,d) and their confidence in those reports, as assessed by the frequencies of high-confidence reports (Fig. 3e,f). Perceptual-error variability tended to be higher for change-point trials, when predictions were irrelevant (CP contrast, $p < 10^{-5}$), and for the high-versus low-noise condition (Noise contrast, $p < 10^{-5}$). Perceptual-error variability also tended to decrease on experiencing more samples from the new distribution, with a reliable effect across individuals in the low-noise condition (Exp_{low} contrast, $p < 0.005$) but not the high-noise condition (Exp_{high} contrast, $p = 0.4$). These dynamics were also apparent in the subjects' confidence report trends (Fig. 3e,f), which reflected trial-by-trial awareness of the changes in perceptual variability and included similar dependencies on CP ($p < 10^{-4}$), Noise ($p = 0.032$), and Exp_{low} ($p = 0.03$) and less reliable dependencies on Exp_{high} ($p = 0.07$). Both the perceptual and confidence report effects were qualitatively similar, in direction and magnitude, to theoretical values computed from optimal combinations of each subjects' changing priors (circles in Fig. 3a,b) and their fixed sensory reliability estimated from the control task (Fig. 2g; see also Fig. 1d–f). These theoretical values also showed strong effects of CP, Noise, and Exp_{low} , and smaller effects of Exp_{high} (Fig. 3c–f, diamonds).

These behavioral dynamics reflected changes in the degree to which the subjects' priors biased their perceptual reports. We quantified perceptual bias as the slope of the relationship between the prediction error and the perceptual error measured on individual trials (Fig. 4a–c). A slope of zero implies no relationship between the prediction error and the perceptual error, and thus no bias towards the prior. In contrast, slope values that increase towards unity imply increasing biases of the perceptual reports towards the prior. This perceptual bias varied systematically as a function of task conditions. Specifically, perceptual bias was lower on change-points (CP contrast, $p < 10^{-5}$) and for the high-versus low-noise condition (Noise contrast, $p = 0.008$). Perceptual bias also tended to increase on experiencing more samples, although these effects were variable across individuals and not statistically reliable in the low-noise condition (Exp_{low} contrast, $p = 0.1$; Exp_{high} contrast, $p = 0.004$). These task-dependent changes in the biases were comparable in direction and magnitude to theoretically computed values given an optimal, reliability-weighted combination of the task-specific predictions on the dynamic task (circles in Fig. 3a) and fixed sensory reliability estimated from the control task (Fig. 2g), computed separately for each subject (diamonds in Fig. 4d,e). Despite these comparable task-dependent trends (compare circles and diamonds in Fig. 4e), the subjects' perceptual biases were on average smaller than the theoretical values (compare circles and diamonds in Fig. 4d). This shift was consistent with their overall worse predictions than the model (compare circles and diamonds in Fig. 3a). However, overall

performance, measured as perceptual-error variability, was relatively insensitive to this overall shift, as compared to task-dependent adjustments, in the magnitude of the perceptual biases (compare circles and triangles in Fig. 3c,d).

Individual differences in the modulation of perceptual biases

The above analyses demonstrated that for individual subjects, dynamic changes in the relevance and reliability of priors within an experimental session were associated with changes in the degree to which those priors biased perception. We identified similar effects across subjects, implying that individual differences in perception can reflect differences in how priors are updated and maintained in dynamic environments. Specifically, we compared subjects' overall biases to the variability of their sensory and prediction errors (linear regression of the mean perceptual biases of individual subjects from non-change-point trials as a function of the STD of perceptual errors from the control task and the STD of prediction errors across non-change-point trials from the dynamic task; F statistic=7.39, $p=0.002$). According to these fits and consistent with Bayesian theory, subjects with higher overall prior-driven perceptual biases tended to have higher sensory variability ($\beta=0.033$, t -test for $H_0: \beta=0$, $p=0.013$; Fig. 5a) and lower prediction variability ($\beta=-0.030$, $p=0.002$; Fig 5b). We also found individual differences in how perceptual biases changed as a function of particular task conditions, and that those differences were predicted by subject-specific changes in priors under those conditions. Subjects whose priors improved (i.e., became less variable) the most also tended to have the largest increases in prior-driven perceptual biases: 1) just after a change-point (Fig. 5e), 2) on experiencing samples from a new distribution (in the low-but not high-noise condition; Figs. 5c and d), or 3) between the high- and low-noise conditions (Fig. 5f). Thus, on average, subjects tended to weigh prior and sensory information according to their relative reliabilities, taking into account variability in the priors across task conditions and individual subjects.

To more quantitatively account for the factors that affected perceptual biases across task conditions and individual subjects, we used a linear model that included normative and non-normative terms that each were weighed according to their contributions to each subject's behavior (Fig. 6). Data generated by a purely normative model could capture some qualitative aspects of behaviour, but it systematically overestimated perceptual biases (Fig 6A). A linear model that included both normative and non-normative terms offered a better description of behaviour (Fig. 6B). The normative terms were extracted from a Bayesian model of perception, which generated perceptual biases that minimized simulated perceptual errors, given each subject's variable predictions and sensory estimates. These terms were: 1) prior relevance, which reflected the probability that the current sound came from the same generative distribution as the previous sound (and thus is related to the CP effects illustrated in Figs. 3 and 4; Fig. 6c); and 2) prior reliability, which reflected changes in the total width of the predictive distribution relative to the likelihood, given new samples (and thus is related to the Exp and Noise effects illustrated in Figs. 3 and 4; Fig. 6d). The non-normative terms included one describing a fixed bias as a function of the prediction error, one to allow the strength of perceptual bias to depend on reported confidence (i.e., whether the subject reported high confidence or not), and spatial terms to account for the subjects' overall tendency to give perceptual reports that were biased slightly towards straight ahead (Fig. 2f).

On average, the linear model captured the behavioral trends well (Fig. 6b), based on contributions of each of the terms described above that tended to vary in magnitude across subjects (Fig. 6e). By comparison, a parameter-free normative model captured some of the behavioral trends (Fig. 6a) but reported higher perceptual biases than subjects (compare red points and bar in Fig. 6e), particularly on change-points (compare green points and bars in Fig. 6e).

Modulations of perceptual biases reflected in pupil diameter

A key question addressed in this work is whether arousal systems, as reflected in pupil diameter, contribute to the dynamic modulation of perceptual biases. Using linear regression at each time-point relative to sound onset (the average sound-evoked pupil response from all probe trials and subjects is shown in Fig. 7a), we found that pupil diameter varied with several of the factors from the linear model that accounted for behavioral biases (Eq. 6; Fig. 7b). Specifically, prior reliability was reflected in the baseline diameter before presentation of the probe sound, with smaller baselines reflecting more reliable priors ($p=0.03$; Fig. 7c,h). However, prior reliability did not modulate the magnitude of the stimulus-evoked pupil response, after accounting for the baseline effect (Fig. 7f,i). In contrast, prior relevance was unrelated to baseline diameter but was robustly encoded by the stimulus-evoked pupil diameter, with larger evoked pupil responses reflecting lower prior relevance (Fig. 7b,e). This effect peaked around the time of the maximum sound-evoked pupil response (permutation test for effect duration: duration=1.0 s, $p=0.02$; Fig. 7i). The pupil response, but not the baseline, also reflected the subjects' upcoming confidence report, with high confidence corresponding to larger pupil diameters, particularly late in the fixation interval (duration=1.8 s, $p=0.01$; Fig. 7d,g,i; note that these duration estimates were limited by the size of our measurement window).

If the arousal system is contributing to the dynamic regulation of the influence of priors on perception, then pupil diameter may co-vary with adjustments in prior influence even after accounting for all of the factors in the behavioral linear model (for example, if variability in internal representations of sound-source location affect both behavior and arousal). We therefore included the residual perceptual bias from our model of behavior (Fig. 6) in our model of pupil diameter. A positive/negative value of the residual biases indicates that the subject was more/less biased by the prior on the given trial than predicted by the linear model. There was a trend toward positive coefficients for this term in explaining baseline pupil diameter (larger baseline diameters corresponded to slightly stronger biases than predicted by the behavioral model; $p=0.06$; Fig. 7h). In addition, there was a robust reflection of the residual bias term in sound-evoked pupil response (smaller responses near the peak of the evoked response corresponded to stronger biases than predicted by the behavioural model; duration=1.2 s, $p=0.02$; Fig. 7i). This residual bias effect implies that pupil diameter reflects not just particular factors like prior reliability and relevance that can be used to make effective predictions in dynamic environments¹¹, but also the extent to which those and other factors are actually used to bias perception from one stimulus to the next.

In addition to these average, within-subject effects, there were also across-subject relationships between pupil diameter and perceptual biases. In particular, stimulus-evoked pupil responses tended to be, on average, smaller in subjects with higher overall perceptual biases (PE term in Fig. 6e; Fig. 8c) or relevance-dependent biases (PE*relevance term in Fig. 6e; Fig. 8d). These effects were not evident for baseline pupil diameter (Fig. 8a,b). However, because the behavioral influences of overall perceptual biases and prior relevance covaried considerably across subjects ($r=0.77$, $p<10^{-9}$), we constructed a new linear model that included two individual-difference variables that corresponded to the shared and unique variance of the two behavioral coefficients. The effects of the shared term were negative for most of the measurement window (Fig. 8e; duration=2.2 s, $p=0.01$). In contrast, the unique-variance term did not show a strong relationship to average pupil traces. This result implies that subjects who had the strongest overall perceptual biases, and modulated them most according to prior relevance, tended also to have the smallest stimulus-evoked pupil responses.

To further quantify these within- and across-subject relationships between pupil diameter and task performance, we used pupil diameter to predict the subjects' perceptual biases. Specifically, we created three normalized variables to reflect within- and across-subject variability in pupil responses at the time of peak response (1.4 s following stimulus onset) along with their multiplicative interaction. Each pupil-derived variable was included as a modulator of prediction errors in three different linear models of perceptual errors. In the simplest model, pupil-derived measures alone predicted systematic differences in perceptual biases observed in the behavioral data (Supplemental Fig. 3a), such that biases were: 1) larger for trials in which pupil responses were smaller than average (t -test, $p<10^{-4}$), 2) larger for subjects who had smaller than average pupil responses ($p<10^{-3}$), and 3) modulated from trial to trial more steeply for subjects with smaller overall pupil responses ($p<0.01$; Supplemental Fig. 3b). Consistent with these relationships, the pupil-based measures offered a substantial improvement to the base model in terms of predicting behavior (likelihood-ratio test, $p<10^{-7}$; Supplemental Fig. 3c). The pupil-based measures also offered an explanatory advantage when added to more complex models that accounted for direct fixed effects (one coefficient for all subjects) or random effects (one coefficient per subject) of relevance, reliability, and confidence reports on perceptual biases ($p<10^{-4}$ for both models; Supplemental Fig. 3c). Taken together, these results imply that fluctuations in pupil diameter, particularly those mediated by stimuli and related to context relevance, can be used to determine the extent to which perception is biased towards pre-existing priors.

Discussion

We used an auditory-localization task to show that the influence of prior expectations on perception is regulated rapidly and adaptively in changing environments. This work combines and extends several lines of research. The first has emphasized the role of priors on the perception of an uncertain sensory stimulus¹². Many of these studies have focused on priors that are related to relatively stable properties of the environment, although several recent studies have shown that certain sensory or sensory-motor priors can be learned relatively rapidly^{9,20–23}. The second has shown that under a variety of conditions, individual variability in decision-making can involve differential use of priors²⁴. The third has

identified how predictions are updated and used to make decisions in dynamic environments^{10,25,26}. The fourth has related this dynamic updating process to changes in physiological arousal^{11,27}. We showed that many of these principles, including dynamic, arousal-related adjustments in predictions, apply to how priors are updated and used to guide perception.

These principles involve ongoing assessments of the relevance and reliability of priors that represent a form of statistical learning^{28,29}. We quantified this learning process using two variables derived from normative theory^{17,30–33}. The first, which we termed prior relevance, is closely related to unexpected uncertainty and reflects the probability that a new observation is consistent with recent history^{10,17}. The second, which we termed prior reliability, is a form of reducible uncertainty that reflects ambiguity, typically resulting from undersampling, about the current generative process^{32,33}. Previous work showed that new information exerts the least influence on existing predictions when those predictions are the most relevant and reliable^{10,25}. We showed analogous effects for perception: new sensory input exerts the least influence on perception, relative to the influence of priors (i.e., perceptual biases are largest), when the priors are the most relevant and reliable.

Both of these normative factors, scaled according to their effects on each subject's behavior, were reflected in modulations of arousal state as measured by pupil diameter. Prior reliability corresponded to changes in baseline pupil diameter, and prior relevance corresponded to changes in the stimulus-evoked change in pupil diameter. These modulations were similar to those that we reported previously for a predictive-inference task, but the different demands of our present task imply a broader relevance to different forms of information processing¹¹. Specifically, our previous findings implicated a role for arousal fluctuations in adjusting bottom-up effects of new sensory input on stored cognitive representations. In contrast, our present findings implicate a role for arousal fluctuations in adjusting top-down control exerted by stored representations on the interpretation of new sensory input.

This result has broad implications for decision-making. For simple sensory-motor tasks, sequential effects of choice and response times can reflect priors inferred from recent task patterns, even when the patterns are spurious and thus the effects are detrimental to overall performance^{34–36}. Our results suggest a role for stimulus-evoked arousal responses in minimizing such suboptimal biases, potentially by reducing the impact of the top-down signals that mediate them. Consistent with this idea, pupil dilations have been shown to be accompanied by reduced individual and sequential-choice biases on perceptual decision-making tasks^{37,38}. For more complex tasks, top-down prior information might be used to select task-relevant feature information and thereby reduce implicit processing biases^{39,40}. This effect might explain why individuals with larger evoked pupil responses tend to be more susceptible to their own implicit processing biases^{41,42}. Future work should address this possibility in paradigms that combine implicit sensory biases with stimulus history-dependent priors such as those used in our task.

These results are also consistent with the idea that transient increases in arousal, in response to surprising events or other factors, may generally correspond to higher sensitivity to

immediate sensory input^{43,44}. In principle, this increased sensitivity could emerge from an enhancement of feed-forward processing, perhaps though an increase in neural gain^{11,18,41,45}. An alternative, but not necessarily mutually exclusive, possibility supported by our results is that enhanced sensitivity to sensory input is afforded by a reduction in the effectiveness of top-down priors in regularizing, and thereby biasing, sensory percepts. Distinguishing and understanding the independent contributions of these alternatives to arousal-mediated information processing will require the development of new paradigms that can separately control the bottom-up and top-down flow of information.

We also found relationships between subjective confidence, perceptual biases, and pupil diameter. We measured confidence using a post-decision binary confidence report (high/low confidence), which previously has been linked to the sensory-driven decision variable that also governs the speed and accuracy of the perceptual decision^{46–48}. We showed that confidence is also modulated according to changes in the relevance and reliability of perceptual priors that affect perceptual errors. This modulation was also evident in pupil diameter, which reflected high confidence-report frequency even after accounting for the normative variables that also governed the perceptual biases. However, this extra effect was in the opposite direction as for the normative factors, relative to the behavioral effect: high confidence-report frequency corresponded to larger pupil diameters but stronger prior influence. This pupil effect is somewhat surprising given that pupil diameter can be enhanced under uncertain, rather than certain, conditions^{11,38,49–52} (but see²⁷). One possible explanation for this discrepancy is that the post-decision confidence report led subjects to anticipate the increased reward or risk associated with high confidence-report trials, leading to stronger arousal responses^{51,53}. This idea is supported by the time course of confidence-related pupil dilations, which had a maximal dilation immediately prior to the perceptual report. This idea also highlights the multiple, possibly interacting roles that the arousal system likely plays in even simple sensory-motor tasks like this one.

These multiple roles undoubtedly result from multiple mechanisms by which arousal affects neural information processing⁵⁴. One such mechanism likely involves cortical levels of norepinephrine (NE), which is controlled via neurons in the midbrain nucleus locus coeruleus (LC)¹⁸. Firing rates of LC neurons correlate with pupil diameter over relatively short timescales, which has prompted the suggestion that pupil diameter can be used as a proxy for LC activity^{18,19,55,56}. Thus, the factors in our task that corresponded to stimulus-evoked pupil dilations, such as more surprising stimuli with lower prior relevance, may also correspond to increased LC activation. This activation, in turn, would increase levels of cortical NE, which have been theorized to signal unexpected context changes and allow neural representations to reorient rapidly to meet changing contextual demands, possibly via modulations of the input/output gain of individual cortical neurons^{14,17,18,45,57}. In contrast, slower changes in pupil diameter, such as those related to our baseline modulations, may be more closely related to levels of acetylcholine released from the basal forebrain, which has been theorized to signal expected uncertainty of task-relevant beliefs^{17,58,59}. More work is needed to determine how these multiple, potentially interacting neuromodulatory systems help to regulate perception and decision-making in dynamic environments.

Experimental Procedures

Human subject protocols were approved by the University of Pennsylvania Internal Review Board. 29 subjects (16 female, 13 male) participated in the study after providing informed consent. Thirty-one additional subjects consented to the study but did not meet the inclusion criterion of participating in at least three experimental sessions. Our sample size was well powered to detect effects of $d' > 0.6$ (statistical power > 0.88 for $d' = 0.6$) providing sufficient sensitivity in the range of previously reported behaviour-pupil relationships.

Auditory-localization task

We used an auditory-localization task in which subjects heard sounds with varying source locations that were simulated using head-related transfer functions (HRTFs) from the IRCAM database (<http://recherche.ircam.fr/equipes/salles/listen/download.html>). Each sound was a sequence of five Gaussian noise pulses bandpass filtered between 100 Hz and 15 KHz. The pulses were 50 ms each with a delay of 10 ms following each pulse, for a total stimulus duration of 300 ms. The latency for the sound to reach the ears following the command execution was ~ 3 ms. For each subject, we tested a number of HRTFs during the initial session by playing sound sequences that circularly traversed the entire horizontal plane in 15° intervals. We picked the HRTF that was reported to give the most uniformly circular percept for the sound sequence. Each subject performed 3–6 total sessions.

Each subject completed two tasks per session. The first was a control localization task that required the subject to indicate the perceived location of simulated sound sources that were sampled independently and uniformly randomly along the frontal, horizontal plane. In each of 72 trials, the subject was required to: 1) fixate for 2.5 s on a central spot while listening to the auditory stimulus; and 2) indicate the perceived location of the sound using a mouse, which controlled a cursor that moved along a semi-circular arc on the computer screen that represented the range of possible sound-source locations (Fig. 1). Failure to maintain fixation resulted in a warning sound and trial break. Feedback was displayed on the screen after the subject reported the perceived location.

The second task was a dynamic localization task that required the subject to report predictions, perceptions, and confidence judgments of sound-source locations that were generated from a change-point process along the same horizontal plane. For this task, the subject listened to extended sequences of sounds generated by the change-point process, with an interval of 150 ms between each sound presentation. Each sound was paired with a visual cue indicating its simulated source location on the semi-circular arc. During the presentation of these sequences, no action was required. Occasionally, however, the sequences stopped, indicating the start of a “probe trial” with the following structure (Fig. 1c). First, the subject was required to predict the angular location of the next, upcoming probe stimulus on the arc using a mouse. Second, following the prediction, the subject was required to maintain fixation for 2.5 s on the same central spot used in the control task. The auditory probe stimulus, with no corresponding visual cue, was presented at the beginning of this fixation period. Fourth, after the fixation period ended, the subject indicated the perceived location of the probe stimulus using the mouse and the visual display. Fifth, the subject then reported confidence (high/low) on the accuracy of the perceptual report (Fig 1).

Each subject performed four blocks of the dynamic task per session, which included ~30 probe trials each. Each session lasted ~90 min, with some variability due to the self-paced nature of the prediction, perceptual report, and confidence reporting periods of the task (median [IQR] reaction times were: 1.72 [1.49–2.35] sec for the prediction, 2.02 [1.50–2.30] sec for the perceptual report, and 1.18 [0.94–1.43] sec for the confidence report).

The sequence of simulated sound-source locations for the dynamic task was determined according to a process that included both irreducible variability (noise) and abrupt discontinuities (change-points). Our goal was to test if and how these manipulations, which can affect the reliability and relevance, respectively, of new information on existing predictions, also affect perceptual reports that can, in principle, use such predictions to improve the perception of ambiguous stimuli¹⁰. Source locations were sampled from a Gaussian distribution with a standard deviation (STD) that was held constant within a block of 600 trials (10° or 20° for the low- or high-noise condition, respectively) and a mean that underwent abrupt change-points with a fixed probability, or hazard rate (H), of 0.15 for each sound sample. At each change-point, the mean of the generative distribution was resampled uniformly across the sound space, such that the newly generated source locations were independent from previous ones. The sequence was interrupted for probe trials at random using a procedure that ensured: 1) a roughly even distribution of probes occurring 1–6 sounds after a change-point (SAC); 2) that probes were separated by at least 8 sounds; and 3) the number of sounds between any two probe trials was the same, on average, regardless of the nature of the two probe trials (SAC 1–6). Thus, on some trials the probe sound-source location was independent of the previous stimulus sequence (SAC=1). On other trials, the probe location was generated from the same distribution that produced the immediately preceding locations (SAC=2–6).

Subjects were motivated to make accurate perceptual reports on each probe trial through an incentive system. They were instructed to report high confidence if they were confident that the true location was within a 16° window centered on their second (perceptual) report, and to report low confidence otherwise. A correct/incorrect high confidence report resulted in a score of (15/–10), and a correct/incorrect low confidence resulted in a score of (5/–3). Subjects were paid a bonus that depended on their total score.

Behavioral data analysis: contrasts

To provide an intuitive understanding of how behavior was affected by particular task conditions, we sorted probe trials into twelve conditions according to the recency of the previous change-point (SAC=1–6) and noise condition (high/low). To emphasize the effects of change-points on the behavioral reports, these analyses included data only from sequences following easily noticeable change-points, corresponding to changes in generative mean of at least twice the generative STD for SAC=1 (note that the linear model described below was used to analyze the full data sets, including all change-points). Perceptual and prediction errors were computed by subtracting reported percepts and predictions from the true (simulated) sound source location for each trial. For each condition, the STD of prediction and estimation errors was used as a metric of average absolute error magnitude.

To quantify how prediction errors, estimation errors, and average confidence reports depended on specific task conditions, we performed four orthogonal linear contrasts over our twelve task conditions. Each contrast was computed by multiplying a weight matrix by the measured prediction errors, estimation errors, or average confidence reports, aggregated according to the task conditions for a single subject. Weight matrices were mean centered and chosen to identify: 1) differences between change-point and non-change-point trials (CP); 2) linear increases with increases in the number of sounds experienced (Exp) following a change-point, from SAC=2 to SAC=6, in the high-noise condition (Exp_{high}); 3) comparable linear increases in the low-noise condition (Exp_{low}); and 4) differences between the high- and low-noise conditions (Noise). Thus, the contrasts provided a per-subject measure of how much each behavioral measurement was modified according to these factors. For Figs. 3–5, we considered only sound sequences following relatively large change-points, corresponding to at least twice the generative STD. Contrasts were normalized for presentation in Figs. 3 and 4 by dividing the contrast value for each subject by the standard deviation of that contrast across all subjects. This procedure allowed all contrasts to be meaningfully displayed on the same set of axes.

Behavioral data analysis: perceptual bias

To quantify the influence of the prior on the perceptual report, we measured the slope of the best-fit line to the relationship between prediction errors (prediction–true location) and perceptual errors (percept–true location) for the given task condition. Slopes close to one indicate a high perceptual bias, and slopes close to zero indicate low perceptual bias. To measure how the perceptual bias evolved as a function of the number of sounds after a change-point (SAC), we used the following linear model and included only data from sequences following noticeable change-points (jumps of at least twice the generative STD):

$$ERR_{percp} = \beta_0 + \beta_1 ERR_{pred,1}^{high} + \dots + \beta_6 ERR_{pred,6}^{high} + \beta_7 ERR_{pred,1}^{low} + \dots + \beta_{12} ERR_{pred,6}^{low} + \beta_{13} Bias_{spatial} \quad [1]$$

where ERR_{percp} is the perceptual error and $ERR_{pred,1}^{high}$ is the prediction error on change-point trials (SAC=1) in the high-noise condition, and so on. The last term, $Bias_{spatial}$ captures the slight bias in the perceptual reports towards center of the screen.

Behavioral data analysis: theoretical benchmarks

The theoretically expected overall perceptual-error STD per subject (abscissa in Fig. 2i) was computed from an optimal, reliability-weighted combination of prior and sensory

information: $\sigma_{theoretical}^{-2} = \sigma_{predictions}^{-2} + \sigma_{sensory}^{-2}$. The theoretically expected perceptual-error STD per subject (diamonds in Fig. 3c,d), given their corresponding predictions for each SAC

condition, was computed using $\left(\sigma_{theoretical}^{SAC}\right)^{-2} = \left(\sigma_{predictions}^{SAC}\right)^{-2} + \sigma_{sensory}^{-2}$. The theoretically expected frequency of high-confidence reports (diamonds in Fig. 3e,f) was computed as the probability mass contained in a 16° window centered at the mean of a Gaussian with a STD of the theoretically expected perceptual errors, $\sigma_{theoretical}^{SAC}$. Thus, the proportion of expected

high-confidence reports increased with narrower perceptual error distributions. The theoretically expected perceptual bias per subject (diamonds in Fig. 4d,e) was computed as $\sigma_{sensory}^2 / [(\sigma_{predictions}^{SAC})^2 + \sigma_{sensory}^2]$. In all of the above, $\sigma_{predictions}$ is the STD of prediction errors on non-change-point trials, $\sigma_{predictions}^{SAC}$ is the STD of prediction errors for the specified number of sounds after a change-point, and $\sigma_{sensory}$ is the STD of perceptual errors on the control task, computed per subject.

Behavioral data analysis: normative model

Auditory localization in a dynamic environment can be posed as a perceptual inference problem with the goal of inferring the location of the sound source on trial t (X_t) according to a noisy internal sensory representation of that sound source (λ_t) and the history of previously experienced sound sources ($X_{1:t-1}$). This problem can be simplified by exploiting the conditional independencies in the Markov change-point process through which sound sources are selected (see Supplementary Fig. 1). In particular, the sound sources locations on the current trial (X_t) are independent of those on previous trials ($X_{1:t-1}$) conditioned on the mean of the generative distribution on the current trial (μ_t). In turn, the mean of the generative distribution on the current trial (μ_t) is also independent of previous observations conditioned on: 1) the mean of the generative distribution on the previous trial (μ_{t-1}), and 2) a latent change-point variable that determines whether the mean is resampled on the current trial (s_t). Applying Bayes rule to invert the generative graph in Supplementary Fig. 1 yields the following inference equation:

$$P(X_t | \lambda_t, X_{1:t-1}) = \frac{P(\lambda_t | X_t) \sum_{\mu_t} P(X_t | \mu_t) \sum_{s_t} \sum_{\mu_{t-1}} P(\mu_t | \mu_{t-1}, s_t) P(s_t) P(\mu_{t-1} | X_{1:t-1})}{\sum_{X_t} P(X_t, \lambda_t | X_{1:t-1})} \quad [2]$$

where the likelihood $P(\lambda_t | X_t)$ reflects the conditional distribution of internal representations across true sound source locations; $P(X_t | \mu_t)$ reflects the conditional probability of a sound source location being generated from a particular generative mean; $P(\mu_t | \mu_{t-1}, s_t)$ reflects the process through which means are resampled on change-point ($s_t = 1$) trials; and $P(s_t)$ is the hazard rate (H), which was fixed to 0.15 for the task and all simulations. The likelihood $P(\lambda_t | X_t)$ was modeled as a normal distribution centered on X_t with a variance that was fixed for each subject to the variance of perceptual reports made by that subject on the control task ($\sigma_{sensory}^2$). $P(\mu_{t-1} | X_{1:t-1})$ is the distribution over possible generative means, which can be updated recursively. Although exact Bayesian solutions to this recursive problem exist^{13,30}, we use a normal approximation to the Bayesian mixture distribution with a mean ($\hat{\mu}$) and variance ($\hat{\sigma}^2$) that capture the key dynamics of normative inference and offers better descriptions of human behaviour¹⁰. As in previous work, predictions made using this approximation were more accurate than subject predictions. To account for this discrepancy, we created a subjective prediction $\hat{\mu}_{subj}$ by sampling a random normal variable with mean equal to $\hat{\mu}$ and a variance that was incremented on each trial according to the difference in variance of subject and normative prediction errors:

$$\hat{\sigma}_{subj}^2 = \hat{\sigma}^2 + \text{Var}(X - \text{subject predictions}) - \text{Var}(X - \hat{\mu}) \quad [3]$$

Thus, the overall model incorporates the prospect of sub-optimal predictions about the sound-source location but implements Bayesian optimal combination of these predictions with incoming sensory information according to environmental dynamics.

Perceptions and predictions from the normative model were simulated by sampling internal representations (λ_t) and subjective predictions ($\hat{\mu}_{subj}$) for each trial according to the actual sequence of sound source locations experienced. Descriptive statistics for model simulations were averaged across 100 such simulated runs.

In addition to simulating behavioral data, the normative model also allowed us to extract latent variables that guide normative adjustments in perceptual bias. In particular, the model adjusts bias towards prior expectations in accordance with the relevance and reliability of those expectations. The relevance of prior expectations (π_t) is, in our generative framework, equal to the probability that a change-point did not occur on this trial given all previous data. This probability was computed on each trial by marginalizing Eq. 2 over all dimensions other than s . The impact of normative priors also depends critically on their reliability relative to that of the likelihood distribution capturing the noisy internal sensory representation (λ_t):

$$\text{prior reliability: } \tau_t = \frac{\sigma_{\text{sensory}}^2}{\sigma_{\text{sensory}}^2 + \hat{\sigma}_{\text{subj}}^2 + \sigma_{\text{noise}}^2} \quad [4]$$

where τ_t is prior reliability, $\sigma_{\text{sensory}}^2$ is the variance of perceptual reports made by that subject on the control task, $\hat{\sigma}_{\text{subj}}^2$ is the variance on subjective assessments of the underlying mean, and σ_{noise}^2 is the expected variance of sound source locations about that mean. The sum of the latter two terms reflects the total variance on the model's predictive distribution over possible sound locations. To ensure that these latent variables best reflected circumstances experienced by the subject, we fixed the model predictions ($\hat{\mu}_{subj}$) to the actual subject predictions from each trial and computed each measure as the expected value across all possible values of λ_t using a grid-point approximation.

Behavioral data analysis: Linear model of perceptual bias

To provide a more complete description of how behavioral data from all conditions, including all generative change-point and non-change-point trials, depended on both normative and non-normative factors, we fit the following linear model to data from all trials in each session:

$$\begin{aligned} ERR_{\text{percp}}(t) = & \beta_0 + \beta_1 ERR_{\text{pred}}(t) + \beta_2 ERR_{\text{pred}}(t) \cdot \pi_t + \beta_3 ERR_{\text{pred}}(t) \cdot \tau_t \\ & + \beta_4 ERR_{\text{pred}}(t) \cdot bet + \beta_5 Bias_{\text{center}} + \beta_6 Bias_{\text{spatial}} \end{aligned} \quad [5]$$

where β_1 describes the overall prior bias; β_2 and β_3 describe the extent to which the overall bias is dynamically modulated by the prior relevance and reliability, respectively (see above); β_4 describes the interaction of prior bias with confidence report (a binary variable); β_5 describes the bias towards the center of the screen; and β_6 captures the angular spatial bias (mean perceptual error at the given angle) measured in the control task. Residuals from the model fit were signed according to the prediction error on each trial to create a residual bias term for use in pupil analysis.

Pupil measurements

Pupil diameter was sampled from both eyes at 60 Hz using an infrared eye-tracker built into the monitor (Tobii T60-XL). Pupil analyses used the mean value from the two eyes at each time point measured during fixation. We excluded from further analyses trials with blinks, which we identified using a custom blink-detection routine on the basis of missing pupil diameter measurements and/or vertical and horizontal eye position that deviated from fixation for at least 10 contiguous samples (median [IQR] percentage of excluded trials across sessions = 5.54 [3.16–9.16]%). For the remaining trials with <10 missing contiguous pupil samples, we linearly interpolate the data before low-pass filtering. Low-pass filtering was done using a Butterworth filter with a cut-off frequency of 4 Hz. These filtered measurements were then z-scored in each session. We also removed a linear trend in the average pupil diameter over the whole session to account for any slow drift. The pupil baseline for each probe trial was defined as the mean of the first three samples immediately preceding the measurement period for that probe trial.

Linear model relating pupil diameter to behavioral parameters

To measure how the variables driving behavior were encoded in pupil diameter, we used the following linear model to explain the fluctuations in: 1) the baseline pupil diameter, and 2) stimulus-evoked pupil response across the 2.5 s following the auditory stimulus:

$$\begin{aligned} \text{Pupil diameter} = & \beta_1 \pi_t + \beta_2 \tau_t + \beta_3 \text{Bet}_n + \beta_4 \text{Bias}_{\text{residual}} + \\ & \beta_5 \text{Previous baseline diameter} + \beta_6 \text{Time since Previous probe} + \\ & \beta_{7-9} (\text{low freq. terms}) \end{aligned} \quad [6]$$

where τ_t and π_t are the reliability and relevance, respectively; β_{1-4} capture relationships between pupil diameter and the computational and behavioral variables of interest; β_{5-6} capture persisting fluctuations in pupil diameter that are attributable to the previous trial; and β_{7-9} includes a set of three low-frequency cosine components for each session that capture task-irrelevant variability due to slow modulations or session-based differences in pupil diameter. The exact form of the cosine components was $(\cos(\pi \cdot k \cdot (2n - 1)/2N))$, where $k=0,1,2$; n is the trial number within the session; and N is the total number of trials in the session. When this model was fit to evoked pupil responses, an additional nuisance variable was added to the explanatory matrix that accounted for trial-by-trial differences in baseline diameter.

Significance testing for evoked pupil coefficients was done through cluster-based permutation testing to account for multiple comparisons over time. In short, t -tests were

performed on each set of coefficient values across subjects separately for each time point. Cluster size was determined according to the number of contiguous time points for which this t -test yielded $p < 0.05$. Cluster corrected p -values were determined by comparing cluster sizes attained in this way to those from a permutation distribution of maximal cluster sizes⁶⁰.

Pupil-predicted perceptual bias

Trial-by-trial pupil measurements were extracted for the time of peak pupil response (1.4 seconds) from the behavioral model. Trial-by-trial measurements from each subject were regressed onto a set of nuisance variables that included explanatory variables β_{5+} from Eq. 6 to remove variance attributable to potentially confounding factors. Residual pupil measurements were concatenated across subjects and then divided into two separate variables: one variable accounted for average measurements for each subject and one that reflected normalized deviations from the average measurement within each subject. An additional term was created through the multiplicative interaction of each subject's mean pupil response and pupil-response variability, to account for the possibility that individual differences in the overall arousal response modulate the extent to which trial-to-trial fluctuations in arousal modulate perceptual bias. The three resulting variable arrays were z-scored and multiplied by trial-by-trial prediction errors to create a predictor matrix. Trial-by-trial perceptual errors were regressed onto three separate models with and without the inclusion of the pupil predictor matrix: 1) a null model including an intercept term and a prediction error term to capture fixed effects of perceptual bias across all subjects as well as the spatial bias terms described above [NE]; 2) a fixed-effects model that also included interaction terms accounting for modulation of perceptual bias by prior relevance and reliability the subjects' confidence report [FE]; and 3) a random-effects model that included all terms in model 2 separately for each subject [RE]. Since the random-effects model used dummy variables to account for individual differences in perceptual bias, the pupil predictor matrix included only within-subject variability and thus only one additional parameter rather than three. The marginal benefit of pupil-predictor terms was evaluated through likelihood-ratio tests (evaluating the additional explanatory power offered by pupil predictors) and quantified using AIC, a likelihood-based measure of goodness-of-fit that applies a penalty for each model parameter.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Ana-Andreea Stoica and Naim Kabir for aiding in data collection. We thank Long Ding and Takahiro Doi for helpful comments. This work was funded by NIH grants F32 MH102009 (MRN) and R01 EY015260 and NSF 1533623 (JIG). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

References

1. Bar M. Visual objects in context. *Nature Reviews Neuroscience*. 2004; 5:617–629. [PubMed: 15263892]

2. Edwards W. Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. *Journal of Mathematical Psychology*. 1965
3. Link SW, Heath RA. A sequential theory of psychological discrimination. *Psychometrika*. 1975
4. Maddox WT, Bohil CJ. Base-rate and payoff effects in multidimensional perceptual categorization. *J Exp Psychol Learn Mem Cogn*. 1998; 24:1459–1482. [PubMed: 9835061]
5. Seriès P. Learning what to expect (in visual perception). 2013; :1–14. DOI: 10.3389/fnhum.2013.00668/abstract
6. Summerfield C, Egnér T. Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*. 2009; 13:403–409. [PubMed: 19716752]
7. Fischer BJ, Peña JL. Owl's behavior and neural representation predicted by Bayesian inference. *Nature Publishing Group*. 2011; 14:1061–1066.
8. Knill DC, Pouget A. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*. 2004; 27:712–719. [PubMed: 15541511]
9. Stocker AA, Simoncelli EP. Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience*. 2006; 9:578–585. [PubMed: 16547513]
10. Nassar MR, Wilson RC, Heasly B, Gold JI. An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *Journal of Neuroscience*. 2010; 30:12366–12378. [PubMed: 20844132]
11. Nassar MR, et al. Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*. 2012; 15:1040–1046. [PubMed: 22660479]
12. Knill, DC., Richards, W. *Perception as Bayesian Inference*. Cambridge University Press; 1996.
13. Wilson RC, Nassar MR, Gold JI. Bayesian online learning of the hazard rate in change-point problems. *Neural Comput*. 2010; 22:2452–2476. [PubMed: 20569174]
14. Bouret S, Sara SJ. Network reset: a simplified overarching theory of locus coeruleus noradrenaline function. *Trends in Neurosciences*. 2005; 28:574–582. [PubMed: 16165227]
15. Ullsperger M, Harsay HA, Wessel JR, Ridderinkhof KR. Conscious perception of errors and its relation to the anterior insula. *Brain Struct Funct*. 2010; 214:629–643. [PubMed: 20512371]
16. Harley CW. Norepinephrine and the dentate gyrus. *Prog Brain Res*. 2007; 163:299–318. [PubMed: 17765726]
17. Yu AJ, Dayan P. Uncertainty, neuromodulation, and attention. *Neuron*. 2005; 46:681–692. [PubMed: 15944135]
18. Aston-Jones G, Cohen JD. An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu Rev Neurosci*. 2005; 28:403–450. [PubMed: 16022602]
19. Joshi S, Li Y, Kalwani RM, Gold JI. Relationships between Pupil Diameter and Neuronal Activity in the Locus Coeruleus, Colliculi, and Cingulate Cortex. *Neuron*. 2016; 89:221–234. [PubMed: 26711118]
20. Adams WJ, Graf EW, Ernst MO. Experience can change the 'light-from-above' prior. *Nature Neuroscience*. 2004; 7:1057–1058. [PubMed: 15361877]
21. Berniker M, Voss M, Kording K. Learning priors for Bayesian computations in the nervous system. *PLoS ONE*. 2010; doi: 10.1371/journal.pone.0012686.g001
22. Burge J, Ernst MO, Banks MS. The statistical determinants of adaptation rate in human reaching. *J Vis*. 2008; 8:20–20.
23. Tassinari H, Hudson TE, Landy MS. Combining Priors and Noisy Visual Cues in a Rapid Pointing Task. *Journal of Neuroscience*. 2006; 26:10154–10163. [PubMed: 17021171]
24. Stanovich KE, West RF. Individual differences in reasoning: implications for the rationality debate? *Behav Brain Sci*. 2000; 23:645–65. discussion 665–726. [PubMed: 11301544]
25. Behrens TEJ, Woolrich MW, Walton ME, Rushworth MFS. Learning the value of information in an uncertain world. *Nature Neuroscience*. 2007; 10:1214–1221. [PubMed: 17676057]
26. Wilson RC, Nassar MR, Gold JI. A mixture of delta-rules approximation to bayesian inference in change-point problems. *PLoS Comput Biol*. 2013; 9:e1003150. [PubMed: 23935472]
27. Preusschoff K. Pupil dilation signals surprise: evidence for noradrenaline's role in decision making. 2011; :1–12. DOI: 10.3389/fnins.2011.00115/abstract

28. Tenenbaum JB, Kemp C, Griffiths TL. How to grow a mind: Statistics, structure, and abstraction. Science. 2011
29. Vapnik, V. The Nature of Statistical Learning Theory. Springer Science & Business Media; 2013.
30. Adams, RP., MacKay, D. Bayesian Online Change-point Detection. 2007.
31. Mathys C. A Bayesian foundation for individual learning under uncertainty. 2011; :1–20. DOI: 10.3389/fnhum.2011.00039/abstract
32. Payzan-LeNestour E, Bossaerts P. Risk, Unexpected Uncertainty, and Estimation Uncertainty: Bayesian Learning in Unstable Settings. PLoS Comput Biol. 2011; 7:e1001048. [PubMed: 21283774]
33. Preusschoff K, Bossaerts P. Adding Prediction Risk to the Theory of Reward Learning. Annals of the New York Academy of Sciences. 2007; 1104:135–146. [PubMed: 17344526]
34. Gold JJ, Law CT, Connolly P, Bennur S. The Relative Influences of Priors and Sensory Evidence on an Oculomotor Decision Variable During Perceptual Learning. Journal of Neurophysiology. 2008; 100:2653–2668. [PubMed: 18753326]
35. Jones M, Curran T, Mozer MC, Wilder MH. Sequential effects in response time reveal learning mechanisms and event representations. Psychological Review. 2013; 120:628–666. [PubMed: 23915086]
36. Zhang S, Huang H, Angela JY. Sequential effects: A Bayesian analysis of prior bias on reaction time and behavioral choice. CogSci. 2014
37. de Gee JW, Knapen T, Donner TH. Decision-related pupil dilation reflects upcoming choice and individual bias. Proceedings of the National Academy of Sciences. 2014; 111:E618–E625.
38. Urai A, Braun A, Donner TH. Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. Nature Communications. In press.
39. Frank MJ, Badre D. Mechanisms of Hierarchical Reinforcement Learning in Corticostriatal Circuits 1: Computational Analysis. Cerebral Cortex. 2012; 22:509–526. [PubMed: 21693490]
40. Collins AGE, Frank MJ. Cognitive control over learning: creating, clustering, and generalizing task-set structure. Psychological Review. 2013; 120:190–229. [PubMed: 23356780]
41. Eldar E, Cohen JD, Niv Y. The effects of neural gain on attention and learning. Nature Publishing Group. 2013; 16:1146–1153.
42. Eldar E, Niv Y, Cohen JD. Do You See the Forest or the Tree? Neural Gain and Breadth Versus Focus in Perceptual Processing. Psychol Sci. 2016; 27:1632–1643. [PubMed: 28195019]
43. Pfaff, DW. Brain arousal and information theory. 2006.
44. Sara SJ, Bouret S. Orienting and Reorienting: The Locus Coeruleus Mediates Cognition through Arousal. Neuron. 2012; 76:130–141. [PubMed: 23040811]
45. Servan-Schreiber D, Printz H, Cohen JD. A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior. Science. 1990; 249:892–895. [PubMed: 2392679]
46. Kepecs A, Uchida N, Zariwala HA, Mainen ZF. Neural correlates, computation and behavioural impact of decision confidence. Nature. 2008; 455:227–231. [PubMed: 18690210]
47. Kiani R, Shadlen MN. Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex. Science. 2009; 324:759–764. [PubMed: 19423820]
48. Persaud N, McLeod P, Cowey A. Post-decision wagering objectively measures awareness. Nature Neuroscience. 2007; 10:257–261. [PubMed: 17237774]
49. Jepma M, Nieuwenhuis S. Pupil diameter predicts changes in the exploration-exploitation trade-off: evidence for the adaptive gain theory. J Cogn Neurosci. 2011; 23:1587–1596. [PubMed: 20666595]
50. Lempert KM, Chen YL, Fleming SM. Relating Pupil Dilation and Metacognitive Confidence during Auditory Decision-Making. PLoS ONE. 2015; 10:e0126588. [PubMed: 25950839]
51. Satterthwaite TD, et al. Dissociable but inter-related systems of cognitive control and reward during decision making: Evidence from pupillometry and event-related fMRI. NeuroImage. 2007; 37:1017–1031. [PubMed: 17632014]
52. Wessel JR, Danielmeier C, Ullsperger M. Error awareness revisited: accumulation of multimodal evidence from central and autonomic nervous systems. J Cogn Neurosci. 2011; 23:3021–3036. [PubMed: 21268673]

53. Manohar SG, Husain M. Reduced pupillary reward sensitivity in Parkinson's disease. *npj Parkinsons Disease*. 2015; 1:15026.
54. Robbins, TW., Everitt, BJ. Arousal systems and attention. The MIT Press; 1995.
55. Boudet S, Richmond BJ. Sensitivity of locus ceruleus neurons to reward value for goal-directed actions. *Journal of Neuroscience*. 2015; 35:4005–4014. [PubMed: 25740528]
56. Nieuwenhuis S, De Geus EJ, Aston-Jones G. The anatomical and functional relationship between the P3 and autonomic components of the orienting response. *Psychophysiology*. 2011; 48:162–175. [PubMed: 20557480]
57. Mather M, Clewett D, Sakaki M. Norepinephrine ignites local hot spots of neuronal excitation: How arousal amplifies selectivity in perception and memory. *Behav Brain Sci*. 2015
58. Yu AJ. change is in the eye of the beholder. *Nature Publishing Group*. 2012; 15:933–935.
59. Reimer J, et al. Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex. *Nature Communications*. 2016; 7:13289.
60. Nichols TE, Holmes AP. Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Hum Brain Mapp*. 2002; 15:1–25. [PubMed: 11747097]

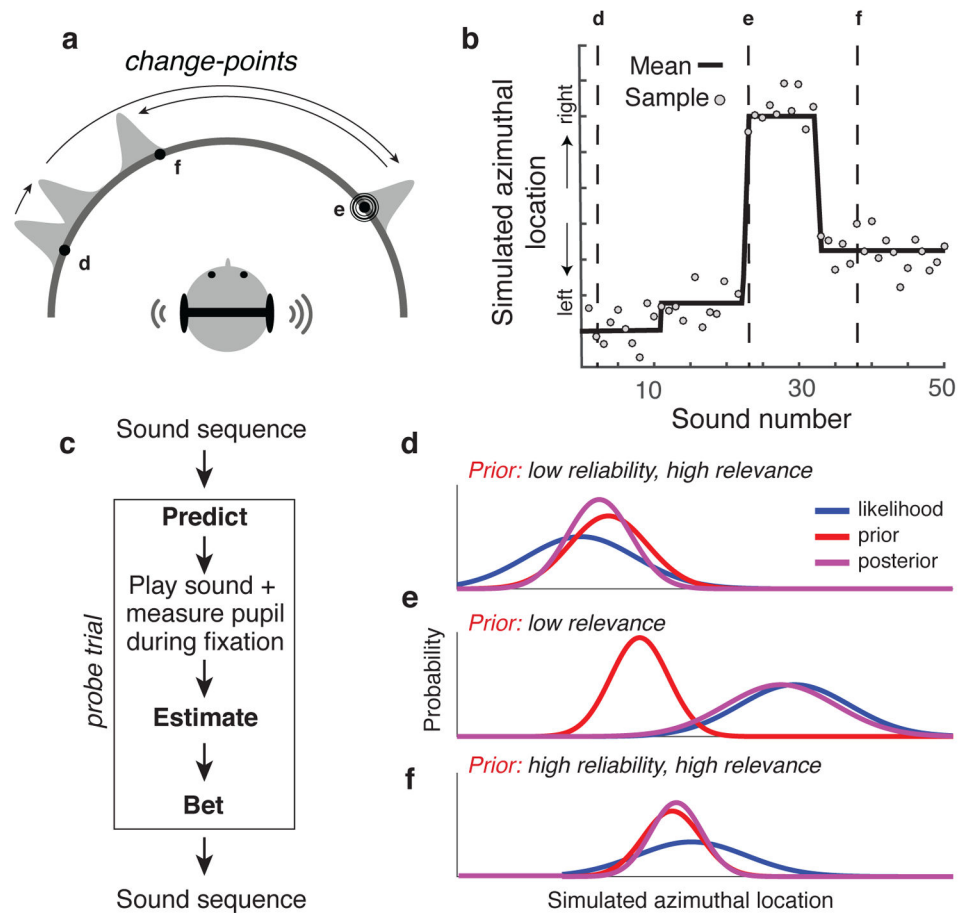


Figure 1. Dynamic sound-localization task

(a) Subjects listened via headphones to noise bursts with virtual source locations that varied along the frontal, azimuthal plane. The locations were sampled (points) from a Gaussian distribution (gray) with a mean that changed abruptly on unsignaled change-points (probability=0.15 for each sound) and a STD of 10° in low-noise blocks, 20° in high-noise blocks. The subjects listened passively to the sound sequence, except for occasional probe trials. All sounds except the probe sound were presented simultaneously with their corresponding locations on a semicircular arc shown on the isoluminant visual display, allowing subjects to develop priors on sound-source location based on both the auditory and visual signals and maintain a stable mapping between the two. (b) An example trial sequence showing the mean (solid line) and sampled (points) locations over 50 trials. Vertical dashed lines indicate randomly selected probe trials. (c) Probe-trial sequence. Using a mouse to control a cursor on the visual display, the subject reported: 1) the predicted location of the upcoming probe sound, followed by 250-ms fixation, presentation of the probe sound, then continued fixation for 2.5 s to allow for pupil measurements; 2) the estimated location of the probe sound; and 3) a high or low confidence report that the true location was within a small window centered on their estimate. The sound sequence then continued until the next probe. (d–f) Schematic illustrating the changing reliability and relevance of priors for the probe sounds in a and b, as indicated. Given a fixed-width

likelihood function, more reliable and relevant priors have a stronger and more beneficial influence on the percept, here represented as the posterior, which is most uncertain (widest) in **e** and least uncertain in **f**.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

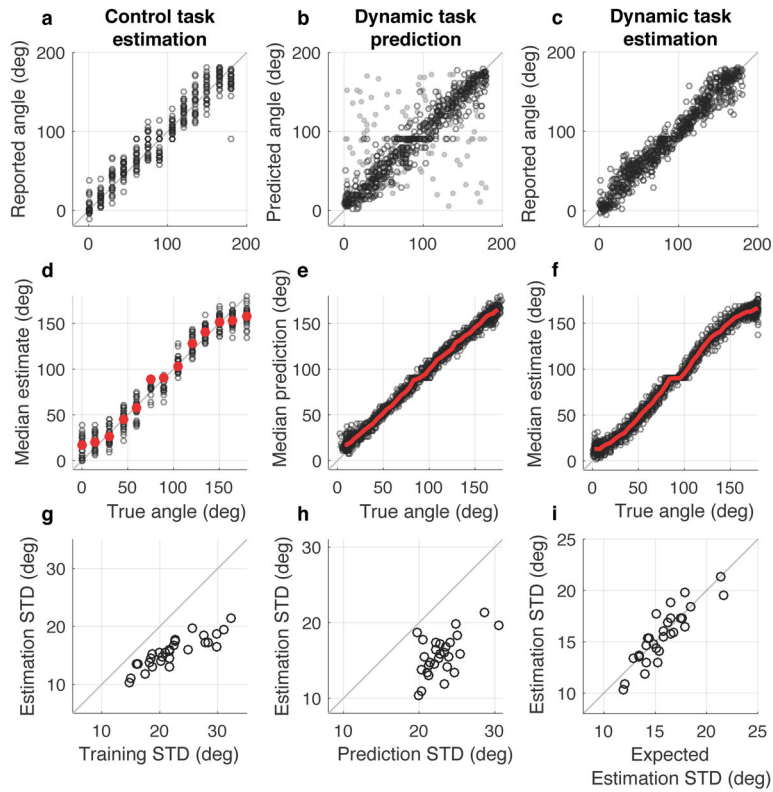


Figure 2. Overall prediction and estimation performance

(a–c) Reported versus true (simulated) sound-source angle for an example subject for: (a) estimations from the control task; (b) predictions from the dynamic task (light gray points indicate change-point trials, on which the probe location was, by design, unpredictable); and (c) estimations from the dynamic task, including all trials. (d–f) Population summaries, plotted as in (a–c), with per-subject median values shown in black and the median of medians shown in red ($n=29$ subjects). For the dynamic tasks, median values were calculated in sliding 20° windows. Non-change-point trials were excluded from the predictions in (e). Note that the subjects' perceptual reports (d and f) were biased slightly towards straight ahead at the far periphery. This bias, which likely reflected learned expectations that sounds were only played in the frontal plane, is accounted for in later analyses (β_5 and β_6 in Eq. 5). (g–i) STD of the perceptual errors from the dynamic task plotted versus the STD of: (g) the perceptual errors from the control task; (h) the prediction errors from the dynamic task; or (i) the expected STD of the perceptual errors, computed from the optimal, reliability-weighted combination of the control perceptual errors and the dynamic prediction errors. Points in g–i represent data from individual subjects. Prediction and perceptual errors were computed with respect to the simulated location of the probe sound.

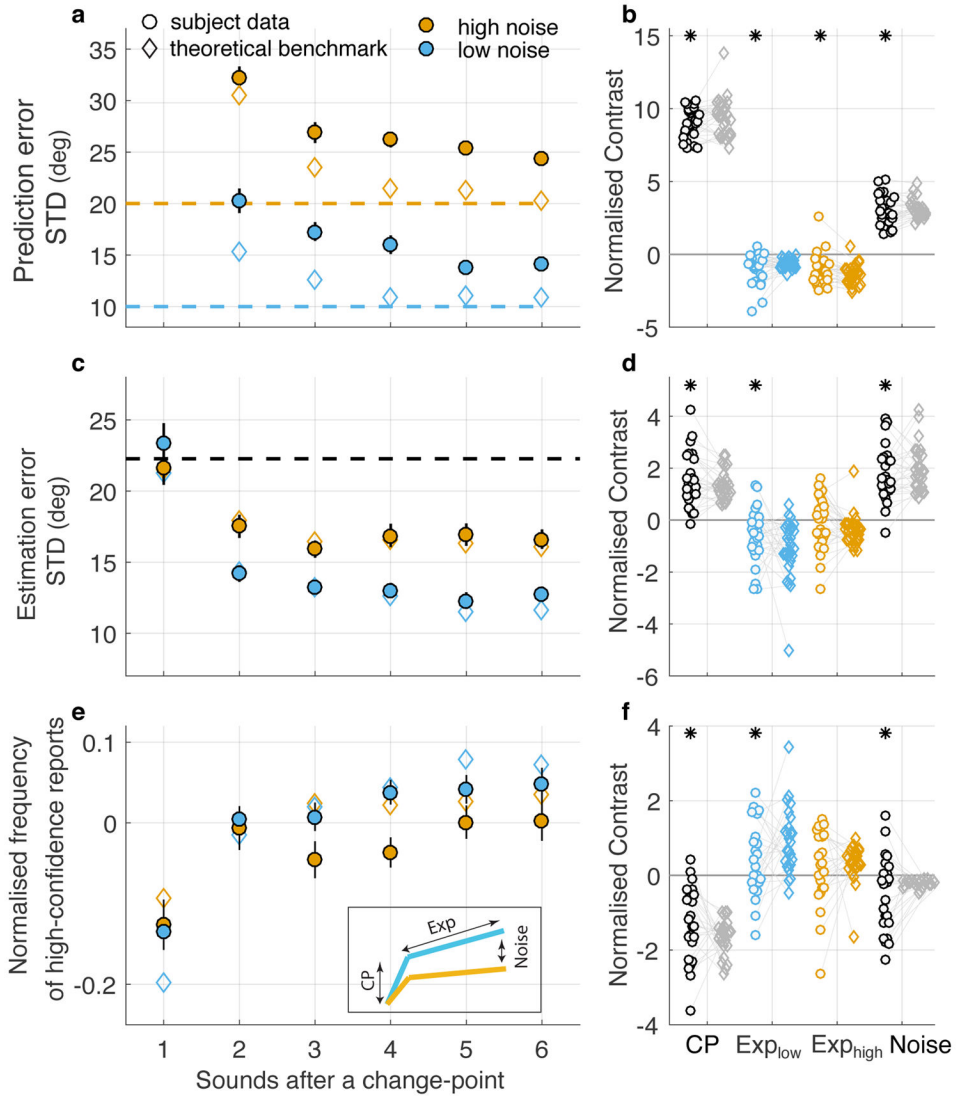


Figure 3. Effects of task dynamics on performance

(a) STD of the subjects' prediction errors (filled circles) as a function of the number of sounds after a change-point (SAC) in the generative mean azimuthal location, plotted separately for the two noise conditions (colors, as indicated; generative STDs are shown as dashed lines). For comparison, prediction-error STDs are shown for an approximately optimal predictive-inference model (open diamonds). Data from change-point trials (SAC=1) are not shown because locations were, by design, unpredictable on those trials. (b) Contrast values from a linear model describing individual subject (circles) and the approximately optimal model (each diamond represents analyses based on the same sound sequence experienced by the subject connected by a line) prediction-error STD in terms of (see inset in e): 1) the difference between change-point and non-change-point trials (CP), 2,3) the linear trend from SAC 2–6 for low-(Exp_{low}) or high-(Exp_{high}) noise trials, and 4) the difference between the two noise conditions (Noise). (c,d) Same conventions as in a,b but for perceptual errors on the dynamic task. Diamonds represent the theoretically predicted

STD of perceptual errors computed from the optimal, precision-weighted combination of the subject- and condition-specific STDs of prior errors (circles in **a**, determined separately for each subject) and the subject-specific estimation-error STDs from the control task (the median value is shown as a horizontal dashed line; see Fig. 2g). (**e,f**) Same conventions as in **a,b** but for the frequency of high-confidence reports relative to overall frequency of high-confidence reports per subject. Diamonds represent the frequency of high-confidence reports corresponding to the theoretical perceptual errors in **c**, computed from the fraction of the theoretical posterior distribution within the confidence window. In **a,c,e**, circles and error bars are mean \pm sem of values measured from all 29 subjects. In **b,d,f**, points are data from individual subjects. Asterisks indicate sign-rank test for H_0 : median value from the subject data=0, $p<0.05$. In each case, paired rank-sum test for H_0 : median difference between subject data and theoretical prediction, $p>0.087$. In all panels, only data from sequences following noticeable change-points (changes in mean of at least twice the generative STD for SAC=1) were included.

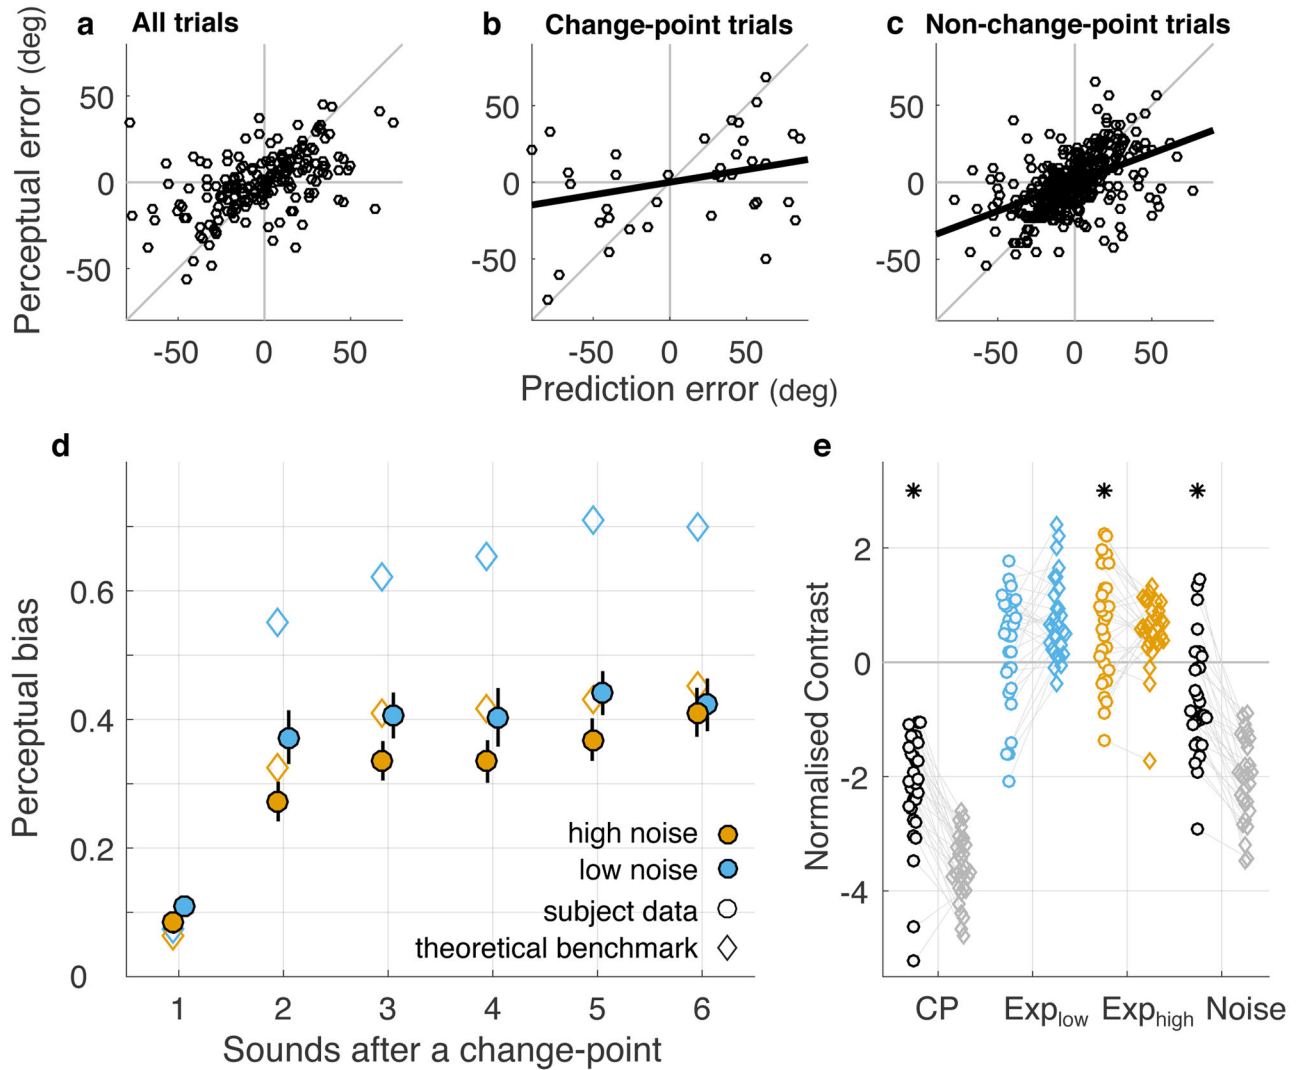


Figure 4. Effects of task dynamics on perceptual bias

(a–c) Example data from a single subject illustrating the quantification of perceptual bias as the slope of the best-fit line to a scatter of the perceptual error versus the prediction error. Slopes close to zero reflect a low perceptual bias (i.e., the percept is unrelated to the prediction), as on change-point trials (b). Slopes closer to unity reflect a higher perceptual bias (i.e., the percept more closely matches the prediction), as on non-change-point trials (c). (d) Perceptual bias as a function of the number sounds after a change-point (SAC) in the generative mean azimuthal location, plotted separately for the two noise conditions (colors, as indicated). Circles and error bars are mean \pm sem of values measured from all 29 subjects. Diamonds indicate the theoretically predicted perceptual bias from an optimal, reliability-weighted combination of the subject- and condition-specific predictions (Fig. 3a) and the subject-specific estimates from the control task (Fig. 2g). (e) Contrast values from a linear model describing individual subject (circles) and model (each diamond represents analyses based on the same sound sequence experienced by the subject connected by a line) perceptual bias in terms of (see inset in Fig. 3e): 1) the difference between change-point and

non-change-point trials (CP), 2,3) the linear trend from SAC 2–6 for low-(Exp_{low}) or high-(Exp_{high}) noise trials, and 4) the difference between the two noise conditions (Noise). Asterisks indicate sign-rank test for H_0 : median value from the subject data=0, $p<0.05$. Paired rank-sum tests for H_0 : median difference between subject data and theoretical prediction, $p<0.01$ for CP, $p=0.16$ for Exp_{low}, $p=0.78$ for Exp_{high}, and $p<0.01$ for Noise. In **d** and **e**, only data from sequences following noticeable change-points (changes in mean of at least twice the generative STD for SAC=1) were included.

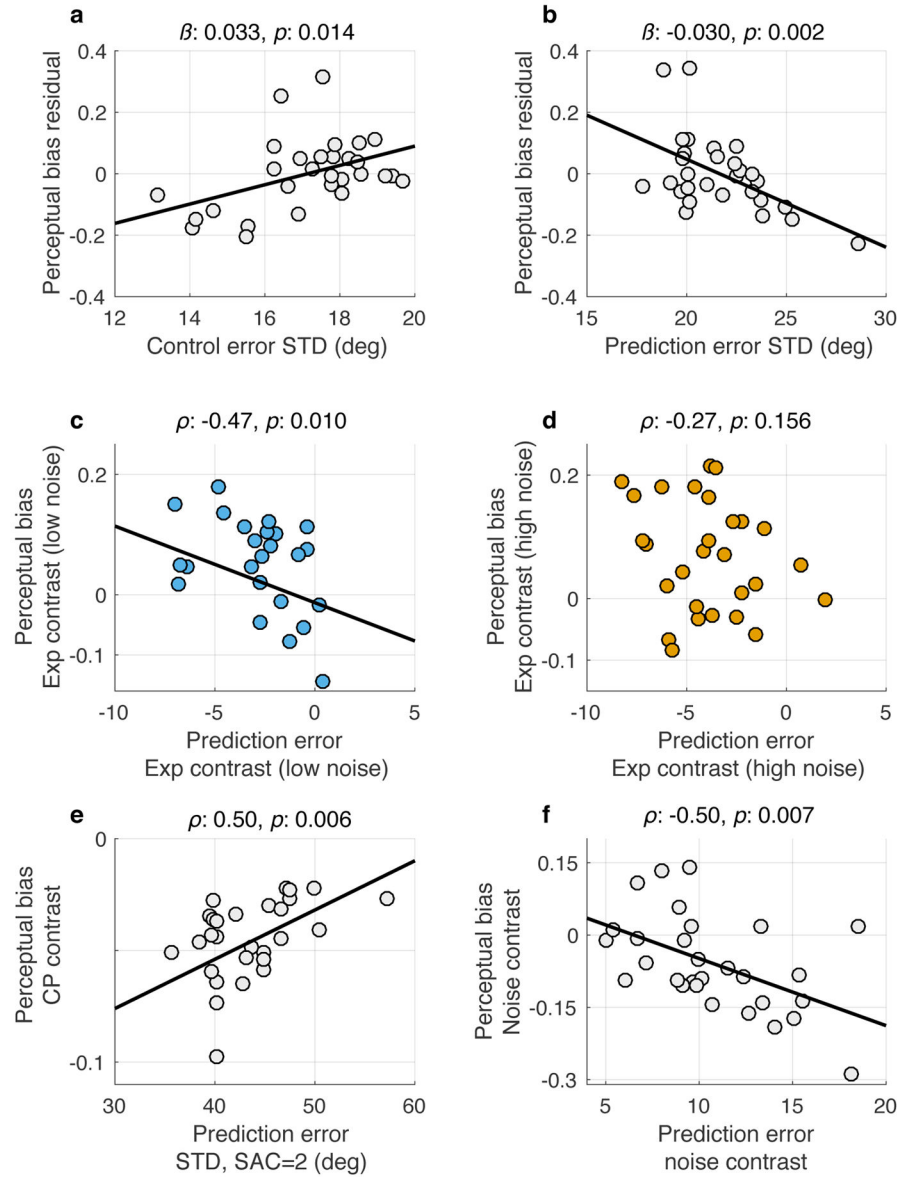


Figure 5. Individual differences in perceptual bias

(a, b) Relationship between overall (mean) perceptual bias and either overall localization ability (STD of perceptual errors on the control task, a) or overall prediction ability (STD of prediction errors from non-change-point trials on the dynamic task, b), after accounting for the other factor (hence “residual”) via linear regression. (c–f) The dependence of perceptual bias on various task conditions, plotted as functions of the dependence of prediction-error STD on the same conditions: c, d) the linear trend from SAC 2–6 in the low-noise (c) and high-noise (d) condition (Exp); e) change-point versus non-change-point trials (CP); and f) high-versus low-noise condition (Noise). In each panel, points represent data from individual subjects. Lines are linear regressions. Only data from sequences following noticeable change-points (changes in mean of at least twice the generative STD for SAC=1) were included.

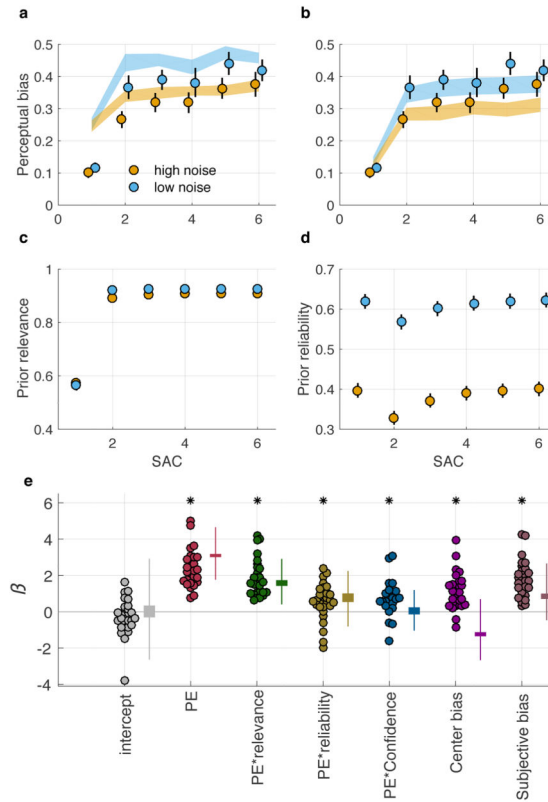


Figure 6. Dynamic modulation of perceptual bias by normative and non-normative factors

(a) Comparison of a parameter-free normative model (ribbons indicate mean \pm SEM simulated perceptual bias for the same task sequences experienced by the subjects) and the subjects' behavior (points and errorbars are mean \pm SEM from 29 subjects), shown as a function of sounds after a change-point (SAC) for the two noise conditions (colors, as indicated). (b) Comparison of the linear model shown in panel e to behavior. Conventions as in panel a. (c,d) Dependence of the normative factors used in both models on task conditions: (c) prior relevance, which measures the probability of the current sound coming from the same distribution as the previous sound; and (d) prior reliability, which measures the anticipated precision of the predictive distribution relative to the likelihood distribution prior to stimulus presentation. (e) Best-fitting parameter estimates from the linear model fit to behavioral data from each subject (points) and to simulations of the parameter-free normative model (thick and thin bars indicate 95% confidence intervals over simulated subjective values and over simulated mean values across subjects, respectively). PE=prediction error. Asterisks indicate coefficients with mean values that differed from zero (t-test, $p < 0.05$).

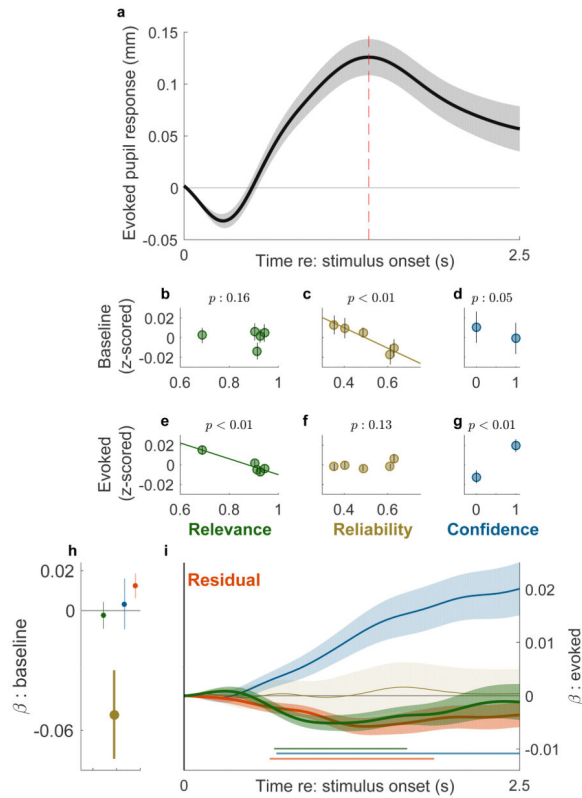


Figure 7. Pupil diameter reflects dynamic modulations of perceptual bias within individual subjects

(a) Mean \pm sem evoked pupil response from 29 subjects, defined as the pupil diameter relative to baseline during the measurement period. Red line indicates the time of the peak mean response (1.38 sec after stimulus presentation). (b–d) Baseline pupil diameter for trials sorted into bins according to relevance (b), reliability (c), and confidence (d). Relevance and reliability were binned in quintiles per subject, then each bin was combined across subjects. Confidence was divided into all trials with a low (0) or high (1) confidence report. Points and errorbars are mean \pm SEM from all values in each bin. (e–g) Same as b–d, but using the pupil diameter measured at the time of the peak response after accounting for the linear baseline dependencies. (h,i) Regression coefficients from a linear model accounting for modulation of baseline pupil diameter (h) or the evoked response (i) at each time-point using as predictors: 1) prior relevance, 2) prior reliability, 3) the upcoming confidence report, and 4) the residual perceptual bias from the linear model in Fig. 6d. Points and error bars in h and lines and ribbons in i represent mean \pm sem of values computed per subject and thus represent within-subject modulations. Points and lines/ribbons corresponding to relevance, reliability, and confidence use the same colors as in (b–g). Bold symbols in h and horizontal lines in i indicate periods for which H_0 : value=0, $p < 0.05$, after accounting for multiple comparisons.

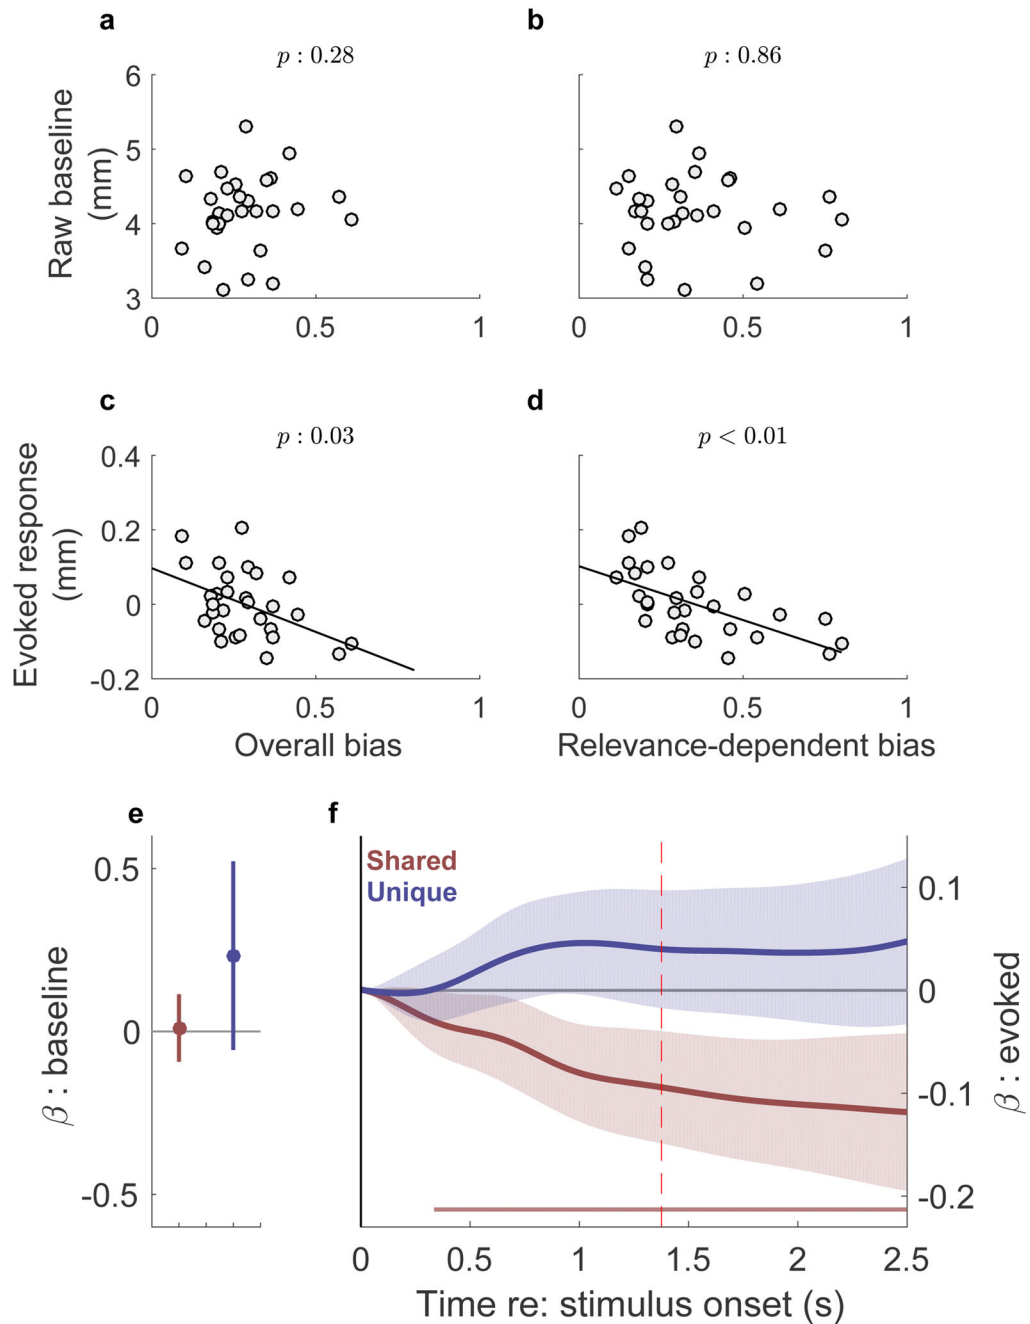


Figure 8. Pupil diameter reflects individual differences in perceptual biases

(a,b) Mean baseline diameter for each subject (points) as a function of the perceptual bias (a; fits to the PE term in Fig. 6e) and relevance-dependent bias (b; fits to the PE*relevance term in Fig. 6e). (c,d) Mean evoked pupil response for each subject as a function of the perceptual bias (a) and relevance-dependent bias (b). Pupil responses were measured at the time of peak response (1.38 sec after stimulus presentation) and orthogonalized to subject baseline pupil measurements. (e,f) Regression coefficients describing the relationship between shared or unique variance (colors, as indicated) in PE and PE*relevance coefficients

from the behavioral model and average baseline (**e**) or stimulus evoked (**f**) pupil diameter. Points and error bars in **d** and lines and ribbons in **e** represent the correlation coefficient and 95% confidence intervals of the estimate and thus represent across-subject modulations. Horizontal lines in **e** indicate periods for which H_0 : value=0, $p<0.05$ after accounting for multiple comparisons.