Adaptive learning through temporal dynamics of state representation

Niloufar Razmi[1,2], Matthew R. Nassar[1,2]

[1]Robert J. & Nancy D. Carney Institute for Brain Science, Brown University, Providence RI 02912-1821, USA

[2]Department of Neuroscience, Brown University, Providence RI 02912-1821, USA

Acknowledgments

We thank Linda Yu, Robert Wilson, Olga Lositsky, Romy Frömer and Cristian Buc Calderon for helpful discussion and comments. This work was supported by R00AG054732 to M.R.N.

Competing interests

The authors declare that no competing interests exist.

Abstract

People adjust their learning rate rationally according to local environmental statistics and calibrate such adjustments based on the broader statistical context. To date, no theory has captured the observed range of adaptive learning behaviors or the complexity of its neural correlates. Here, we attempt to do so using a neural network model that learns to map an internal context representation onto a behavioral response via supervised learning. The network shifts its internal context upon receiving supervised signals that are mismatched to its output, thereby changing the "state" to which feedback is associated. A key feature of the model is that such state transitions can either increase learning or decrease learning depending on the duration over which the new state is maintained. Sustained state transitions that occur after changepoints facilitate faster learning and mimic the evolution of neural patterns in orbitofrontal cortex. In contrast, state transitions after one-off outlier events are short-lived, thereby limiting the impact of outlying observations on future behavior. State transitions in our model provide the first mechanistic interpretation for bidirectional learning signals, such the p300, that relate to learning differentially according to the source of surprising events, and may also shed light on discrepant observations regarding the relationship between transient pupil dilations and learning. Taken together, our results demonstrate that dynamic latent state representations can afford normative inference and provide a coherent framework for understanding neural signatures of adaptive learning across different statistical environments.

Introduction

People and animals exist in complex environments and are often required to update behavior in the face of new information. While standard supervised learning or reinforcement learning models have shown great success in performing particular tasks and explaining general trends in behavior, they lack the flexibility of biological systems, which seem to adjust the influence of new information dynamically, especially in environments that evolve over time (Behrens, Woolrich, Walton, & Rushworth, 2007; Donahue & Lee, 2015; Farashahi, Donahue, Hayden, Lee, & Soltani, 2019; Li, Nassar, Kable, & Gold, 2019; Massi, Donahue, & Lee, 2018; Nassar & Gold, 2010). Recent advances in understanding these adaptive learning behaviors have relied on probabilistic modeling to better understand the computational problems that organisms face for survival in their everyday life (Soltani & Izquierdo, 2019).

Bayesian probability theory has been extensively applied to describing adaptive learning algorithms in changing environment to provide normative accounts for learning behavior. Probabilistic models prescribe learning that is more rapid during periods of environmental change and slower during periods of stability (Adams & MacKay, 2007; Behrens et al., 2007; Nassar & Gold, 2010; Wilson, Nassar, & Gold, 2010). These models have provided insight into why people seem to adjust learning according to their level of uncertainty (Browning, Behrens, Jocham, O'Reilly, & Bishop, 2015; Muller, Mars, Behrens, & O'Reilly, 2019) and the probability with which an observation reflects a changepoint (Adams & MacKay, 2007; Nassar, Wilson, Heasly, & Gold, 2010). In this framework, the human brain is viewed as implementing an optimal learning algorithm that embodies the statistical properties of the world it operates in (Meyniel & Dehaene, 2017; O'Reilly, 2013).

While probabilistic modeling provides an ideal observer account for many of the adjustments in learning rate observed in humans and animals (Behrens et al., 2007; Nassar, Bruckner, & Frank, 2019; Nassar &

Gold, 2010), it has thus far failed to clarify the underlying neural mechanisms. One issue is that exact Bayesian inference can be closely approximated by many qualitatively different algorithms (Bernacchia, Seo, Lee, & Wang, 2011; Farashahi et al., 2017; Iigaya, 2016; Mathys, Daunizeau, Friston, & Stephan, 2011; Nassar et al., 2010; Wilson, Nassar, & Gold, 2013; A. J. Yu & Dayan, 2005). One such approximation that relies on a single dynamic learning rate can capture behavior across a wide range of statistical environments (Nassar, Waltz, Albrecht, Gold, & Frank, 2021). However, direct implementation of this model requires a dynamic learning rate signal that is invariant to statistical context. Such a learning rate signal has yet to be observed in the brain, despite several attempts to do so (D'Acremont & Bossaerts, 2016; Li et al., 2019; Nassar, McGuire, Ritz, & Kable, 2019). In contrast, brain signals that predict more learning in changing environments (Behrens et al., 2007; Jepma et al., 2016; McGuire, Nassar, Gold, & Kable, 2014; Nassar et al., 2010; O'Reilly et al., 2013), do not do so consistently across different statistical conditions (D'Acremont & Bossaerts, 2016), and in at least some cases positively relate to learning in one condition and negatively relate to learning in another (Nassar, Bruckner, et al., 2019; Nassar et al., 2012; O'Reilly et al., 2013). Other approximations to normative learning have been more closely connected to specific neural signals, but fail to capture the range of behaviors displayed by people, for example the ability to immediately discount past experience after a changepoint (Bernacchia et al., 2011; Farashahi et al., 2017; Mathys et al., 2011), or the ability to calibrate learning across different statistical environments (D'Acremont & Bossaerts, 2016). Thus, while there has been a good deal of work exploring potential neural mechanisms for adaptive learning, existing algorithms lack the generality necessary to capture the range of human behavior and its underlying brain signals.

Here we explore such a generalized framework for adaptive learning using a feed-forward neural network model that maps an internal context representation (which can be thought of as its "mental context" and serves to organize learning across events much like the state in a reinforcement learning model) onto a continuous action space in order to perform a predictive inference task. We show that the effective learning rate of the model is proportional to the rate at which its internal context evolves in time, and that better model performance can be achieved when context transitions are discontinuous and elicited by surprising events. Furthermore, we show that context transitions can speed learning after changepoints, or slow them after oddball events, assuming appropriate state transitions occur between trials (L. Yu, Wilson, & Nassar, 2020). Our model produces these behaviors without an explicit representation of learning rate, and instead relies on an *internal context* that transitions rapidly after surprising events much like patterns of activity previously observed in orbitofrontal cortex (Nassar, McGuire, et al., 2019). Furthermore, it requires *context transition signals* that bidirectionally affect learning according to statistical context (changepoint versus oddball), providing a mechanistic explanation for feedback-locked P300 signals that show the same complex relationship to learning (Nassar, Bruckner, et al., 2019), and potentially shedding light on discrepant relationships between pupil diameter and learning that have been reported (compare Nassar et al., 2012 to O'Reilly et al., 2013). Taken together, our results support the idea that adaptive learning behavior emerges through abrupt transitions in mental context. Under this view, we argue that learning rate dynamics emerge as a consequence of changes in the internal representations to which learning is bound, and that the brain has no need to represent a global learning rate signal directly.

Results

Predictive inference task to measure adaptive learning

In order to test whether changes to latent state representations can facilitate adaptive learning behavior we modeled a predictive inference task designed to measure adaptive learning in humans (figure 1)(McGuire et al., 2014). In the task a helicopter, which is hidden behind clouds, drops visible bags containing valuable contents from the sky (figure 1a, right). On each trial, the subject moves a bucket to the location where they believe the helicopter to be, such that they can catch potentially valuable bag contents. Subjects can move the bucket to a new position on each trial in order to update and improve their prediction (figure 1a, left; figure 1b orange arrow). In the "changepoint" variant of the task, bag locations are sampled from a Gaussian distribution centered on the helicopter, which occasionally relocates to a new position on the screen. Such abrupt transitions in helicopter location lead to changes in the statistical context defining the bag locations (context shifts), which can be inferred by monitoring the size of prediction errors (figure 1b, red arrow). Therefore, the helicopter position is a dynamic latent variable that needs to be inferred from noisy observations (i.e. dropped bags) on each trial to yield optimal task performance. Previous work has shown that behavior of human subjects can be captured by a normative learning model that relies on a dynamic "learning rate" adjusted from trial-to-trial according to changepoint probability (CPP) and uncertainty (figure1b&c), but failures to identify neural signals that reflect this dynamic learning rate consistently across conditions cast doubt on its biological relevance (D'Acremont & Bossaerts, 2016; Nassar, Bruckner, et al., 2019; Nassar et al., 2012; O'Reilly et al., 2013). Here we explore whether normative learning may instead be achieved in the brain by a neural network that undergoes dynamic transitions in the "mental context" to which associates are bound, thereby adjusting where information is stored, rather than the degree to which storage occurs.

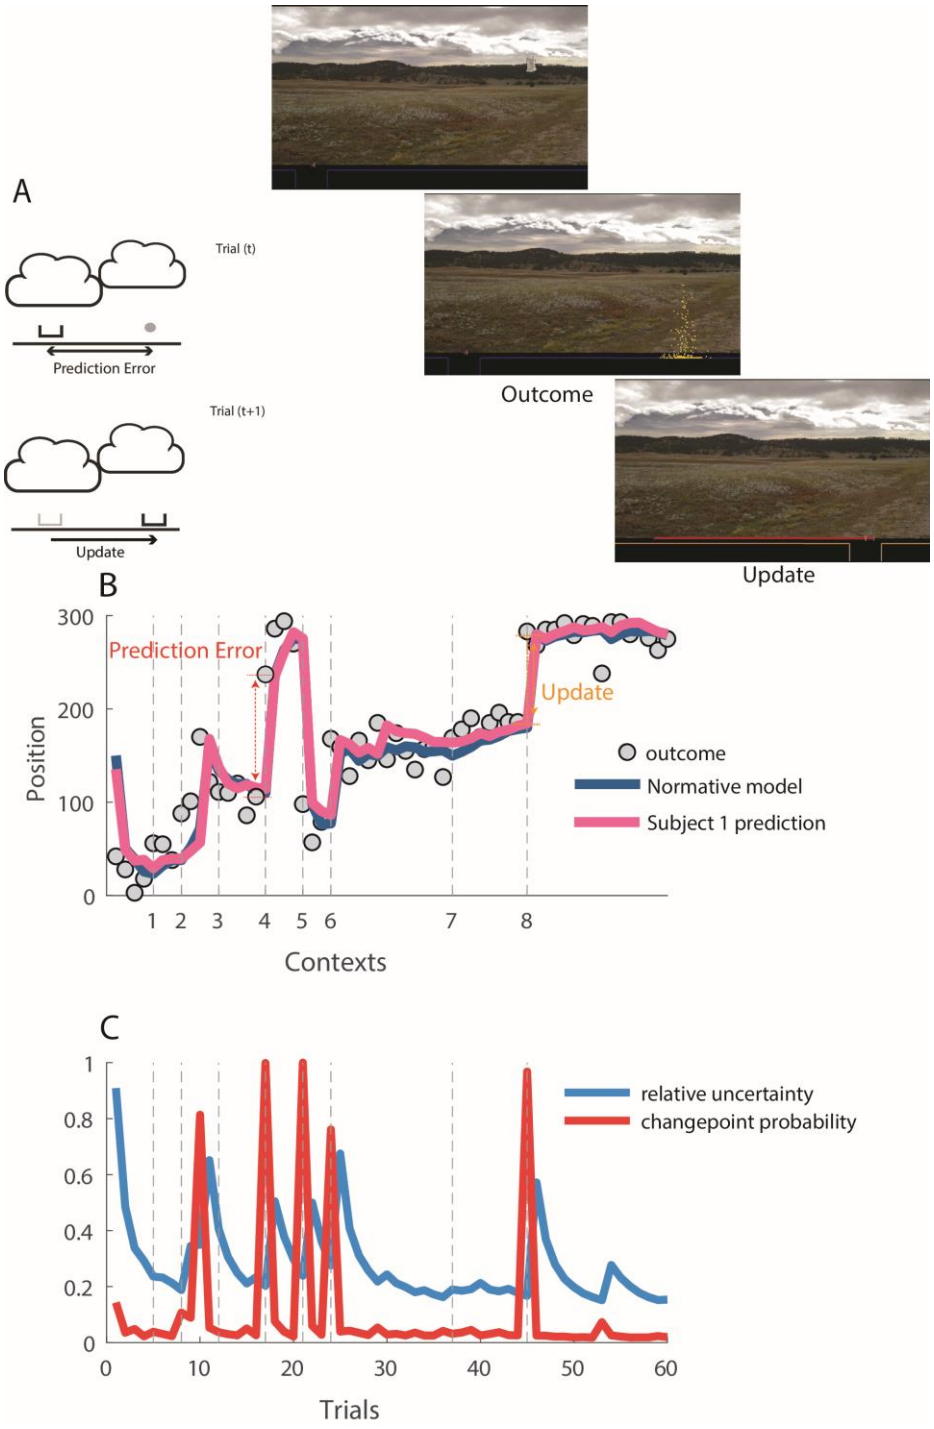

Figure 1: Predictive inference task to measure dynamics of adaptive learning.

A) Schematic Illustration of task design and screenshots of the task. Human subjects place a bucket at the bottom of the screen to catch a bag of coins subsequently dropped from a hidden helicopter. After observing the outcome at the end of each trial (prediction error), subject could improve their response by moving the bucket to a new location most likely to receive the bag in the subsequent trial (update). In the changepoint condition, the helicopter on occasion moves to a new location. B) The sequence of bag locations (outcome) is plotted over context shift designated by sudden changes in outcome mean (dotted vertical lines). Bucket placements made by a subject (pink) and normative model (navy) are shown with a representation of an example prediction and outcome. [Prediction = outcome (t) –

estimate (t) and Update = estimate (t+1) – estimate (t)]. (C) The learning rate describing update magnitude of normative model depends on two factors, changepoint probability (CPP; red) and relative uncertainty (RU; blue), which combine to prescribe learning that is highest at changepoints (CPP) and decays slowly thereafter (RU).

A neural network test bed for exploring adaptive learning

To examine how normative updating could be implemented in a neural network, we devised a two-layer feedforward neural network in which internal representations of context are mapped onto bucket locations by learning weights using a supervised learning rule (figure 2b; see methods). Units in the output layer of the network represent different possible bucket locations in the predictive inference task where the activity of the output layer corresponding to the most likely helicopter position in the models point of view and a linear readout of this layer is used to guide bucket placement and to make a prediction for the next trial. After each trial, a supervised learning signal corresponding to the bag location is provided to the output layer and weights corresponding to connections between input and output units are updated accordingly.

The input layer of our model is designed to reflect the mental context of learning, and its activity is given by a Gaussian activity bump with a mean denoting the position of the neuron with the strongest activity and a standard deviation denoting the width of the activity profile. The primary goal of this work is to understand how changes to the mean of the activity bump, across trials, affect learning within our model. Since the input layer of the network reflects mental context, it does not receive any explicit sensory information, and we can manipulate its activity across trials to provide a flexible test bed for how different task representations (i.e. mental context dynamics) might affect performance of the model. In the simplest case, a non-dynamic network, the mean of the activity bump in the input layer is constant across all trials -- reflecting learning onto a fixed "state". A slightly more complex mental context might be one that drifts slowly over time, such that the mean of the activity bump changes a fixed amount from one trial to the next leading trials occurring close in time to be represented more similarly. In this case, learning would occur onto an evolving temporal state representation. In a more complex (but maybe more intuitive) case, the subset of active neurons in the input layer could correspond to the current "helicopter context" (figure 1b), or period of helicopter stability. In this case, the mean of the activity bump would only transition on trials where the helicopter changes position and thus could be thought of as representing the underlying latent state of the helicopter (e.g. this is the third unique helicopter position I have encountered) – albeit without any explicit encoding of its position.

Context shifts facilitate faster learning

We first examined performance of models in which the mean of the input activity bump transitioned by some fixed amount on each trial. This set included 0 (fixed stimulus representation), small values in which nearby trials had more similar input activity profiles (timing representation) and extreme cases where there was no overlap between the input layer representations on successive trials (individual trial representation). We defined the fixed shift in the mean of the activity profile as the "context shift" of our model. This shift is depicted in figure 2c as the nodes shown in "hot colors" (i.e. active neurons) in the input layer of the neural network moving to the right; Note how the size of rightward shift in the schematic neural network is constant in all four trials shown. We used increments starting from zero (the same input layer population) to a number corresponding to a complete shift (completely new population) in each trial. Learning leads to fine tuning of the weights by strengthening connections between active input neurons and the output neurons nearby the outcome location (bag position) on each trial. We observed that moderate shifts of in the input layer (context shifts) led to the best performance in our task (figure 2e), and that the effective learning rate describing the model's behavior monotonically scaled with context shift (figure 2d). The higher performance of the models with moderate context shift on average over all trials might be thought of intuitively as the classic trade-off between stability and flexibility in delta-rule learning models. (figure 2e, red line) where a higher learning rate leads to better performance in response to environmental change, whereas a smaller learning rate has a performance advantage in a stable but noisy environments. More concretely, when the helicopter remains in the same location minimal network context shifts improve performance by pooling learning over a greater number of bag locations in order to more closely approximate their mean, but maximal context shifts improve performance after changes in helicopter location by reducing the interference between outcomes before and after the helicopter relocation.

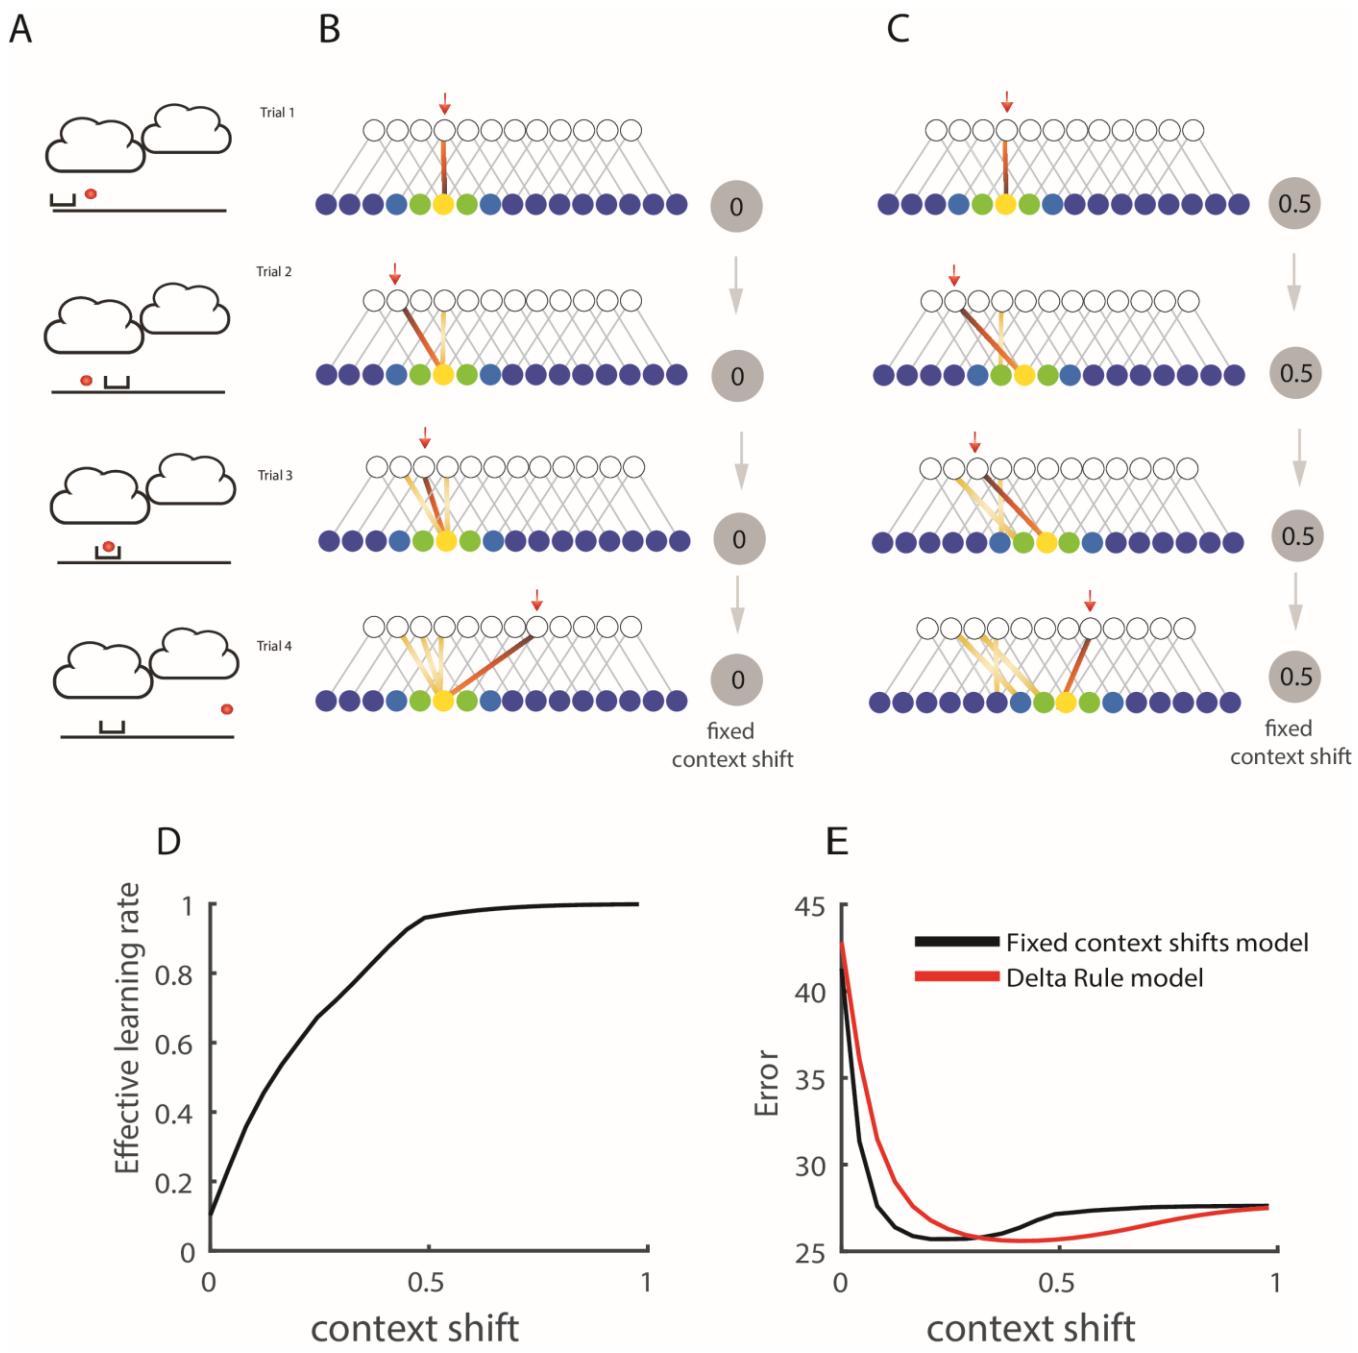

Figure 2: A neural network with fixed context shifts can approximate any constant learning rate. A-C) Network structure and weight updates for two fixed context shift models (B, C) are depicted across four example trials of a predictive inference task (A). For all networks, feedback was provided on each trial corresponding to the observed bag position (circle in panel A, red arrow in B&C) and weights of network were updated using supervised learning. Only a subset of neurons (circles) and connections between them (lines) are shown in neural network schematic. Activation in the input layer was normally distributed around a mean value that was constant in (B) and shifted rightward by a fixed amount on each trial in (C) (context shift). Learned weights (colored lines) were all assigned to the same input neuron when context shift was set to zero (B) but assigned to different neurons when the context shift was substantial (C). D) The effective learning rate (ordinate), characterizing the influence of an unpredicted bag position on the subsequent update in bucket position, increased when the model was endowed with faster internal

context shifts (abscissa). E) Mean prediction error (ordinate) was minimized by neural network models (black line) that incorporated a moderate rate of fixed context shifts (abscissa). Mean error of a simple delta rule model using various learning rates is shown in red (x-axis values indicate the context shift equivalent to the fixed delta rule learning rate derived from panel D). Note that the neural networks with the fixed context shift achieve similar task performance to more standard models delta-rule models that employ a constant learning rate.

Dynamic context shifts can improve task performance

Inspired by the observed relationship between context shift and accuracy, we next modified the model to include discontinuous context shifts in order to optimize performance. In principle, based on the intuitions above, we might improve on our fixed context shift models by only shifting the activity profile of the input layer at a true context shift in the task (i.e. allow the input layer to represent the latent state). Since such a model requires pre-existing knowledge of changepoint timings we refer to it as the ground truth model (figure 3, top). Indeed, we observed that the ground truth model performance is better than any of the fixed context shift models both after changepoints (figure 3c) and during periods of stability (figure 3d) yielding better overall performance (figure 3e).

Needless to say, the brain does not have access to perfect information regarding whether a given trial is a changepoint or not. Is it possible to make a more realistic version of this optimal model, utilizing information that the brain does have access to? To answer this question, we built models that infer changepoint probability based on experienced prediction errors. We built two versions of this model, one that computed changepoint probability (CPP) explicitly according to Bayes rule (Nassar & Gold, 2010), and one that approximated CPP according to the mismatch between output activity in the network and the observed outcome (i.e. supervised signal). Both models achieved good performance after changepoints by elevating context shifts (figure 3c) and during periods of stability by reducing context shifts (figure 3d), although they understandably achieved slightly worse performance than the ground truth model (figure 3e). These results were consistent across different noise conditions (figure 3 – 2).

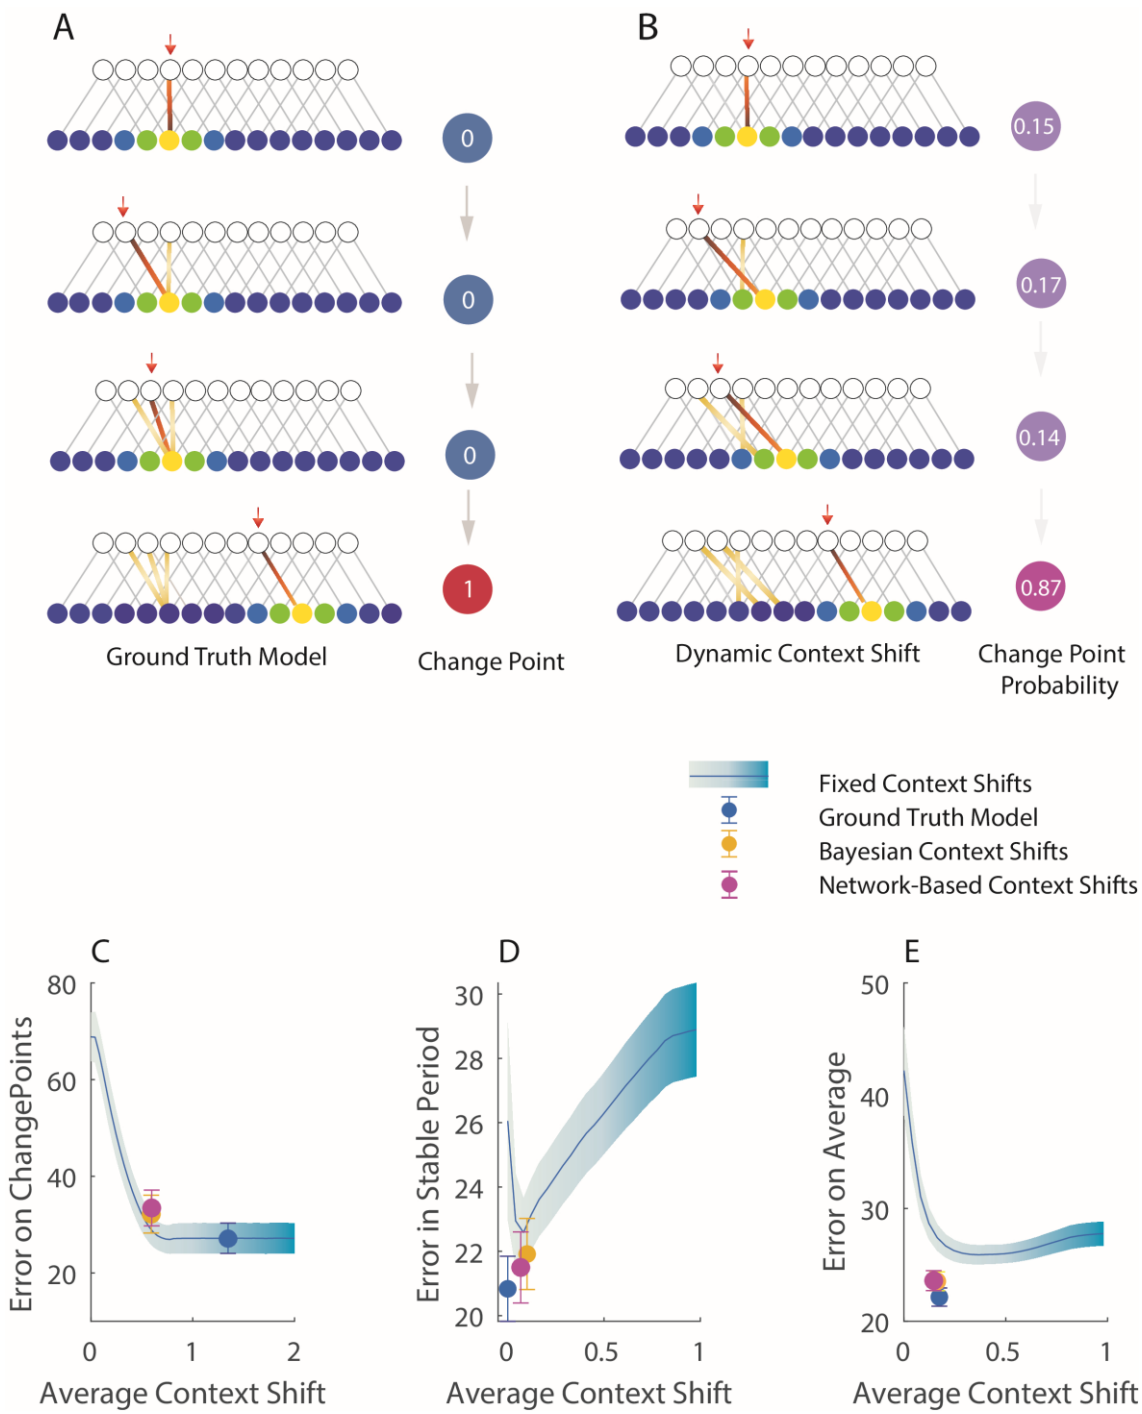

Figure -3: Dynamic context shifts facilitate better performance.

A) Schematic diagram of ground truth model network which is provided with objective information about whether a given trial is changepoint or not and uses that knowledge to shift the context only on changepoint trials. B) The dynamic context shift network uses a subjective estimate of changepoint probability based a statistical model (Bayesian) or the network output (Network-based) to adjust its context shift on each trial. All of these models shift context to a greater degree on changepoint trials (bottom row) than on non-changepoint trials (top 3 rows). C)

Performance on trials immediately following a changepoint was best for models employing the largest context shifts. The ground truth model was also able to achieve this level of performance by undergoing large context shifts specifically on changepoint trials (blue dot) and dynamic context shift models were also able to approach this level of performance by increasing context shifts in proportion to subjective measures of changepoint probability (yellow & pink). D) Smallest errors on trials during periods of stability (> 5 trials after changepoint; ordinate) were achieved by models with less context shifting (abscissa). All dynamic context shift models (ground truth, Bayesian, network-based) made relatively small context shifts for stable trials, leading to good performance. E) Across all trials, subjective dynamic context shift models yielded better average performance than the best fixed context shift model ($Bayesian\ Context\ Shift: t = 21.9. df = 31. p < 10^{-16}$Network − Based Context Shift: $t = 20.48. df = 31. p < 10^{-16}$, and approached the performance of the ground truth model.

Dynamic context shifts capture key behavioral and neural signatures of adaptive learning in humans.

Not only was the dynamic context shift model able to outperform fixed context shift models, it did so by capturing behaviors that are observed in people. The model updates predictions according to prediction errors, but relies more heavily on prediction errors from certain trials (figure 4a). We can quantify the effective learning rate as the slope of the relationship between the model update and its previously observed prediction error in order to compare the behavior of different models. Looking at this *effective learning rate* in more detail, we observe that immediately after a changepoint learning rate becomes maximal for the ground truth model and dynamic context shift models while gradually decreasing during the more stable periods. (figure 4c). A regression analysis, previously used in explaining humans' responses in a similar task, determined the contribution of changepoint probability and relative uncertainty to updates in each model (figure 4d) and indicated that, like human subjects, the dynamic context shift model learned more rapidly during periods of change or uncertainty (figure 4d). Note that most participants (gray dots in figure 4d) fall between the range of behaviors spanning from the fixed context shift model (green dot) and the network based context shift model (pink dot) suggesting that people may use a mental context representation that is somewhere between a purely temporal one (i.e. fixed context shift) and our subjective approximation of latent state (network-based context shift). Taken together, these results suggest that our dynamic context shift models capture the primary behavioral features of adaptive learning in changing environments – but unlike previous such models they do so by adjusting an internal context, rather than a learning rate per se.

These context adjustments provide a potential explanation for rapid changes in the patterns of activation in orbitofrontal cortex that are observed during periods of rapid learning (OFC; Nassar, McGuire, et al., 2019). This phenomenon was previously observed by looking for representations that evolved more rapidly during periods of rapid learning while controlling for other features of the task (eg. predictions, outcomes, or higher level computational variables). However, one important question left open by previous work is why such representations exist at all; as they would not be necessary for implementing existing models of adaptive learning (Nassar et al., 2012, 2010). Given that the dynamic context shift model accomplishes faster learning by changing the context representations, and that orbitofrontal cortex is thought to reflect representations of internal context to guide reinforcement learning (Schuck, Cai, Wilson, & Niv, 2016; Wilson, Takahashi, Schoenbaum, & Niv, 2014), we asked whether our context layer might show trial-to-

trial patterns of representation consistent with what has been observed in OFC – thereby providing a function to the previously observed phenomenon.

We used the RSA approach applied in (Nassar, McGuire, et al., 2019) to create a dissimilarity matrix reflecting differences in the input layer activation across pairs of trials for our dynamic context shift model (figure 5a&b). By using the activity profile of the input layer the dynamic context shift model we were able to obtain a pattern of dissimilarity for adjacent and nonadjacent trials that is nearly identical to the RSA matrix that was used to identify abrupt representational changes in periods of faster learning in the orbitofrontal cortex of humans (figure 5c&d). This, in simpler terms, means that two trials within the same context have more similar representations than two trials from different contexts. Such representations are not necessary for models that adjust learning rate directly, since these models use a dynamic learning rate to directly translate prediction errors into updated predictions (Nassar et al., 2010). Taken together, these results are consistent with OFC encoding the dynamic context signal that is used in our model to adjust the rate of learning.

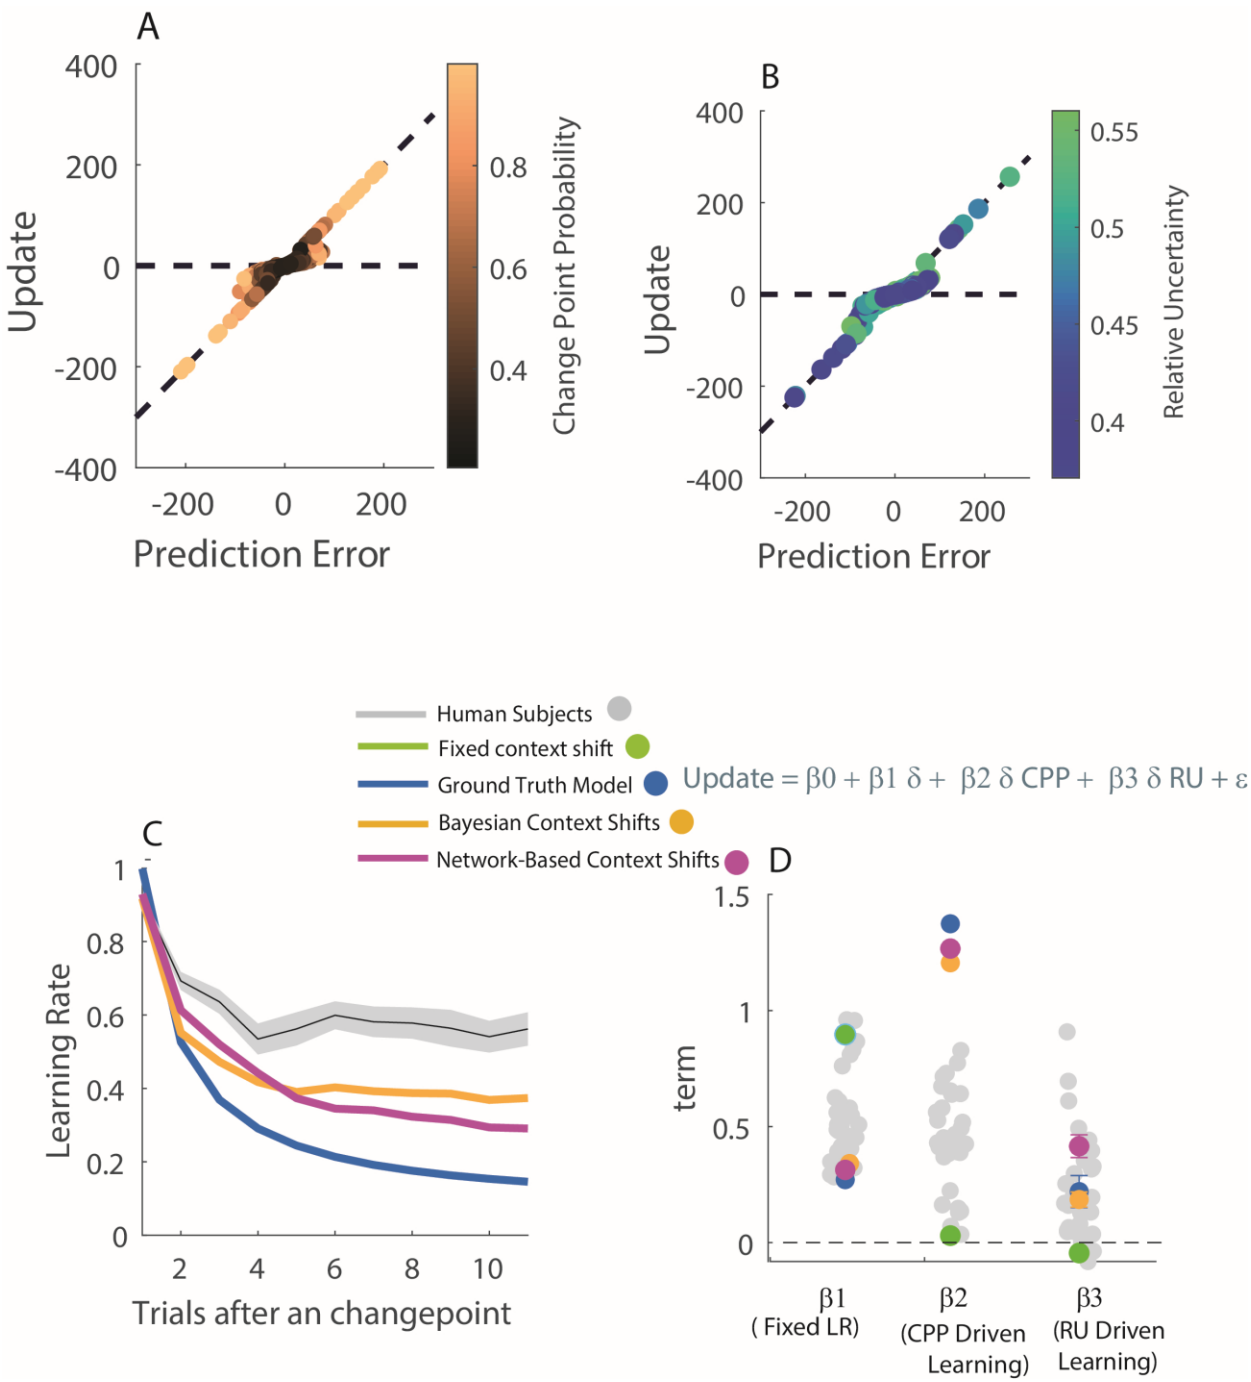

Figure 4: Dynamic context shifts facilitate adaptive learning. A&B: Dynamic context shift model single trial update (ordinate) is plotted against prediction error (abscissa) for each single trial of a simulated session with points colored according to the normative changepoint probability (A) and relative uncertainty (B). C) Effective learning rate as function of trial after change point. Both ground truth and dynamic context shifts models show adjustments in their effective learning rate relative to changepoints, with the dynamic context shift model closely approximating human subjects (gray). D) Coefficients from a regression model (top equation) fit to single trial updates in order to characterize the degree of overall learning (fixed LR), adjustments in learning at likely changepoints (CPP Driven Learning), and adjustments in learning according to normative uncertainty (RU driven learning). The grey circles represent data from human subjects doing the same task.

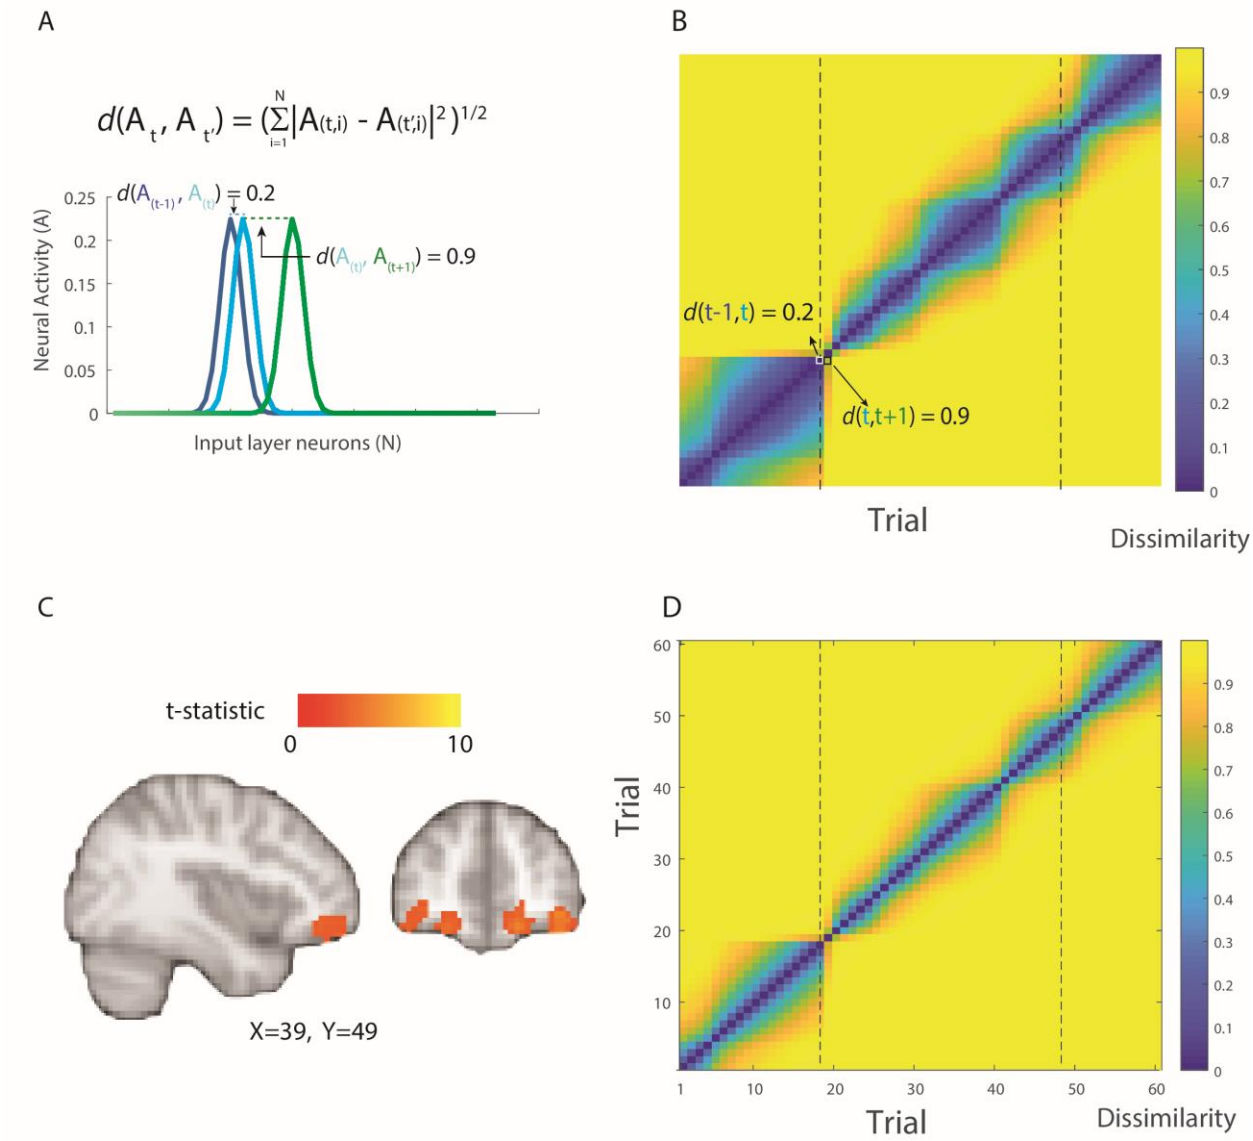

Figure 5: Dissimilarity in input layer representations across trials is increased across changepoints and resembles patterns of activity observed in orbitofrontal cortex.

(A) Dissimilarity in the input representation between pairs of trials was computed according to the Euclidean distance between those trials in the space of population activity. Trial (t) is an example changepoint. B) The dissimilarity matrix represent that input layer activity for trials within the same context (i.e. trial (t-1) and trial (t)) are more similar than for consecutive trials belonging to two different contexts (trial t and trial (t+1)). C) Regions within OFC display patterns of voxel activation that transition more rapidly during periods of rapid learning (reproduced from Nassar et al. 2019, Journal of Neuroscience). D) Dissimilarity matrix used for fMRI analysis that identified the brain regions in (C) for the same example block that were used to characterize the dissimilarity in the input layer of the dynamic context shift model (B). Note the similarity between the two matrices (B&D), and in particular that in both cases trials separated by a changepoint (depicted by vertical dotted lines) are maximally dissimilar to one another. Across subjects the correlation between adjacent trial elements in the two matrices was consistently positive (mean[SEM] Pearson correlation = 0.884[0.011]).

Dynamic context shifts can reduce learning from oddballs

In order to understand how dynamic context shifts might be employed to improve learning in an alternate statistical environment we considered a set of "oddball" generative statistics that have recently been employed to investigate neural signatures of learning (D'Acremont & Bossaerts, 2016; Nassar, Bruckner, et al., 2019). In the oddball condition, the mean of the output distribution does not abruptly change but instead gradually drifts according to a random walk. However, on occasion a bag might drop in an unexpected location uniformly sampled across the width of the screen, and is thus is an outlier unrelated to both past and future outcomes. In the presence of such oddballs, large prediction errors should lead to less, rather than more, learning. This normative behavior has been observed in human subjects (D'Acremont & Bossaerts, 2016; Nassar, Bruckner, et al., 2019; Nassar & Troiani, 2020).

In order to extend our model to this condition we created a network that adjusted active input units based on the oddball condition transition structure and ground truth information regarding whether or not a given trial is an oddball (figure 6a right). Throughout the experiment, the model would shift the context with a small constant rate, corresponding to the drift rate in the generative process (i.e. the helicopter position slowly drifting from trial to trial) but on oddball trials, the model would undergo a large context shift to ensure that the oddball trial assigned learning to a non-overlapping set of input layer neurons. The model was also endowed with knowledge of the transition structure of the task, which includes that oddballs are typically followed by non-oddball trials, and as such, the input layer activity bump would transition to its previous non-oddball location subsequent to learning from the oddball outcome (L. Yu et al., 2020). Consequently, the results of the oddball trials would not be stored in the same context as the ordinary trials, but predictions would always be made from the previous "non-oddball" context.

Similar to the changepoint condition, we also created versions of the model in which oddballs were inferred using either a Bayesian inference model or the activity profile of the output units. Oddball probabilities (computed either from the normative model or the network's output activity itself) were then used to guide transitions of the active input layer units. In these models the probability of an oddball event drove immediate transitions of the active input layer units to facilitate storage of information related to oddballs in a separate location, but subsequent predictions were always made from the input units corresponding to the most recent non-oddball event (plus a constant expected drift). These models achieved significantly better overall performance than the best fixed context shift model and similar performance to the ground truth context shift model (figure 6d). It is noteworthy that the dynamic context shift model approximates normative behavior in changepoint conditions, where it increases learning by sustaining the new context but decreases learning from oddball trials by immediately abandoning the new context in favor of the more "typical" one.

Dynamic context shifts explain bidirectional learning signals observed in the brain

A primary objective in this study was to identify the missing link between the algorithms that afford adaptive learning in dynamic environments and their biological implementations. One key challenge to forging such a link has been the contextual sensitivity of apparent "learning rate" signals observed in the brain. For example, in EEG studies the P300 associated with feedback onset positively predicts behavioral adjustments in static or changing environments (Fischer & Ullsperger, 2013; Jepma et al., 2018, 2016), but negatively predicts behavioral adjustments in the oddball condition that we describe above (Nassar, Bruckner, et al., 2019). These bidirectional relationships are strongest in people who adjust their learning strategies most across conditions, and persist even after controlling for a host of other factors related to

behavior, suggesting that they are actually playing a role in learning, albeit a complex one (Nassar, Bruckner, et al., 2019).

Here we propose that the P300 does not reflect learning rate, but the need for a context shift implied by the observed trial outcome. Thus, a stronger P300 signal would produce a larger context shift, which would result in a stronger partition between current learning and previously learned associations. In changing environments, this would effectively increase learning, as it would decrease the degree to which prior experience is reflected in the weights associated with the currently active input units. In the oddball environment, where context changes *prevent* oddball events from affecting weights of the relevant input layer units, we would make the opposite prediction. We formalized this idea by measuring the effective learning rate in the dynamic context shift models for bins of trials sorted according to the context shift that was experienced on them. As displayed in figure 6e, the measurement network-based context shifts employed by our model reflect the same relationship to effective learning rate as the P300 signal in human behavioral results: context shifts in the changepoint condition correspond to faster learning while the same context shift signal in the oddball condition negatively corresponds to learning rate (figure 6f). Thus, our model suggests that the P300 may relate to learning indirectly by signaling or promoting transitions in an internal representation of latent state.

Relationship between context shifts and pupil diameter response:

One major theory of learning has suggested that adaptive learning is facilitated by fluctuations in arousal mediated by the LC/NE system (A. J. Yu & Dayan, 2005). This idea has been supported by evidence from transient pupil dilations, which in animals are linked to LC/NE signaling (Joshi & Gold, 2020; Reimer et al., 2016), and are positively related to learning in changing environments (Nassar et al., 2012). Nonetheless, these results are difficult to interpret in light of another study that employed both changepoints and oddballs and observed the opposite relationship between pupil dilation and learning (O'Reilly et al., 2013). In contrast to learning theories, another prominent theory has suggested that the LC/NE system plays a role in resetting ongoing context representations (Network reset hypothesis; Bouret & Sara, 2005), which maps well onto the context shift signals that our model requires to adjust effective learning rates.

Our model provides a framework to examine the possible implications of the network reset hypothesis for the relationship between pupil diameter and learning. Specifically, we consider the possibility that LC/NE system is related to the instantaneous context shifts in our model, and that pupil dilations occur as a delayed and temporally smeared version of this LC/NE signal. In this framework we might consider two distinct influences on the pupil diameter. First, the context shifts elicited by observations that deviate substantially from expectations, which might reflect either changepoints or oddballs depending on the statistical context (figure 6a&b). Second, the context shift required to "return" to the previous context after a likely oddball event, which must occur after processing feedback from a given trial, but before the start of the next trial (figure 6a&c).

Thus, we modeled pupil diameter as the degree of context shift predicted by our model in two discrete events occurring at specific time points after outcome presentation and before the start of the subsequent trial. Resulting discrete signals were then convolved with a gamma probability density function to produce simulated pupil responses for three types of trials: Neutral (no-change), changepoints and oddballs (Figure 7c). We then used measures of learning and surprise from the Bayesian dynamic learning model in a regressor to see how surprise (changepoint/oddball probability) and learning (trial-wise learning rate) are related to our simulated pupil responses. Our results predict an initial positive correlation with surprise and a later negative correlation with learning. These results provide a qualitative agreement with data from

Nassar 2012 showing that transient pupil dilations reflect changepoint probability in a context with changepoints, but also capture the relationships to early relationship to surprise and late negative relationship to learning observed in O'Reilly 2013.

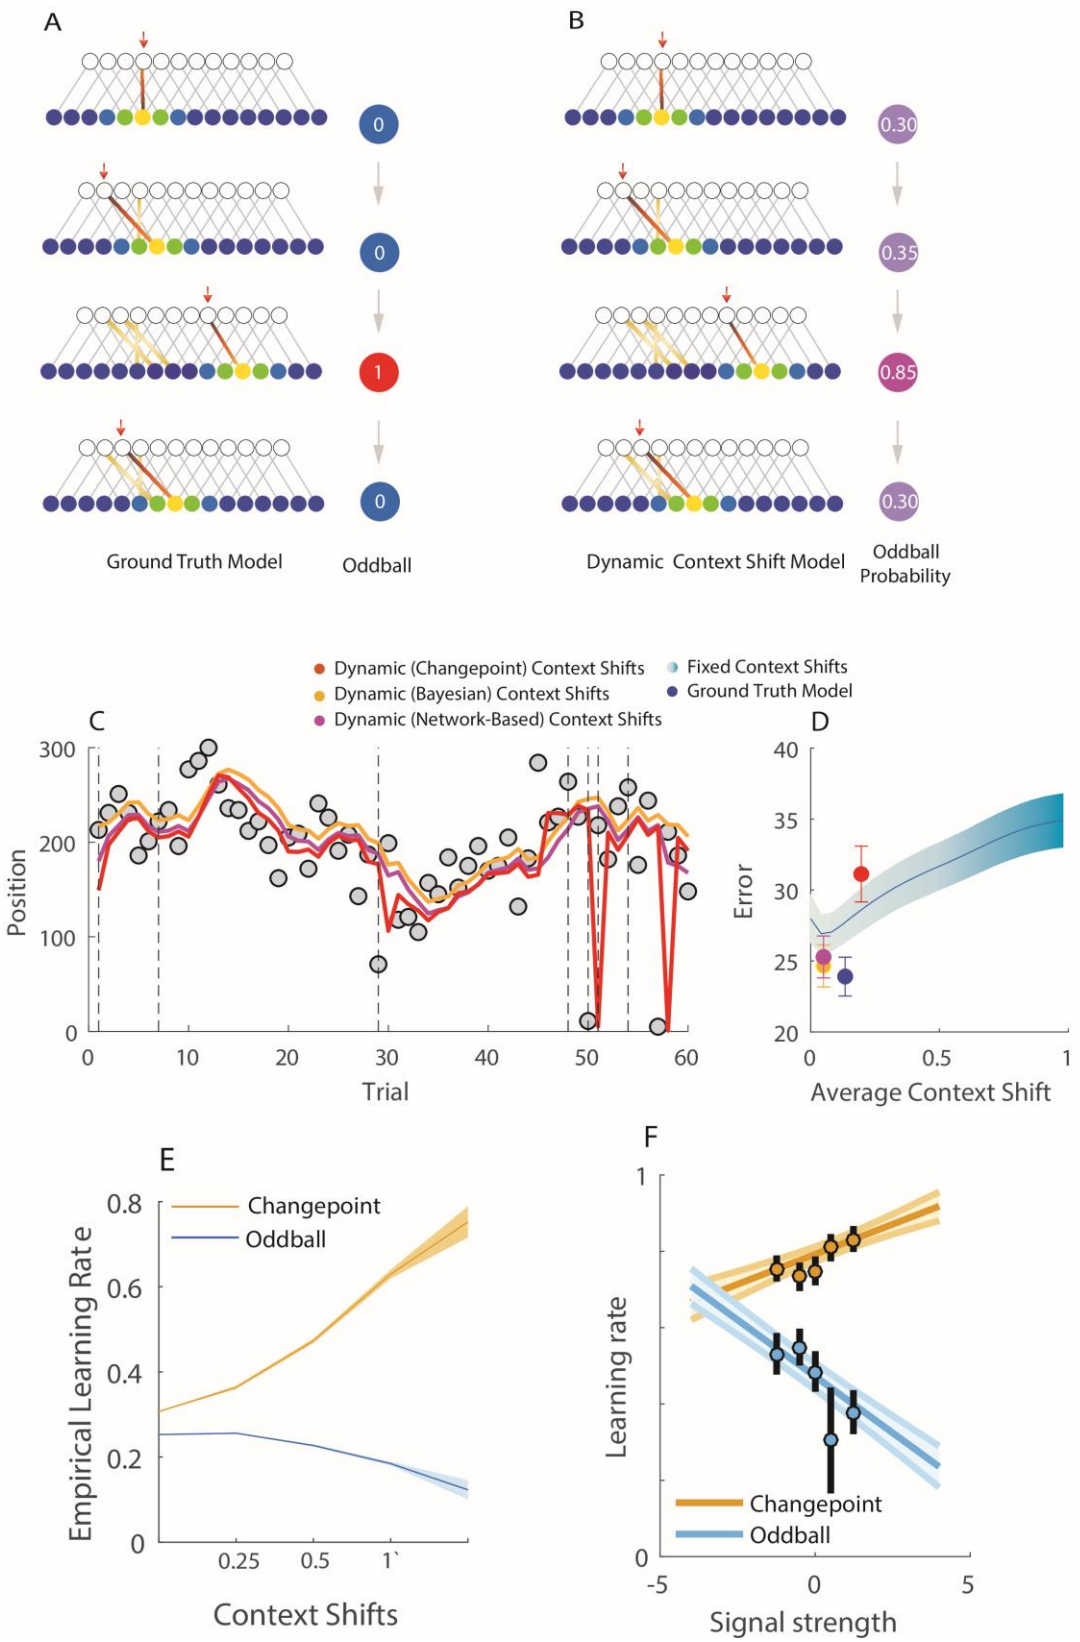

Figure 6 – Neural network structure is able to perform the oddball condition. A) Schematic representation of the ground truth model, which has a constant context shift proportionate to the environmental drift.

On oddball trials (third row) there is a large context shift, but context on next trial returns to its pre-oddball activity pattern. B) Schematic of the dynamic latent state context shift model, which on each trial shifts the context according to α_t term the same as the changepoint condition: OPP but after receiving the supervised learning signal from the outcome, returns back the context to where it was before and shift the context again according to the constant drift rate Thus the net effect would be a constant context shift proportional to the constant drift rate of the helicopter position on every trial. However, a high oddball probability will cause the supervised signal to be stored in a completely separate context. C) Example predictions of the two dynamic models across 60 trials of the oddball condition compared to the changepoint version of dynamic model. Outcomes (i.e. bag locations) are shown in grey circles. Note that because these two models which parametrically changes context according to the probability of oddball events but transitions back to its previous context on the next trial, are less influenced by extreme outcomes than the version of the model that was used in changepoint condition. D) The dynamic context shift models had better performance than the best fixed context shift model ($Bayesian\ Model: t = 9.22.\ df = 31.\ p = 2.13 \times 10^{-10}. Network - Basaed\ Model: t = 7.85.\ df = 31.\ p = 7.2 \times 10^{-9}$) but not as good as the ground truth model, in terms of performance across all simulations. E & F) dynamic context shift model shows the same positive and negative relationship of effective learning rate and measures of context shifts in changepoint and oddball trials respectively as the relationship between learning rate and P300 surprise signal in human subjects during the same task.

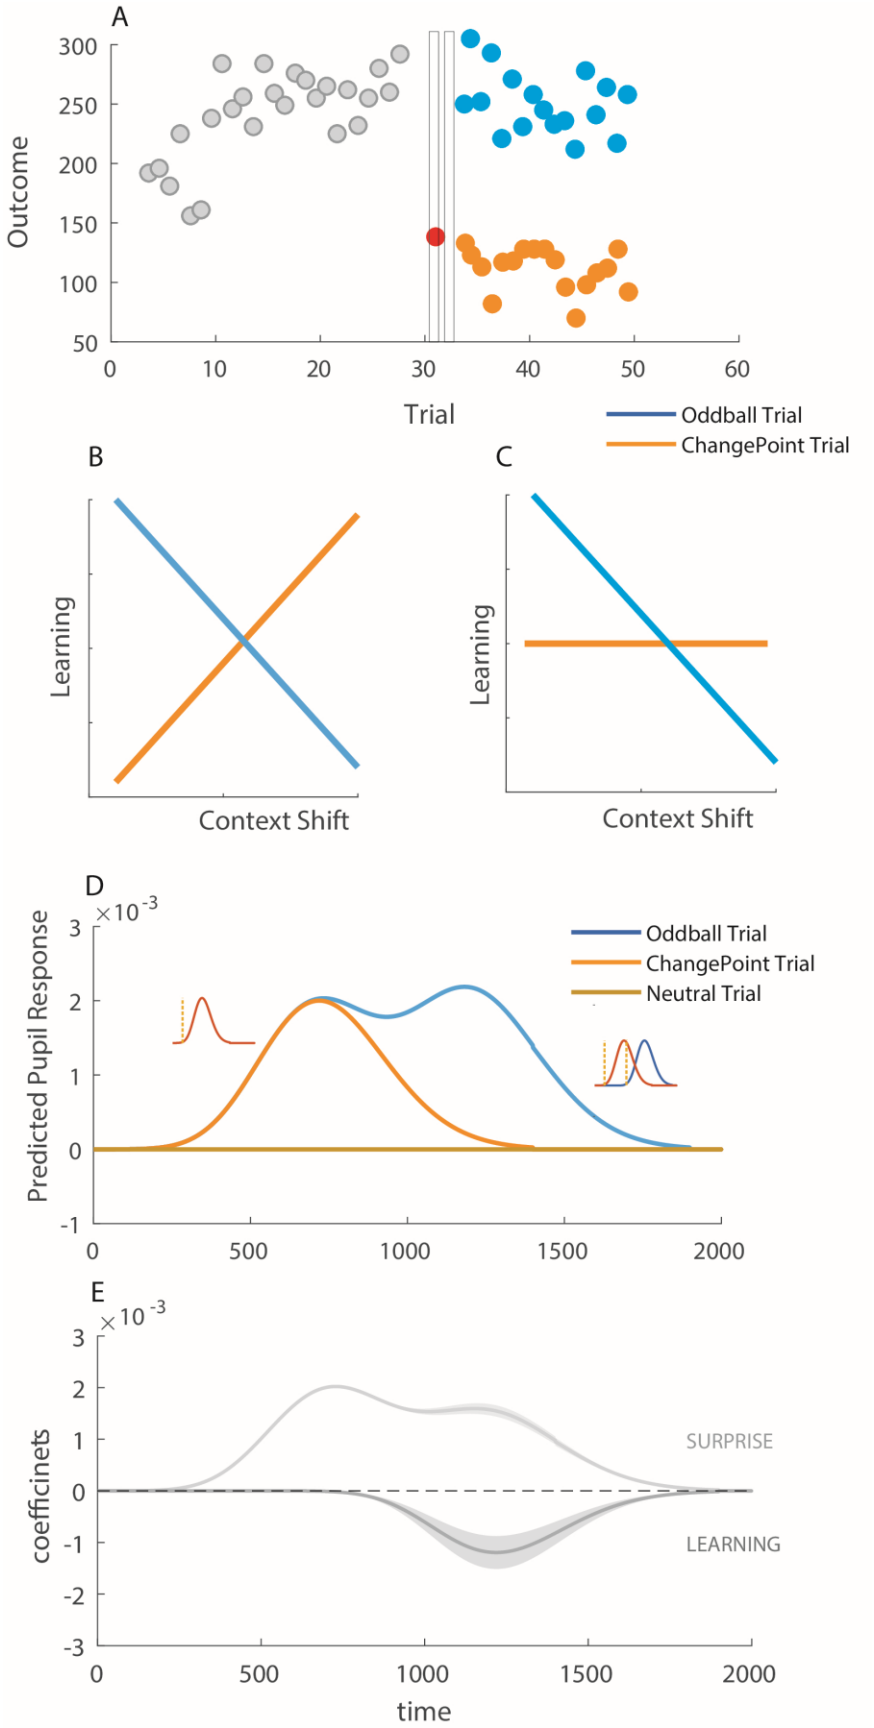
A
Outcome
300
250
200
150
100
50
0
10
20
30
40
50
60
Trial
Oddball Trial
ChangePoint Trial
B
Learning
Context Shift
C
Learning
Context Shift
D
Predicted Pupil Response
×10^-3
3
2
1
0
-1
0
500
1000
1500
2000
Oddball Trial
ChangePoint Trial
Neutral Trial
E
coefficinets
×10^-3
3
2
1
0
-1
-2
-3
0
500
1000
1500
2000
time
SURPRISE
LEARNING

Figure 7 Simulation results predicting the effects of learning and surprise on pupil dilation as a function of time. A) During the first context shift of changepoint/oddball trials we expect positive/negative relationships between learning and changepoint/oddball probability respectively. However, during the second context shift, which is only predicted to happen on oddball trials, producing the same negative effect, no effect would be observed for changepoint trials. C) The result of an example simulated pupil response for the three trial types. A gamma probability density function was convolved with a feedback signal at two discrete time points: The first signal happens at 400s when the subjects are about to see the outcome of the current trial; the second signal happens at 900s around the time before subjects make a prediction for the next trial outcome. Based on the predictions of the dynamic context shift model (A&B), a context shift happens at the first time point in both changepoint and oddball trials while a second context shift is only to happen at the second time point
C) The pupil responses were then combined in one matrix which was then used in a regressor containing condition type, surprise (Change-point/Odd ball probability) and learning (dynamic trial-by-trial learning rate) and the resulting regression coefficients are plotted for each time step in order to show the effect of each variable.

Discussion

Existing models of adaptive learning have failed to capture the range of behaviors in humans across different statistical environments and their underlying neural correlates. Here we developed a neural network framework and demonstrated that internal context shifts within this framework provide a flexible mechanism through which learning rate can be adjusted. Within this test bed we demonstrate that abrupt transitions in context, triggered by unexpected outcomes, can facilitate improved performance in two different statistical environments that differ in the sort of adaptive learning that they require, and do so in a manner that mimics human behavior. Context representations from this dynamic model provide a mechanistic interpretation of activity patterns previously observed in orbitofrontal cortex that abruptly change during periods of rapid learning. The dynamic context shift signal itself provides a mechanistic interpretation for feedback locked P300 signals that conditionally predict learning, and may also resolve a contradiction in different studies examining the relationship between pupil dilation and learning. Taken together, our results provide a mechanistic explanation for adaptive learning behavior and the signals that give rise to it, and furthermore suggest that apparent adjustments in "how much" to learn may actually reflect the dynamics controlling "where" learning takes place.

The context layer that our model employs for flexible learning builds on the notion of latent states for representation learning. Through this lens, our work can be thought of as an extension to a larger body of research on structure learning, much of which has focused on identifying commonalities across stimulus categories (Collins & Frank, 2013; Gershman & Niv, 2010). In cases where temporal dynamics have been explored, the focus has been on the degree to which latent states allow efficient pooling of information across similar contexts that are separated in time (Collins & Frank, 2013; Gershman, Blei, & Niv, 2010; Wilson et al., 2014). Here we highlight another advantage of using temporal dynamics to control active state representations: efficient partitioning of information in time to prevent interference. In addition to highlighting this advantage, our results highlight a shared anatomical basis for state representations across different types of tasks. Patterns of input layer activity in our model transition like neural representations in the orbitofrontal cortex (Nassar, McGuire, et al., 2019), an area that has been suggested to represent latent states for sharing knowledge across common structures(Schuck et al., 2016; Wilson et al., 2014). The ability of our model to account for changes in the pattern of OFC BOLD responses provides some support for our assumption that associations are controlled through changes in the pattern of active input units over time (e.g. figure 3), rather than alternative accounts in which associations are selectively attributed to only

a subset of active units through synchronization (Verbeke & Verguts, 2019), although these two mechanisms need not be mutually exclusive.

Our model description shares some mechanistic similarities with temporal context models (TCM) of episodic and sequential memory recall. (DuBrow, Rouhani, Niv, & Norman, 2017; Franklin et al., 2020; Howard & Kahana, 2002; Kornysheva et al., 2019; Polyn, Norman, & Kahana, 2009; Shankar, Jagadisan, & Howard, 2009). In temporal context models, there is a gradual change in context activity that occurs through passage of time or a through learned linear mapping of the stimuli to contexts,, however our dynamic model relies on discontinuous changes in context more analogous to the underlying latent state dynamics and provides a normative rationale for such abrupt transitions at surprising events, namely that such transitions promote pooling of relevant information within a context (figure 3d) and partitioning of information across contexts (figure 3c) in order to improve inference in complex and dynamic environments (figure 3e & figure 6d). These modeling assumptions allowed us to capture prediction behavior in changepoint and oddball conditions – but to capture a more general set generative statistics – our model would also need to incorporate the possibility of returning to a previous context, and thus considering a hybrid between the assumptions in our model and those of the temporal context models might be an interesting avenue for future study.

More recently, extensions of the temporal context models have suggested the existence of event boundaries which cause discontinuity in temporal context (Zacks, Speer, Swallow, Braver, & Reynolds, 2007). The emergence of these boundaries has been attributed to errors in predictions which, analogous to detected outliers in our model, cause the subsequent observations to be stored in a different context (DuBrow & Davachi, 2013; Rouhani, Norman, Niv, & Bornstein, 2020). Such segmented events also lead to more dissociable representations in fMRI (Antony et al., 2020; Baldassano et al., 2017; Lositsky et al., 2016). While these interpretations of discontinuity in memory are closely related to our model, we take a step further by assigning a key role to such segmentations. In particular, our model shows that it is useful to segment internal context representation after a surprising event in order to improve predictions.

An important question here is how to quantitatively control the transition to new contexts, particularly when such context transitions are not overtly signaled. In previous computational models of event segmentation, surprise has been suggested as the main factor controlling such transition probabilities (Schapiro, Rogers, Cordova, Turk-Browne, & Botvinick, 2013). Our dynamic context shift model uses surprise, as indexed by the probability of an unexpected event (changepoint/oddball), to control context shifts. Such probabilities can be inferred using a Bayesian learning model calibrated to the environmental structure, however, we show that they could also be estimated from output layer of our network itself. Previous work has suggested that changepoint and oddball probability are reflected by BOLD activations in both cortical and subcortical regions (D'Acremont & Bossaerts, 2016; Kao et al., 2020; McGuire et al., 2014; Meyniel & Dehaene, 2017; Nassar, McGuire, et al., 2019; Nassar et al., 2012; O'Reilly et al., 2013; A. J. Yu & Dayan, 2005). While such signals have previously been interpreted as early stage computations performed in the service of computing a learning rate, our work suggests that they serve another purpose, namely in signaling the need to change the active context representation. This interpretation would be consistent with the observation that in at least one case, BOLD responses to surprising events look quite similar across behavioral contexts in which such events should be either learned from, or ignored (D'Acremont & Bossaerts, 2016).

Our model also requires knowledge of the transition structure between trials, in particular facilitating maintenance of the active state in the changepoint condition and return to the previous internal state representation after an oddball event. In principle, this transition structure could be learned implicitly (L. Yu et al., 2020), but in practice we model tasks where the transition structure is provided to participants directly through explicit instruction. We speculate that this transition structure might be provided by a

separate set of neural systems that includes the medial temporal lobe (MTL). This speculation is based on 1) the observation that our context representations mirror the dynamics of representations in orbitofrontal cortex (figure 5), 2) that OFC receives strong inputs from the medial temporal lobe (MTL) (Wikenheiser & Schoenbaum, 2016), and 3) the important role played by the MTL in model based learning and planning (Mattar & Daw, 2018; Schuck & Niv, 2019; Vikbladh et al., 2019). However, future work examining adaptive learning behavior in the face of ambiguous transition structures may help to tease apart the functional roles of different brain signals that occur at surprising task events (Bakst & McGuire, 2020).

Of particular interest in this regard is the feedback-locked P300 signal, an EEG-based correlate of surprise in humans (Kolossa, 2016; Kopp et al., 2016; Mars et al., 2008). A recent study showed that this signal positively related to learning in a changing environment and negatively related to learning in one containing oddballs(Nassar, Bruckner, et al., 2019). Here we show that the context shift variable in our dynamic model has the exact same bidirectional relationship to learning. In our model this reflects a causal relationship, whereby context transitions that persist in the changepoint condition lead new observations to have greater behavioral impact (i.e. more learning; figure 6e), and transient context transitions in the oddball condition limit the behavioral impact of oddball events by associating them with a different context from the one in which predictions are generated (i.e. less learning; figure 6e). We note that this distinction relies in part on our definition of learning. In reality, our model learns makes the same sorts of weight adjustments for both situations, yet the situations differ in the degree to which those weight adjustments impact future predictions.

This bidirectional adjustment of learning rate is a key prediction of our model. We also predict that other physiological measures of surprise that have previously been related to learning, such as pupil diameter, should also provide similar results in environments with different sources of surprising outcomes. However, a key difference of pupil dilation predictions is that given the slow time course of the pupil signal, we predict that it will aggregate multiple state transitions that can occur on an oddball trial (i.e. the transition away from the original state to a new one, and the transition back to the original state). This aspect of the signaling predicts heightened pupil dilations on oddball relative to changepoint trials, which agrees qualitatively with previous observations (O'Reilly et al., 2013), and may help to resolve confusion in the existing literature regarding the relationship between pupil dilations and behavioral adjustment (Nassar et al., 2012; O'Reilly et al., 2013). Our model predicts that such a signal should also drive changes in state representations in OFC. This prediction, at least in part, is consistent with another recent experiment on neuromodulatory control of uncertainty (Muller et al., 2019), in which the strength of pupil dilation predicts the level of uncertainty regarding the current state of the environment, represented in medial orbitofrontal cortex. Our model predicts that these relationships should also depend on the task structure, with state transitions driving OFC representations toward an alternative state in reversal tasks (Muller et al., 2019) toward a completely new persisting state in changepoint tasks (Nassar, McGuire, et al., 2019) and toward a transient state after oddball events. These relationships between state transition signals and neural representations have yet to be measured across the range of contexts that would be necessary to fully test our models predictions, and thus is an interesting avenue for future empirical work.

A major implication of our findings is that behavioral markers of learning rate adjustment may be produced by a network that relies on a fixed learning rate (i.e. the rate of synaptic weight changes), so long as that network adjusts its own internal representations according to the structure of the environment. This is also what distinguishes our model from other accounts of behavior (Nassar et al., 2012, 2010) that adjust learning rate directly, or from computational models that have used surprise detection signals to control learning rate at the synaptic level (Iigaya, 2016). By introducing context shifts in our model we were able to build a mechanistic role for surprise in a learning algorithm that can explain the conditional nature of

heretofore identified learning rate signals: they are actually signaling state transitions, rather than learning per se.

Our model opens the door for a number of future investigations. We catered our analysis to the behavioral experiments of Nassar et al 2019 (Nassar, Bruckner, et al., 2019; Nassar, McGuire, et al., 2019), and therefore only considered changepoint and oddball conditions but did not study the case where context could either shift to a new context or return to a previous context it has learned before. Recognizing that a new observation actually comes from a previously learned context would involve additional pattern recognition and memory retrieval mechanisms (Redish, Jensen, Johnson, & Kurth-nelson, 2007), which might be thought of as part of a more general model-based inference framework as described above (Franklin, Norman, Ranganath, Zacks, & Gershman, 2020; Whittington et al., 2019). That is to say, in order to solve all types of real-world problems, our model would be required to know not only that an observation is different from the recent past, but also which previously encountered state would provide the best generalization to this new situation. Doing so effectively would require organization of states based on similarity, such that similar states shared learning to some degree, in the same way that states which occur nearby in time pool learning in our current model.

In summary, we suggest that flexible learning emerges from dynamic internal context representations that are updated in response to surprising observations in accordance with knowledge of task structure. Our model requires representations consistent with those that have previously been observed in orbitofrontal cortex as well as state transition signals necessary to update them. We suggest that biological signals previously thought to reflect "dynamic learning rates" actually signal the need for internal state transitions, and provide the first mechanistic explanation for the context-dependence with which these signals relate to learning. Taken together, our results support the notion that adaptive learning behaviors may arise through dynamic control of representations of task structure

Methods

Experimental task:

We examine human and model behavior in a predictive inference task that has been described previously (McGuire et al., 2014; Nassar & Troiani, 2020). The predictive inference task is a computerized task in which an animated helicopter drops bags in an open field. In the pre-training session, human subjects learned to move a bucket with a joystick beneath the helicopter to catch bags that could contain valuable contents. During the main phase of the experiment, the helicopter was occluded by clouds and the participants were forced to infer its position on based on the locations of bags it had previously dropped.

Our initial simulations focus on dynamic environments in which surprising events often signal a change in the underlying generative structure (changepoint condition; figures 1-5). In the chanagepoint condition, bag locations were drawn from a distribution centered on the helicopter with a fixed standard deviation of 25 (unless otherwise specified in the analysis). The helicopter remained stationary on most trials, but occasionally and abruptly changed its position to a random uniform horizontal screen location. The probability of moving to a new location on a given trial is controlled by the hazard rate ($H = 0.1$). Unless otherwise noted, our modeling results are presented with 32 simulated subjects, to correspond to the sample size in (McGuire et al., 2014).

We also considered a complementary generative environment in which surprising events were unrelated to the underlying generative structure (oddball condition; figure 6)(Nassar & Troiani, 2020). In the oddball

condition, the helicopter would gradually move in the sky according to a Gaussian random walk (drift rate $(DR) = 10$). In the oddball condition bags were typically drawn from a normal distribution centered on the helicopter as described above, but on occasion a bag would be dropped in random location unrelated to the position of the helicopter. The location of an oddball bag was sampled from a uniform distribution that spanned the entire screen. The probability of an oddball event was controlled by a hazard rate ($H = 0.1$).

Normative learning model:

The simplest delta rule model in the predictive inference task states that:

$$B_{t+1} = B_t + \alpha\delta \quad (1)$$

$$\delta = Bag\ Position(t) - B_t(t) \quad (2)$$

here B is belief regarding helicopter position on each trial , δ is the prediction error and α is the learning rate. With a constant α, the model assigns the same weight to all predictions and outcomes. Previous work has shown that Bayesian optimal inference can be reduced to a delta rule learning under certain approximations, leading to normative prescriptions for learning rate that are adjusted dynamically (Nassar, Bruckner, et al., 2019; Nassar et al., 2010). The resulting normative learning model takes information which human subjects would normally obtain during the pre-training sessions including Hazard rate and standard deviation, but also computes two latent variables, by using the trial-by-trial prediction error: 1) changepoint probability which is computed after an outcome is observed and indicates the probability that the observed outcome has reflects a change in the helicopter location, and 2) relative uncertainty which is computed before making the next prediction and indicates the models uncertainty about the location of the helicopter. Detailed information regarding how CPP and RU are calculated can be found in appendix-1.

In the changepoint condition the normative learning rate α_t is defined by:

$$\alpha_t = CPP + RU - CPP \times RU \quad (3)$$

Where CPP is changepoint probability and RU is relative uncertainty. Using these two latent variables, which both track the prediction error, but with different temporal dynamics (McGuire et al., 2014), the model computes a dynamic learning rate that increases after a changepoint and gradually decreases in the following stable period after a changepoint.

The same approximation to Bayesian inference can be applied in the oddball condition to produce a normative learning model that relies on oddball probability and relative uncertainty to guide learning. While the latent variables and form of the model mimic that in the changepoint condition, the learning rate differs in that it is reduced, rather than enhanced, in response to outcomes that are inconsistent with prior expectations:

$$\alpha_t = RU - OBP \times RU \quad (4)$$

Where OBP is the models posterior probability estimate that an outcome was an oddball event and RU reflects the model's uncertainty about the current helicopter location. Thus, normative inference in the oddball condition requires decreasing learning according to the probability of an extreme event (oddball), whereas normative inference in the changepoint condition required increasing it.

Neural network models:

In order to better understand how normative learning principles might be applied in a neural network we created a series of neural network models that use supervised learning rules to generate predictions in the predictive inference task. Specifically, we created a two-layer feed forward neural network that can perform the predictive inference task.

Network architecture includes two layers:

The input layer is composed of N neurons with responses characterized by a von Mises (circular) distribution with mean m and fixed concentration equal to 16 We implemented several versions of this model depending on how the mean m changes on a trial-by-trial basis.

The output layer contains neurons corresponding to spatial location of the bucket on the screen. The response of output layer neurons was computed by the weighted sum of input layer:

$$r_j = \sum_{i=1}^{N_{in}} x_i w_{ij} \quad (5)$$

Where x_i is the activation of neuron i in the input layer, r_j is the response of neuron j in the output layer and w_{ij} is the connection weight between neuron i and neuron j. The bucket position chosen by the model on each trial was computed as a linear readout of the output layer:

$$estimate = \sum_{j=1}^{N_{out}} L_j r_j \quad (6)$$

Where L_j is the location encoded by each corresponding unit r_j in the output layer. Weight matrix is randomly initialized with a uniform distribution of mean zero and SD equal to 5×10^{-4}. The network is then trained on each trial by modifying the weight matrix according to:

$$w_{ij} = (1 - \eta) w_{ij} + \eta y_j x_i \ (7)$$

Where y_j is the probability on a normal distribution centered on the observed outcome evaluated at L_j, and η is a constant synaptic learning rate controlling the weight changes of the neural network and was set to 0.1 for all models simulations.

The fixed context shifts models:

In the first models we consider, *fixed context shift* models, The mean m is computed on each trial as follows:

$$m_{(t+1)} = m_{(t)} + \Delta m_f \ (8)$$

Here, Δm_f takes a fixed value for all trials throughout the simulation (figure 2b&c). We considered 50 different Δm_f values ranging from 0 to 2 in order to study the effect of context shifts on model performance. The word "context" refers to the subpopulation of input layer neurons that are firing above the threshold on each trial. By incrementally increasing the mean of response distribution of the input layer, we can think of this context being changed on each trial. The architecture of the input layer is arranged in a circle so that hypothetically the context would be able to shift clockwise indefinitely. In order to minimize interference from previous visits to a set of context neurons we implemented weight decay (WD) on each time step according to the following rule:

$$W_{t+1}(x_t < threshold) = W_t(x_t < threshold) \times WD \quad (9)$$

$$WD = 0.1$$

Note that this weight decay is not intended as a biological assumption, but rather a convenient simplification to allow the model to represent a large number of contexts with a small pool of neurons.

Therefore, on each trial, first the model would make a prediction based on weighted sum of the active input, observe an outcome, shift the context by the assigned context shift and store the supervised signal in the new context. This new context is in turn used at the beginning of the next trial to produce a response.

Table 1- Summary of the parameters used for simulation of the probabilistic inference task and neural network training.

Neural Network Parameter:	***Value***	***Description***
Number of neurons in the input layer (N_{in})	63	Equally-spaced points between $[-\pi. +\pi]$ incrementing by 0.1
Concentration (κ)	32	Concentration of the von Mises pdf used in the input layer
Number of neurons in the output layer (N_{out})	41	Equally-spaced points between -50 and 350 , incrementing by 10
Synaptic learning Rate (η)	0.1	
Weight Decay Threshold	0.01	
Weight Decay Rate (*WD*)	0.1	
Model Hazard Rate	0.7	The model uses a higher value compared to the actual hazard rate for optimal performance
Task Parameter:		
Hazard Rate (*H*)	0.1	Probability of a changepoint/oddball trial
Noise (σ_N)	25	Standard Deviation of random process generating outcomes
Standard Deviation of Drift Rate (σ_{drift})	10	Standard Deviation of the random process generating drift rate in oddball condition

Ground Truth context shifts neural network model:

To leverage the benefits of different context shifts which we observed in the fixed context shifts models we designed a model that uses optimal context shift for each trial. The ground truth context shifts model has the same design of a fixed context shift model except instead of the constant term Δm, the model uses a function for computing Δm:

$$\Delta m = \begin{cases} max(\Delta m_f). & if \;\; t \; is \; changepoint \\ 0. & otherwise \end{cases} \qquad (10)$$

Dynamic context shifts neural network model:

In order to make the ground truth context shifts model more biologically plausible, we estimated the probability of changepoints subjectively, and used this estimate of changepoint probability (CPP) to control context shifts:

$$\Delta m = f(CPP) \qquad (11)$$

The function, f, provides a fixed level of context shift according to the estimated changepoint probability by inverting the relationship between context shift and effective learning rate observed in the fixed context shift models and plotted in figure 2d. Thus, on each trial, the model will choose a context shift belonging to a fixed context shift model that has the closest effective learning rate to CPP. This way, more surprising outcomes will consistently result in larger context shifts with CPPs of zero and one resulting in 0 or maximal context shifting, as in the ground truth model. CPP was computed either using the Bayesian model described above "Bayesian context shift" or from an approximation derived from the neural network itself. In the network-based version, the probability of a state transition is subjectively computed by the following equation:

$$\frac{H/41}{H/41+ r_{X_t}(1-H)} \quad (12)$$

which can be interpreted as a network-based approximation to Bayesian CPP estimation (See Appendix-1), or in terms of a non-linear activation over prediction errors such as has been proposed in various conflict models (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Cockburn & Frank, 2013). H can be thought of in Bayesian terms as a hazard rate, or in neural network terms as controlling the threshold of the activation function, and r_{X_t} is the firing rate of the output unit corresponding to the location X_t, which can be thought of as providing a readout of the outcome probability based on a Bayesian population code. The 41 reflects the total number of output units in our population, and since outcomes could occur that were in between the tuning of these units, in practice we used linear interpolation to estimate r_{X_t}based the two output units closest to the actual outcome location. The hazard rate H was set to 0.7 for the changepoint condition in order to achieve optimal performance (see figure 3-1). Note that this fixed hazard rate, which maximized model performance, is considerably higher than the true rate of changepoints in the task (0.1).

Extension of network models to the oddball condition:

To test our proposed models in a variation of the task where prediction errors are not indicative of a change in context i.e. oddball condition we use the same design of neural network but with a simple modification in temporal dynamics of context shifts.

The task involved the same paradigm described above, but with outcomes (i.e. bag locations) determined by a different generative structure. In particular, the helicopter location gradually changed its position in the sky with a constant drift rate, and bags were occasionally sampled from a uniform distribution spanning the range of possible outcomes, rather than being "dropped" from the helicopter itself(Nassar & Troiani, 2020; Nassar et al., 2021).

The ground truth neural network model was modified to incorporate the alternate generative structure of the oddball condition. In particular, on each trial, input activity mean m was changed by 1) maximally context shifting in response to oddballs at the time of feedback, 2) "returning" from the oddball induced context shift at the end of the feedback period, prior to the subsequent trial, and 3) adding a constant value (0.05) proportional to the fixed drift rate of the random walk process prior to making the prediction. Thus after a prediction is made on trial context mean changes according to:

$$\Delta m_1 = \begin{cases} max(\Delta m_f), & if\ \ t\ is\ oddball \\ 0 & otherwise \end{cases} \quad (13)$$

$$m_{t+1} = m_t + \Delta m_1 \quad (14)$$

But, after the model receives the supervised signal and stores it the new context, context transition back to:

$$m_{t+1} = m_t - \Delta m_1 + \Delta m_2 \quad (15)$$

Where Δm_2 is a constant (here 0.05) is proportional to the drift rate of the random walk process. This leads the information from oddball trial to be stored in a different context that won't influence the upcoming prediction of the model.

The dynamic context shift models were constructed to follow the same logic, but using subjective measures of oddball probability rather than perfect knowledge about whether a trial is an oddball. Specifically, we updated context upon observing feedback according to the probability that the feedback reflects an oddball (OP):

$$\Delta m_1 = f(OP) \quad (16)$$

$$m_{t+1} = m_t + \Delta m_1 \quad (17)$$

And prior to making a prediction for the subsequent trial returned to the previous context except with a slight shift modeling to account for the drift in the helicopter position due to the random walk:

$$m_{t+1} = m_t - \Delta m_1 + \Delta m_2 \quad (18)$$

This model captures the intuition that if an outcome is known to be an outlier, it should be partitioned from knowledge that pertains to the helicopter location, rather than combined with it. To accomplish this, the model changes the context first according to the oddball probability or Δm_1 in above equation, after storing the supervised learning signal in the new context, the model transition back to its previous context by subtracting the first context shift term $\Delta m_1$and move the context according to a constant shift proportional to the drift rate.Δm_2. The Δm_1 term causes significant shifts on oddball trials, but after that the model transition back to previous context and shifts according to the Δm_2 which would not be influenced by oddball trials. Similar to the changepoint condition, here, we also made a version of the dynamic Bayesian context shift model, which used network output layer activity to compute subjective measures of oddball probability.

Representational similarity analysis:

We computed a trial-by-trial dissimilarity matrix where each cell in the matrix represent the number corresponding to the dissimilarity between two trials. The dissimilarity matrix (D) of the dynamic context shifts model uses Euclidean distance and is computed by:

$$D_{ij} = \sqrt{\sum_{q=1}^{N_{in}} (Act_{(i.q)} - Act_{(j.q)})^2} \quad (19)$$

Where $Act_{(i,q)}$ and $Act_{(j,q)}$ are the activities of neuron q of the input layer on trials i and j respectively. The dissimilarity matrix of normative learning model is constructed by using normative learning rates on each trial:

$$M_{ij} = 1 - \prod_{i=t}^{j-1} 1 - \alpha_t \quad (20)$$

Where the similarity between two trials i and j is inversely correlated with the product of the learning rates α_t of all trials between them. Subtracting similarity from 1 gives us the trial-by-trial dissimilarity matrix (M). Using the RSA approach we were able to directly compare the results of our neural network input layer activity to the dynamic learning rate computed from CPP and RU in a normative model.

Behavioral analysis:

Behavioral analyses are aimed at understanding the degree to which we revise our behavior in response to new observations. In order to quantify this, we define an "effective learning rate" as the slope of the relationship between trial-to-trial predictions errors and trial-to –trial updates (i.e. the change in bucket position from one trial to the next). The adjective "effective" is chosen here so that this learning rate won't be mistaken by the reader with two other learning rates used in this paper: 1) the fixed learning rate of the neural network 2) the normative learning rate from the reduced Bayesian model. To measure effective learning rate, we regressed updates (UP) onto the prediction errors (PE) that preceded them:

$$UP = \beta_0 + \beta_1 \times PE \quad (21)$$

The resulting slope term, β_1 captures the effective learning rate, or the amount of update expected for a given prediction error. We also performed a more extensive regression analysis that included terms for 1) prediction error 2) prediction error times changepoint probability 4) prediction error times relative uncertainty (figure 4d).

Comparison to P300 analysis:

For analyzing the effect of trial-to-trial variability in context shifts from the dynamic context shift model on effective learning rate produced by that model, we fit the regression model above to simulated predictions for the dynamic context shift models, but did so while splitting data into quartiles according to the size of the context shift size that the model underwent on a given trial. The corresponding figure (figure 6e) of P300 signal and learning rate are from ref (Nassar, Bruckner, et al., 2019).

Pupil Response Simulation:

We modeled 480 trials of a predictive inference task for each of the two conditions (oddball, changepoint). We created synthetic pupil traces by defining time points for feedback-locked context shifts, which occurred 400ms after oddball or changepoint, and pre-prediction context shifts at 900ms after oddball events (see eq. 10 & 13). We used measurements of context shift for the respective changepoint and oddball trials (see eq. 11 & 16) at these time points and convolved these measurements with a gamma distribution to create simulated time courses of a pupil response under the assumption that the pupil signal reflects the need for a context shift. We analyzed this signal with a regression model that was applied to all synthetic data in sliding windows of time. Explanatory variables in our model included surprise (changepoint/oddball probability computed from normative model) and learning (trial-by-trial learning rate computed from the normative model).

Data Availability

All analysis and modeling code (including code for generating the figures) has been made available on GitHub: https://github.com/NiloufarRazmi/dynamicContextShift .

References:

Adams, R. P., & MacKay, D. J. C. (2007). *Bayesian Online Changepoint Detection*. Retrieved from http://arxiv.org/abs/0710.3742

Antony, J. W., Hartshorne, T. H., Pomeroy, K., Gureckis, T. M., Hasson, U., McDougle, S. D., & Norman, K. A. (2020). Behavioral, physiological, and neural signatures of surprise during naturalistic sports viewing. *BioRxiv*, 2020.03.26.008714. https://doi.org/10.1101/2020.03.26.008714

Bakst, L., & McGuire, J. (2020). Eye movements reflect adaptive predictions and predictive precision. *Journal of Experimental Psychology: General*. https://doi.org/10.1037/xge0000977

Baldassano, C., Chen, J., Zadbood, A., Pillow, J. W., Hasson, U., & Norman, K. A. (2017). Discovering Event Structure in Continuous Narrative Perception and Memory. *Neuron*, *95*(3), 709-721.e5. https://doi.org/10.1016/j.neuron.2017.06.041

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). *Learning the value of information in an uncertain world*. *10*(9), 1214–1221. https://doi.org/10.1038/nn1954

Bernacchia, A., Seo, H., Lee, D., & Wang, X.-J. (2011). A reservoir of time constants for memory traces in cortical neurons. *Nature Neuroscience*, *14*(3), 366–372. https://doi.org/10.1038/nn.2752

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, Vol. 108, pp. 624–652. https://doi.org/10.1037/0033-295X.108.3.624

Bouret, S., & Sara, S. J. (2005). Network reset: a simplified overarching theory of locus coeruleus noradrenaline function. *Trends in Neurosciences*, *28*(11), 574–582. https://doi.org/10.1016/j.tins.2005.09.002

Browning, M., Behrens, T. E., Jocham, G., O'Reilly, J. X., & Bishop, S. J. (2015). Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nature Neuroscience*, *18*(4), 590–596. https://doi.org/10.1038/nn.3961

Cockburn, J., & Frank, M. (2013). Reinforcement Learning, Conflict Monitoring, and Cognitive Control: An Integrative Model of Cingulate-Striatal Interactions and the ERN. *Neural Basis of Motivational and Cognitive Control*, 310–331. https://doi.org/10.7551/mitpress/9780262016438.003.0017

Collins, A. G. E., & Frank, M. J. (2013). Cognitive control over learning: creating, clustering, and generalizing task-set structure. *Psychological Review*, *120*(1), 190–229. https://doi.org/10.1037/a0030852

D'Acremont, M., & Bossaerts, P. (2016). Neural Mechanisms behind Identification of Leptokurtic Noise and Adaptive Behavioral Response. *Cerebral Cortex*, *26*(4), 1818–1830. https://doi.org/10.1093/cercor/bhw013

Donahue, C. H., & Lee, D. (2015). Dynamic routing of task-relevant signals for decision making in

dorsolateral prefrontal cortex. *Nature Neuroscience*, *18*(2), 295–301. https://doi.org/10.1038/nn.3918

DuBrow, S., & Davachi, L. (2013). The influence of context boundaries on memory for the sequential order of events. *Journal of Experimental Psychology: General*, *142*(4), 1277–1286. https://doi.org/10.1037/a0034024

Farashahi, S., Donahue, C. H., Hayden, B. Y., Lee, D., & Soltani, A. (2019). Flexible combination of reward information across primates. *Nature Human Behaviour*, *3*(11), 1215–1224. https://doi.org/10.1038/s41562-019-0714-3

Farashahi, S., Donahue, C. H., Khorsand, P., Seo, H., Lee, D., & Soltani, A. (2017). Metaplasticity as a Neural Substrate for Adaptive Learning and Choice under Uncertainty. *Neuron*, *94*(2), 401-414.e6. https://doi.org/10.1016/j.neuron.2017.03.044

Fischer, A. G., & Ullsperger, M. (2013). Real and Fictive Outcomes Are Processed Differently but Converge on a Common Adaptive Mechanism. *Neuron*, *79*(6), 1243–1255. https://doi.org/10.1016/j.neuron.2013.07.006

Franklin, N. T., Norman, K. A., Ranganath, C., Zacks, J. M., & Gershman, S. J. (2020). Structured Event Memory: A neuro-symbolic model of event cognition. *Psychological Review*, *127*(3), 327–361. https://doi.org/10.1037/rev0000177

Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, Learning, and Extinction. *Psychological Review*, Vol. 117, pp. 197–209. Retrieved from https://nivlab.princeton.edu/sites/default/files/nivlab/files/gershmanetal2009.pdf

Gershman, S. J., & Niv, Y. (2010). Learning latent structure: Carving nature at its joints. *Current Opinion in Neurobiology*, *20*(2), 251–256. https://doi.org/10.1016/j.conb.2010.02.008

Iigaya, K. (2016). Adaptive learning and decision-making under uncertainty by metaplastic synapses guided by a surprise detection system. *ELife*, *5*, e18073. https://doi.org/10.7554/eLife.18073

Jepma, M., Brown, S. B. R. E., Murphy, P. R., Koelewijn, S. C., de Vries, B., van den Maagdenberg, A. M., & Nieuwenhuis, S. (2018). Noradrenergic and Cholinergic Modulation of Belief Updating. *Journal of Cognitive Neuroscience*, *30*(12), 1803–1820. https://doi.org/10.1162/jocn_a_01317

Jepma, M., Murphy, P. R., Nassar, M. R., Rangel-Gomez, M., Meeter, M., & Nieuwenhuis, S. (2016). Catecholaminergic Regulation of Learning Rate in a Dynamic Environment. *PLOS Computational Biology*, *12*(10), e1005171. Retrieved from https://doi.org/10.1371/journal.pcbi.1005171

Joshi, S., & Gold, J. I. (2020). Pupil Size as a Window on Neural Substrates of Cognition. *Trends in Cognitive Sciences*, *24*(6), 466–480. https://doi.org/10.1016/j.tics.2020.03.005

Kao, C.-H., Khambhati, A. N., Bassett, D. S., Nassar, M. R., McGuire, J. T., Gold, J. I., & Kable, J. W. (2020). Functional brain network reconfiguration during learning in a dynamic environment. *Nature Communications*, *11*(1), 1682. https://doi.org/10.1038/s41467-020-15442-2

Kolossa, A. (2016). *A New Theory of Trial-by-Trial P300 Amplitude Fluctuations*. https://doi.org/10.1007/978-3-319-32285-8_3

Kopp, B., Seer, C., Lange, F., Kluytmans, A., Kolossa, A., Fingscheidt, T., & Hoijtink, H. (2016). P300 amplitude variations, prior probabilities, and likelihoods: A Bayesian ERP study. *Cognitive, Affective, & Behavioral Neuroscience*, *16*(5), 911–928. https://doi.org/10.3758/s13415-016-0442-3

Li, Y. S., Nassar, M. R., Kable, J. W., & Gold, J. I. (2019). Individual Neurons in the Cingulate Cortex

Encode Action Monitoring, Not Selection, during Adaptive Decision-Making. *The Journal of Neuroscience*, *39*(34), 6668 LP – 6683. https://doi.org/10.1523/JNEUROSCI.0159-19.2019

Lositsky, O., Chen, J., Toker, D., Honey, C. J., Shvartsman, M., Poppenk, J. L., … Norman, K. A. (2016). Neural pattern change during encoding of a narrative predicts retrospective duration estimates. *ELife*, *5*, e16070. https://doi.org/10.7554/eLife.16070

Mars, R. B., Debener, S., Gladwin, T. E., Harrison, L. M., Haggard, P., Rothwell, J. C., & Bestmann, S. (2008). Trial-by-Trial Fluctuations in the Event-Related Electroencephalogram Reflect Dynamic Changes in the Degree of Surprise. *The Journal of Neuroscience*, *28*(47), 12539 LP – 12545. https://doi.org/10.1523/JNEUROSCI.2925-08.2008

Massi, B., Donahue, C. H., & Lee, D. (2018). Volatility Facilitates Value Updating in the Prefrontal Cortex. *Neuron*, *99*(3), 598-608.e4. https://doi.org/https://doi.org/10.1016/j.neuron.2018.06.033

Mathys, C., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A bayesian foundation for individual learning under uncertainty. *Frontiers in Human Neuroscience*, *5*, 39. https://doi.org/10.3389/fnhum.2011.00039

Mattar, M. G., & Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neuroscience*, *21*(11), 1609–1617. https://doi.org/10.1038/s41593-018-0232-z

McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally dissociable influences on learning rate in a dynamic environment. *Neuron*, *84*(4), 870–881. https://doi.org/10.1016/j.neuron.2014.10.013

Meyniel, F., & Dehaene, S. (2017). Brain networks for confidence weighting and hierarchical inference during probabilistic learning. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(19), E3859–E3868. https://doi.org/10.1073/pnas.1615773114

Muller, T. H., Mars, R. B., Behrens, T. E., & O'Reilly, J. X. (2019). Control of entropy in neural models of environmental state. *ELife*, *8*, 1–30. https://doi.org/10.7554/eLife.39404

Nassar, M. R., Bruckner, R., & Frank, M. J. (2019). Statistical context dictates the relationship between feedback-related EEG signals and learning. *ELife*, *8*, 1–26. https://doi.org/10.7554/eLife.46975

Nassar, M. R., & Gold, J. I. (2010). *Supplementary Material for : Bayesian On-line Learning of the Hazard Rate in Change-Point Problems*. *22*(9), 2452–2476.

Nassar, M. R., McGuire, J. T., Ritz, H., & Kable, J. W. (2019). Dissociable forms of uncertainty-driven representational change across the human brain. *Journal of Neuroscience*, *39*(9), 1688–1698. https://doi.org/10.1523/JNEUROSCI.1713-18.2018

Nassar, M. R., Rumsey, K. M., Wilson, R. C., Parikh, K., Heasly, B., & Gold, J. I. (2012). Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*, *15*(7), 1040–1046. https://doi.org/10.1038/nn.3130

Nassar, M. R., & Troiani, V. (2020). The stability flexibility tradeoff and the dark side of detail. *Cognitive, Affective, & Behavioral Neuroscience*. https://doi.org/10.3758/s13415-020-00848-8

Nassar, M. R., Waltz, J. A., Albrecht, M. A., Gold, J. M., & Frank, M. J. (2021). All or nothing belief updating in patients with schizophrenia reduces precision and flexibility of beliefs. *Brain*. https://doi.org/10.1093/brain/awaa453

Nassar, M. R., Wilson, R. C., Heasly, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *Journal of*

Neuroscience, *30*(37), 12366–12378. https://doi.org/10.1523/JNEUROSCI.0822-10.2010

O'Reilly, J. X. (2013). Making predictions in a changing world-inference, uncertainty, and learning. *Frontiers in Neuroscience*, *7*(7 JUN), 1–10. https://doi.org/10.3389/fnins.2013.00105

O'Reilly, J. X., Schüffelgen, U., Cuell, S. F., Behrens, T. E. J., Mars, R. B., & Rushworth, M. F. S. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(38). https://doi.org/10.1073/pnas.1305373110

Redish, A. D., Jensen, S., Johnson, A., & Kurth-nelson, Z. (2007). *Reconciling Reinforcement Learning Models With Behavioral Extinction and Renewal : Implications for Addiction , Relapse , and Problem Gambling*. *114*(3), 784–805. https://doi.org/10.1037/0033-295X.114.3.784

Reimer, J., McGinley, M. J., Liu, Y., Rodenkirch, C., Wang, Q., McCormick, D. A., & Tolias, A. S. (2016). Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex. *Nature Communications*, *7*(1), 13289. https://doi.org/10.1038/ncomms13289

Rouhani, N., Norman, K. A., Niv, Y., & Bornstein, A. M. (2020). Reward prediction errors create event boundaries in memory. *Cognition*, *203*, 104269. https://doi.org/https://doi.org/10.1016/j.cognition.2020.104269

Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature Neuroscience*, *16*(4), 486–492. https://doi.org/10.1038/nn.3331

Schuck, N. W., Cai, M. B., Wilson, R. C., & Niv, Y. (2016). Human Orbitofrontal Cortex Represents a Cognitive Map of State Space. *Neuron*, *91*(6), 1402–1412. https://doi.org/10.1016/j.neuron.2016.08.019

Schuck, N. W., & Niv, Y. (2019). Sequential replay of nonspatial task states in the human hippocampus. *Science*, *364*(6447), eaaw5181. https://doi.org/10.1126/science.aaw5181

Soltani, A., & Izquierdo, A. (2019). Adaptive learning under expected and unexpected uncertainty. *Nature Reviews Neuroscience*, *20*(10), 635–644. https://doi.org/10.1038/s41583-019-0180-y

Verbeke, P., & Verguts, T. (2019). Learning to synchronize: How biological agents can couple neural task modules for dealing with the stability-plasticity dilemma. *PLOS Computational Biology*, *15*(8), e1006604. Retrieved from https://doi.org/10.1371/journal.pcbi.1006604

Vikbladh, O. M., Meager, M. R., King, J., Blackmon, K., Devinsky, O., Shohamy, D., … Daw, N. D. (2019). Hippocampal Contributions to Model-Based Planning and Spatial Memory. *Neuron*, *102*(3), 683-693.e4. https://doi.org/10.1016/j.neuron.2019.02.014

Whittington, J. C. R., Muller, T. H., Mark, S., Chen, G., Barry, C., Burgess, N., & Behrens, T. E. J. (2019). The Tolman-Eichenbaum Machine: Unifying space and relational memory through generalisation in the hippocampal formation. *BioRxiv*, 770495. https://doi.org/10.1101/770495

Wikenheiser, A., & Schoenbaum, G. (2016). Over the river, through the woods: cognitive maps in the hippocampus and orbitofrontal cortex. *Nature Reviews Neuroscience*, *17*. https://doi.org/10.1038/nrn.2016.56

Wilson, R. C., Nassar, M. R., & Gold, J. I. (2010). Bayesian online learning of the hazard rate in change-point problems. *Neural Computation*, *22*(9), 2452–2476. https://doi.org/10.1162/NECO_a_00007

Wilson, R. C., Nassar, M. R., & Gold, J. I. (2013). A mixture of delta-rules approximation to bayesian

inference in change-point problems. *PLoS Computational Biology*, *9*(7), e1003150–e1003150. https://doi.org/10.1371/journal.pcbi.1003150

Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal cortex as a cognitive map of task space. *Neuron*, *81*(2), 267–279. https://doi.org/10.1016/j.neuron.2013.11.005

Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, *46*(4), 681–692. https://doi.org/10.1016/j.neuron.2005.04.026

Yu, L., Wilson, R., & Nassar, M. (2020). *Adaptive learning is structure learning in time*. https://doi.org/10.31234/osf.io/r637c

Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., & Reynolds, J. R. (2007). Event perception: A mind-brain perspective. *Psychological Bulletin*, Vol. 133, pp. 273–293. https://doi.org/10.1037/0033-2909.133.2.273

Appendix – 1

Changepoint condition:

The probability that a changepoint has occurred in the current trial is computed as the likelihood of current trial observation coming from a new helicopter position, which is randomly chosen from a uniform distribution between 0 and 300, over the sum of the likelihoods of the current observation coming from a non-changepoint distribution (Gaussian distribution centered around the most recent belief about helicopter positions) and changepoint distribution:

$$\text{CPP} = \Omega_t = \frac{U(X_t|0.300)\, H}{U(X_t|0.300)H + N(\delta_t; 0.\sigma_\mu^2 + \sigma_N^2)(1-H)}$$

Where Ω_t is changepoint probability, H is the hazard rate and is the same hazard rate used in the generative process for simulating outcomes, δ_t is the difference between the predicted and actual outcome, σ_N^2 is the variance of the distribution of bags around the helicopter, and σ_μ^2 is the variance on predictive distribution (estimation uncertainty) and is computed recursively after observing an outcome and using the changepoint probability of the current trial:

$$\sigma_\mu^2 = \Omega_t \sigma_N^2 + (1-\Omega_t)\, \tau_t \sigma_N^2 + \Omega_t (1-\Omega_t)(\tau_t + B_t(1-\tau_t) - X_t)$$

Where B_t is the current belief about helicopter position and X_t is the current trial outcome and τ_t is the relative uncertainty.

Relative uncertainty is computed as estimation uncertainty about bag locations as a fraction of total uncertainty that is sum the of estimated uncertainty and noise.

$$RU = \tau_{t+1} = \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_N^2}$$

Oddball condition:

The probability of an observation being an oddball is computed in a similar manner using the prediction error of the model on each trial:

$$\text{OBP} = \Omega_t = \frac{U(X_t|0.300)\, H}{U(X_t|0.300)H + N(\delta_t; 0.\sigma_\mu^2 + \sigma_N^2)(1-H)}$$

Where H is the hazard rate of the oddball condition. Here, the numerator is the likelihood of current outcome being an oddball and the second term in the denominator is the likelihood of current outcome coming from a normal distribution with the mean of predicted outcome and variance of total uncertainty.

The estimation uncertainty is computed similar to change point condition, albeit with a few modifications:

$$\sigma_{\mu}^{2} = \Omega_t \frac{\sigma_N^2 \tau_t}{1-\tau_t} + (1-\Omega_t)\,\tau_t \sigma_N^2 + \Omega_t (1-\Omega_t)(\delta_t \tau_t)^2 + \sigma_{drift}^2$$

Where the first term represents the contribution of an oddball observation to the uncertainty, the second term is the contribution of a nonoddball observation to the uncertainty, the third term is uncertainty proportional to the prediction error received from either an oddball or a nonodball trial and the last term reflects the uncertainty expected from the drift rate of the helicopter in between two trials.

Again similar to changepoint condition, relative uncertainty is computed as estimation uncertainty over total uncertainty:

$$RU = \tau_{t+1} = \frac{\sigma_{\mu}^2}{\sigma_{\mu}^2 + \sigma_N^2}$$

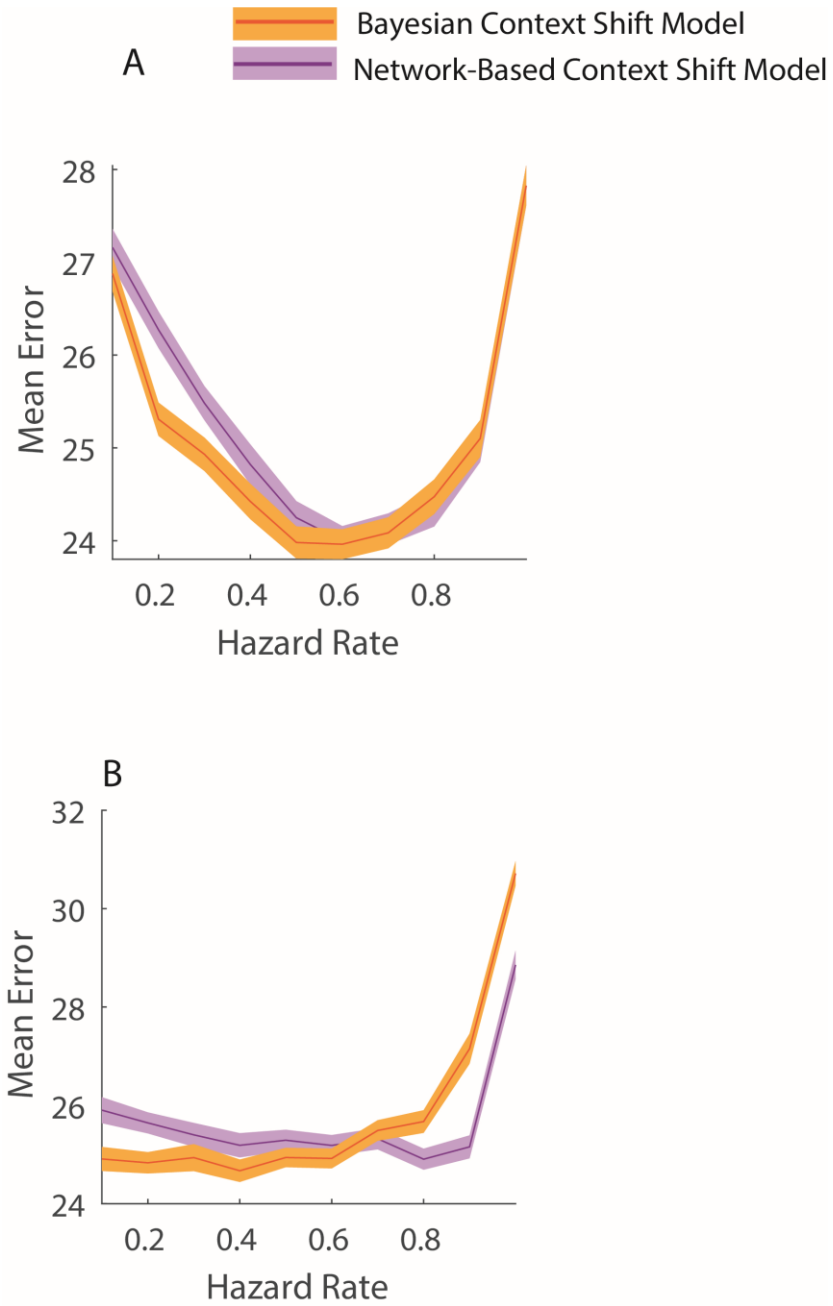

Figure 3-1 Performance of the Bayesian and network-based dynamic context shift models with different levels of hazard rate for both (A) changepoint and (B) oddball conditions. A fixed level of 0.7 hazard rate was used in both models for all the analysis in the paper.

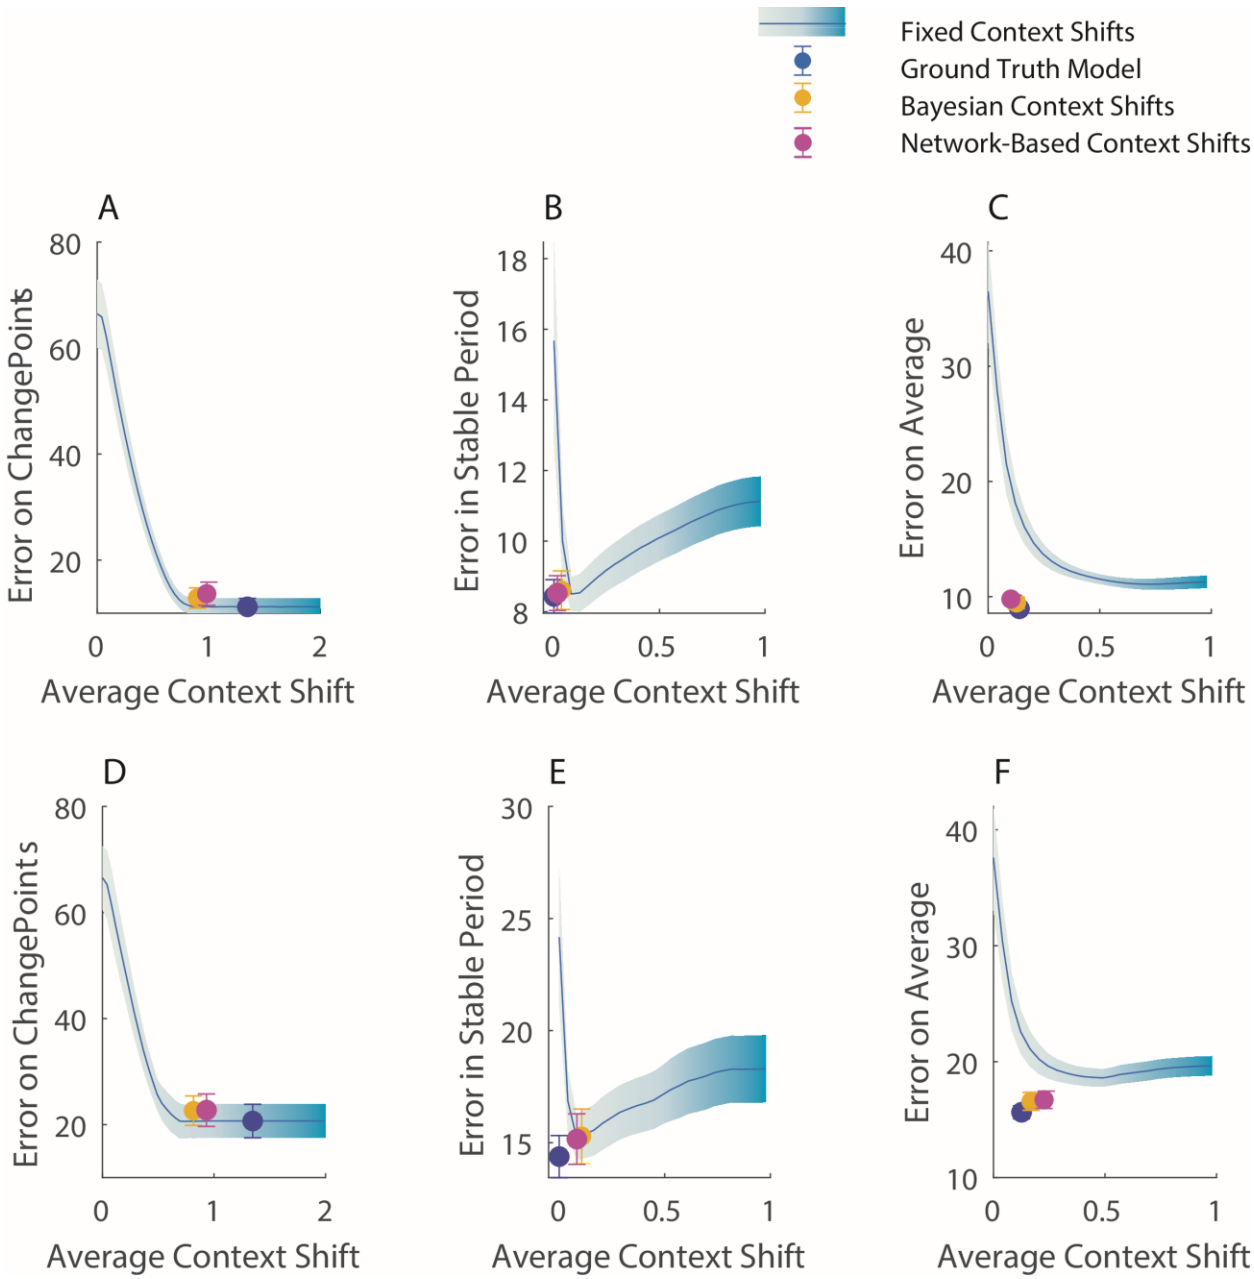

Figure 3- 2. The same analysis as figure 3 c-e with different noise levels.
Performance of the three types of models with a lower noise parameter (standard deviation of dropped bags locations around helicopter position), here noise=10 on A) trials immediately following a changepoint B) trials during periods of stability (> 5 trials after a changepoint) C) Averaged across all trials, the performance of the dynamic context shift models are better than the best fixed context shift model ($Bayesian\ Context\ Shift: t = 21.9. df = 31. p < 10^{-16}$ D-F. Network-Based Context Shift: t = 20.48, , $df = 31. p < 10^{-16}$ D-F). Similar analysis with a combination of the two different noise levels in different blocks (noise= 10 or 25) the same as what was used in the behavioral tasks with human participants in (McGuire et al., 2014). The same analysis for alternating noise levels condition ($Bayesian\ Context\ Shift: t = 20.62. df = 31. p < 10^{-16}. Network - Based\ Context\ Shift: = t = 15.91. df = 31. p < 10^{-16}$)) shows that the predictions of the dynamic model are consistent across different noise levels.

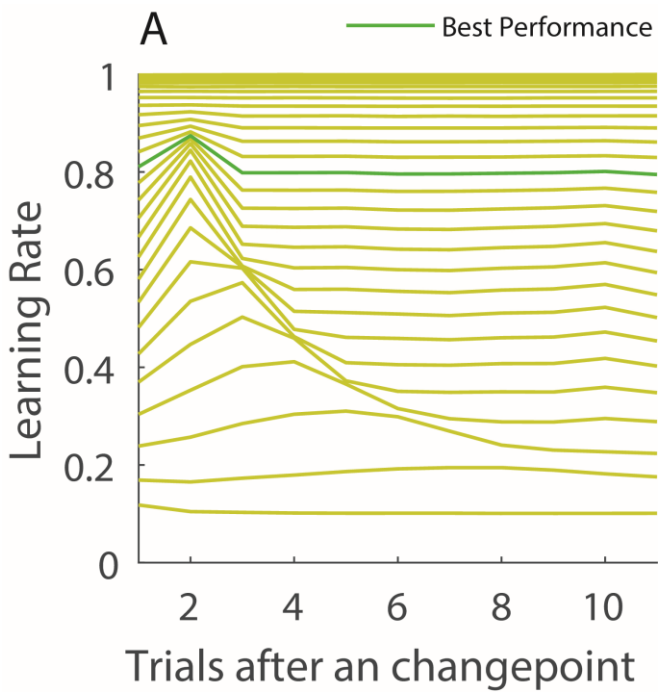

Figure 4-1: Effective learning rate dynamics for different fixed context shift models. A) In extreme cases where the context shift is close to zero or one, effective learning rate is independent of trial after a changepoint. However, with moderate levels of context shift, the model shows increased learning one (or a few) trial(s) after a changepoint which is consistent with the fact that it takes a few time steps for the model to dissociate from previous context and transition to a new context.