

1 Gated recurrence enables simple and accurate sequence prediction 2 in stochastic, changing, and structured environments

3 Cedric Foucault^{1,2,*} and Florent Meyniel^{1,*}

4 1. Cognitive Neuroimaging Unit, INSERM, CEA, Université Paris-Saclay, NeuroSpin center, 91191 Gif/Yvette, France.

5 2. Sorbonne Université, Collège Doctoral, F-75005 Paris, France.

6 * Corresponding authors: cedric.foucault@gmail.com and florent.meyniel@cea.fr

7 **Abstract**

8 From decision making to perception to language, predicting what is coming next is crucial. It is also
9 challenging in stochastic, changing, and structured environments; yet the brain makes accurate predictions
10 in many situations. What computational architecture could enable this feat? Bayesian inference makes
11 optimal predictions but is prohibitively difficult to compute. Here, we show that a specific recurrent neural
12 network architecture enables simple and accurate solutions in several environments. This architecture relies
13 on three mechanisms: gating, lateral connections, and recurrent weight training. Like the optimal solution
14 and the human brain, such networks develop internal representations of their changing environment
15 (including estimates of the environment's latent variables and the precision of these estimates), leverage
16 multiple levels of latent structure, and adapt their effective learning rate to changes without changing their
17 connection weights. Being ubiquitous in the brain, gated recurrence could therefore serve as a generic
18 building block to predict in real-life environments.

19 Introduction

20 Being able to correctly predict what is coming next is advantageous: it enables better decisions (Dolan &
21 Dayan, 2013; R. S. Sutton & Barto, 1998), a more accurate perception of our world, and faster reactions
22 (De Lange et al., 2018; Dehaene et al., 2015; Saffran et al., 1996; Sherman et al., 2020; Summerfield & de
23 Lange, 2014). In many situations, predictions are informed by a sequence of past observations. In that case,
24 the prediction process formally corresponds to a statistical inference that uses past observations to estimate
25 latent variables of the environment (e.g. the probability of a stimulus) that then serve to predict what is likely
26 to be observed next. Specific features of real-life environments make this inference a challenge: they are
27 often partly random, changing, and structured in different ways. Yet, in many situations, the brain is able to
28 overcome these challenges and shows several aspects of the optimal solution (Dehaene et al., 2015; Dolan
29 & Dayan, 2013; Gallistel et al., 2014; Summerfield & de Lange, 2014). Here we aim to identify the
30 computational mechanisms that could enable the brain to exhibit these aspects of optimality in these
31 environments.

32 We start by unpacking two specific challenges which arise in real-life environments. First, the joint
33 presence of randomness and changes (i.e., the non-stationarity of the stochastic process generating the
34 observations) poses a well-known tension between stability and flexibility (Behrens et al., 2007; Soltani &
35 Izquierdo, 2019; R. Sutton, 1992). Randomness in observations requires integrating information over time
36 to derive a stable estimate. However, when a change in the estimated variable is suspected, it is better to
37 limit the integration of past observations to update the estimate more quickly. The prediction should thus be
38 adaptive, i.e. dynamically adjusted to promote flexibility in the face of changes and stability otherwise. Past
39 studies have shown that the brain does so in many contexts: perception (Fairhall et al., 2001; Wark et al.,
40 2009), homeostatic regulation (Pezzulo et al., 2015; Sterling, 2004), sensorimotor control (Berniker &
41 Kording, 2008; Wolpert et al., 1995), and reinforcement learning (Behrens et al., 2007; Iglesias et al., 2013;
42 Soltani & Izquierdo, 2019; R. S. Sutton & Barto, 1998).

43 Second, the structure of our environment can involve complex relationships. For instance, the sentence
44 beginnings "what science can do for you is ..." and "what you can do for science is ..." call for different
45 endings even though they contain the same words, illustrating that prediction takes into account the ordering
46 of observations. Such structures appear not only in human language but also in animal communication
47 (Dehaene et al., 2015; Hauser et al., 2001; Robinson, 1979; Rose et al., 2004), and all kinds of stimulus-
48 stimulus and stimulus-action associations in the world (Saffran et al., 1996; Schapiro et al., 2013; Soltani &
49 Izquierdo, 2019; R. S. Sutton & Barto, 1998). Such a structure is often latent (i.e. not directly observable)
50 and it governs the relationship between observations (e.g. words forming a sentence, stimulus-action
51 associations). These relationships must be leveraged by the prediction, making it more difficult to compute.

52 In sum, the randomness, changes, and latent structure of real-life environments pose two major
53 challenges: that of adapting to changes and that of leveraging the latent structure. Two commonly used
54 approaches offer different solutions to these challenges. The Bayesian approach allows to derive statistically
55 optimal predictions for a given environment knowing its underlying generative model. This optimal solution
56 is a useful benchmark and has some descriptive validity since, in some contexts, organisms behave close
57 to optimally (Ma & Jazayeri, 2014; Tauber et al., 2017) or exhibit several qualitative aspects of the optimal
58 solution (Behrens et al., 2007; Heilbron & Meyniel, 2019; Meyniel et al., 2015). However, a specific Bayes-
59 optimal solution only applies to a specific generative model (or class of models (Tenenbaum et al., 2011)).
60 This mathematical solution also does not in general lead to an algorithm of reasonable complexity (Cooper,
61 1990; Dagum & Luby, 1993). Bayesian inference therefore says little about the algorithms that the brain
62 could use, and the biological basis of those computations remains mostly unknown with only a few proposals
63 highly debated (Fiser et al., 2010; Ma et al., 2006; Sahani & Dayan, 2003).

64 Opposite to the Bayes-optimal approach is the heuristics approach: solutions that are easy to compute
65 and accurate in specific environments (Todd & Gigerenzer, 2000). However, heuristics lack generality: their
66 performance can be quite poor outside the environment that suits them. In addition, although simple, their
67 biological implementation often remains unknown (besides the delta-rule (Eshel et al., 2013; Rescorla &
68 Wagner, 1972; Schultz et al., 1997)).

69 Those two approaches leave open the following questions: Is there a general, biologically feasible
70 architecture that enables, in different environments, solutions that are simple, effective, and that reproduce
71 the qualitative aspects of optimal prediction observed in organisms? If so, what are its essential mechanistic
72 elements?

73 Our approach stands in contrast with the elegant closed-form but intractable mathematical solutions
74 offered by Bayesian inference, and the simple but specialized algorithms offered by heuristics. Instead, we
75 look for general mechanisms under the constraints of feasibility and simplicity. We used recurrent neural
76 networks because they can offer a generic, biologically feasible architecture able to realize different
77 prediction algorithms (see (LeCun et al., 2015; Saxe et al., 2021) and Discussion). We used small network
78 sizes in order to produce simple (i.e. low-complexity, memory-bounded) solutions. We tested their generality
79 using different environments. To determine the simplest architecture sufficient for effective solutions and
80 derive mechanistic insights, we considered different architectures that varied in size and mechanisms. For
81 each one, we instantiated several networks and trained them to approach their best possible prediction
82 algorithm in a given environment. We treated the training procedure as a methodological step without
83 claiming it to be biologically plausible. To provide interpretability, we inspected the networks' internal model
84 and representations, and tested specific optimal aspects of their behavior—previously reported in humans
85 (Heilbron & Meyniel, 2019; Meyniel et al., 2015; Nassar et al., 2010, 2012)—which demonstrate the ability
86 to adapt to changes and leverage the latent structure of the environment.

87 **Results**

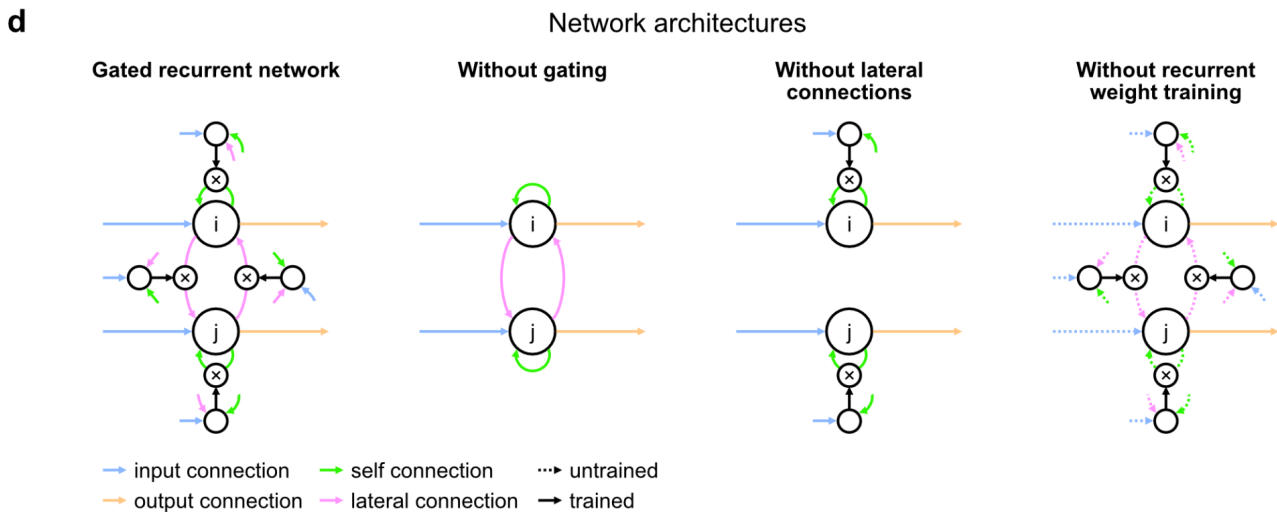

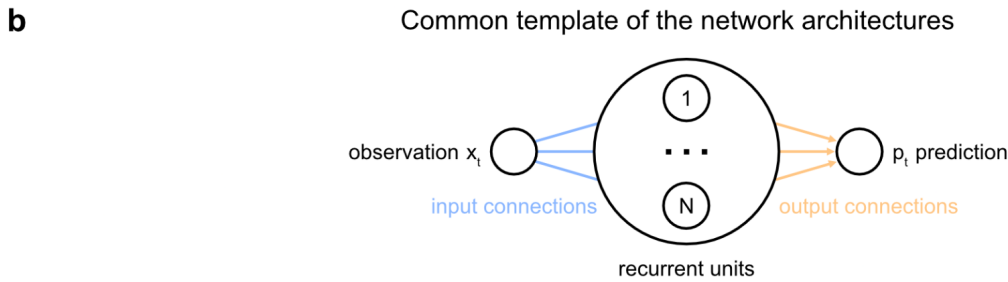

88 **The framework: sequence prediction and network architectures**

89 All our analyses confront simulated agents with the same general problem: sequence prediction. It
90 consists in predicting, at each time step in a sequence where one time step represents one observation, the
91 probability distribution over the value of the next observation given the previous observations (here we used
92 binary observations coded as '0' and '1') (**Fig. 1a**). The environment generates the sequence, and the
93 agent's goal is to make the most accurate predictions possible in this environment. Below, we introduce
94 three environments. All of them are stochastic (observations are governed by latent probabilities) and
95 changing (these latent probabilities change across time), and thus require dynamically adapting the stability-
96 flexibility tradeoff. They also feature increasing levels of latent structure that must be leveraged, making the
97 computation of predictions more complex.

98 How do agents learn to make predictions that fit a particular environment? In real life, agents often do
99 not benefit from any external supervision and must rely only on the observations. To do so, they can take
100 advantage of an intrinsic error signal that measures the discrepancy between their prediction and the actual
101 value observed at the next time step. We adopted this learning paradigm (often called unsupervised, self-
102 supervised, or predictive learning in machine learning (Elman, 1991; LeCun, 2016)) to train our agents in
103 silico. We trained the agents by exposing them to sequences generated by a given environment and letting
104 them adjust their parameters to improve their prediction (see Methods).

105 During testing, we kept the parameters of the trained agents frozen, exposed them to new sequences,
106 and performed targeted analyses to probe whether they exhibit specific capabilities and better understand
107 how they solve the problem.

108 Our investigation focuses on a particular class of agent architectures known as recurrent neural networks.
109 These are well suited for sequence prediction because recurrence allows to process inputs sequentially
110 while carrying information over time in recurrent activity. The network architectures we used all followed the
111 same three-layer template, consisting of one input unit whose activity codes for the current observation, one
112 output unit whose activity codes for the prediction about the next observation, and a number of recurrent
113 units that are fed by the input unit and project to the output unit (**Fig. 1b**). All architectures had self-recurrent
114 connections.



115

116 **Figure 1. Problem to solve and network architectures.** (a) Sequence prediction problem. At each time step t , the environment
 117 generates one binary observation x_t . The agent receives it and returns a prediction p_t : its estimate of the probability that the next
 118 observation will be 1 given the observations collected so far. The agent's goal is to make the most accurate predictions possible.
 119 The agent can measure its accuracy by comparing its prediction p_t with the actual value observed at the next time step x_{t+1} ,
 120 allowing it to learn from the observations without any external supervision. (b) Common three-layer template of the recurrent
 121 neural network architectures. Input connections transmit the observation to the recurrent units and output connections allow the
 122 prediction to be read from the recurrent units. (c) Three key mechanisms of recurrent neural network architectures. Gating allows
 123 for multiplicative interaction between activities. Lateral connections allow the activities of different recurrent units i and j to interact.
 124 Recurrent weight training allows the connection weights of recurrent units to be adjusted to the training environment. i' may be
 125 equal to i . (d) The gated recurrent architecture includes all three mechanisms: gating, lateral connections, and recurrent weight
 126 training. Each alternative architecture includes all but one of the three mechanisms.

127 **Figure supplement 1. Graphical model of the generative process of each environment.**

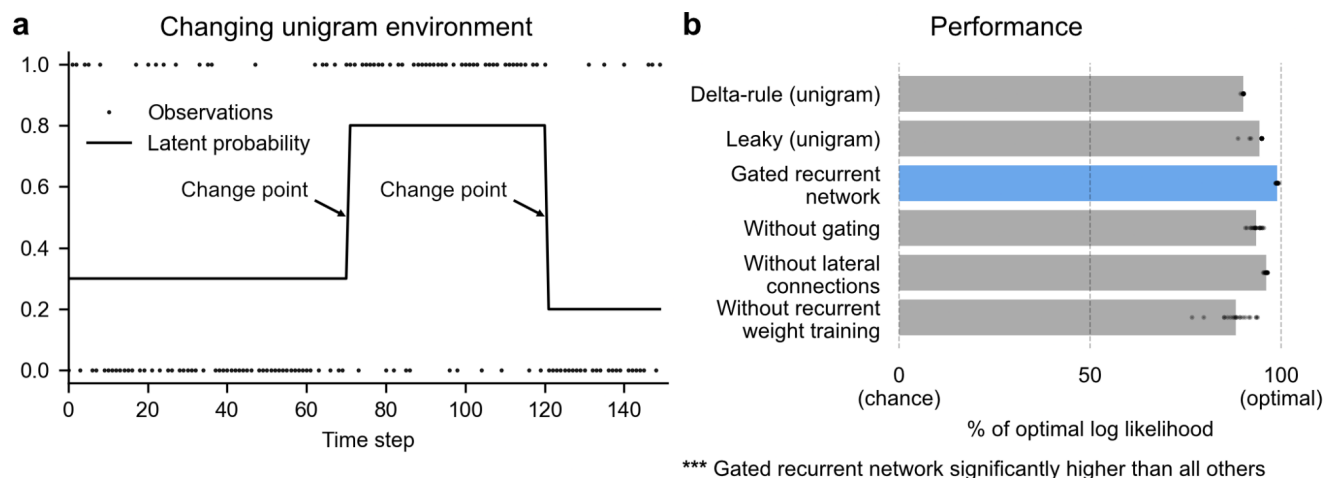
128 We identified three mechanisms of recurrent neural network architectures that endow a network with
129 specific computational properties which have proven advantageous in our environments (**Fig. 1c**). One
130 mechanism is gating, which allows for multiplicative interactions between the activities of units. A second
131 mechanism is lateral connectivity, which allows the activities of different recurrent units to interact with each
132 other. A third mechanism is the training of recurrent connection weights, which allows the dynamics of
133 recurrent activities to be adjusted to the training environment.

134 To get mechanistic insight, we compared an architecture that included all three mechanisms, to
135 alternative architectures that were deprived of one of the three mechanisms but maintained the other two
136 (**Fig. 1d**; see Methods for equations). Here we call an architecture with all three mechanisms ‘gated
137 recurrent’, and the particular architecture we used is known as GRU (Cho et al., 2014; Chung et al., 2014).
138 When deprived of gating, multiplicative interactions between activities are removed, and the architecture
139 reduces to that of a vanilla recurrent neural network also known as the Elman network (Elman, 1990). When
140 deprived of lateral connections, the recurrent units become independent of each other, thus each recurrent
141 unit acts as a temporal filter on the input observations (with possibly time-varying filter weights thanks to
142 gating). When deprived of recurrent weight training, the recurrent activity dynamics become independent of
143 the environment and the only parameters that can be trained are those of the output unit; this architecture
144 is thus one form of reservoir computing (Tanaka et al., 2019). In the results below, unless otherwise stated,
145 the networks all had 11 recurrent units (the smallest network size beyond which the gated recurrent network
146 showed no substantial increase in performance in any of the environments), but the results across
147 architectures are robust to this choice of network size (see the last section of the Results).

148 **Performance in the face of changes in latent probabilities**

149 We designed a first environment to investigate the ability to handle changes in a latent probability (**Fig.**
150 **2a**; see **Fig. 1—figure supplement 1** for a graphical model). In this environment we used the simplest kind
151 of latent probability: $p(1)$, the probability of occurrence (or base rate) of the observation being 1 (note that
152 $p(0)=1-p(1)$), here called ‘unigram probability’. The unigram probability suddenly changed from one value to
153 another at so-called ‘change points’, which could occur at any time, randomly with a given fixed probability.

154 This environment, here called ‘changing unigram environment’, corresponds for instance to a simple
155 oddball task (Aston-Jones et al., 1997; Kaliukhovich & Vogels, 2014; Ulanovsky et al., 2004), or the
156 probabilistic delivery of a reward with abrupt changes in reward probabilities (Behrens et al., 2007; Vinckier
157 et al., 2016). In such an environment, predicting accurately is difficult due to the stability-flexibility tradeoff
158 induced by the stochastic nature of the observations (governed by the unigram probability) and the possibility
159 of a change point at any moment.



160

161 **Figure 2. Gated recurrent networks perform quasi-optimally in the face of changes in latent probabilities.** (a) Sample
 162 sequence of observations (dots) and latent unigram probability (line) generated in the changing unigram environment. At each
 163 time step, a binary observation is randomly generated based on the latent unigram probability, and a change point can occur with
 164 a fixed probability, suddenly changing the unigram probability to a new value uniformly drawn in [0,1]. (b) Prediction performance
 165 in the changing unigram environment. For each type of agent, 20 trained agents (trained with different random seeds) were tested
 166 (dots: agents; bars: average). Their prediction performance was measured as the % of optimal log likelihood (0% being chance
 167 performance and 100% optimal performance, see equation (1) for the log likelihood) and averaged over observations and
 168 sequences. The gated recurrent network significantly outperformed every other type of agent ($p < 0.001$, two-tailed two
 169 independent samples t-test with Welch's correction for unequal variances).

170

171 To assess the networks' prediction accuracy, we compared the networks with the optimal agent for this
 172 specific environment, i.e. the optimal solution to the prediction problem determined using Bayesian
 173 inference. This optimal solution knows the environment's underlying generative process and uses it to
 174 compute, via Bayes' rule, the probability distribution over the possible values of the latent probability given
 175 the past observation sequence, $p(p_{t+1}^{env} | x_0, \dots, x_t)$, known as the posterior distribution. It then outputs as
 176 prediction the mean of this distribution. (For details see Methods and (Heilbron & Meyniel, 2019).)

177 We also compared the networks to two types of heuristics which perform very well in this environment:
 178 the classic 'delta-rule' heuristic (Rescorla & Wagner, 1972; R. S. Sutton & Barto, 1998) and the more
 179 accurate 'leaky' heuristic (Gijssen et al., 2021; Heilbron & Meyniel, 2019; Meyniel et al., 2016; Yu & Cohen,
 180 2008) (see Methods for details). To test the statistical reliability of our conclusions, we trained separately 20
 181 agents of each type (each type of network and each type of heuristic).

182 We found that even with as few as 11 units, the gated recurrent networks performed quasi-optimally.
 183 Their prediction performance was 99% of optimal (CI $\pm 0.1\%$), 0% corresponding to chance level (**Fig. 2b**).
 184 Being only 1% short of optimal, the gated recurrent networks outperformed the delta rule and leaky agents,
 185 which performed 10 times and 5 times further from optimal, respectively (**Fig. 2b**).

186 For mechanistic insight, we tested the alternative architectures deprived of one mechanism. Without
 187 either gating, lateral connections, or recurrent weight training, the average performance was respectively 6

188 times, 4 times, and 12 times further from optimal (**Fig. 2b**), i.e. the level of a leaky agent or worse. The drops
189 in performance remain similar when considering only the best network of each architecture instead of the
190 average performance (**Fig. 2b**, compare rightmost dots across rows).

191 These results show that small gated recurrent networks can achieve quasi-optimal predictions and that
192 the removal of one of the mechanisms of the gated recurrent architecture results in a systematic drop in
193 performance.

194 **Adaptation to changes through the adjustment of the effective learning rate**

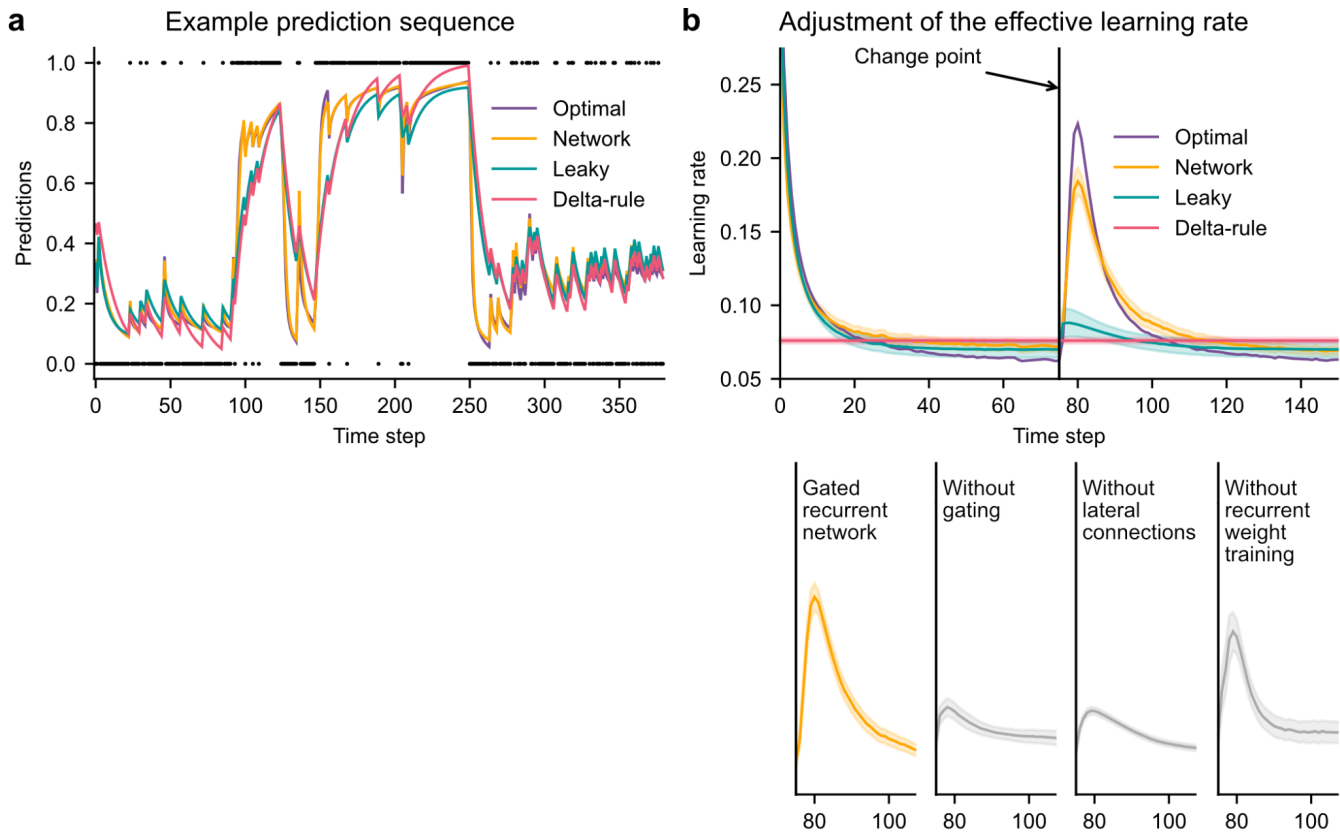
195 In a changing environment, the ability to adapt to changes is key. Networks exposed to more changing
196 environments during training updated their predictions more overall during testing, similarly to the optimal
197 agent (see **Fig. 3—figure supplement 1**) and, to some extent, humans (Behrens et al., 2007, Figure 2e;
198 Findling et al., 2021, Figure 4c). At a finer timescale, the moment-by-moment updating of the predictions
199 also showed sensible dynamics around change points.

200 **Fig. 3a** illustrates a key difference in behavior between, on the one hand, the optimal agent and the gated
201 recurrent network, and on the other hand, the heuristic agents: the dynamics of their update differ. This
202 difference is particularly noticeable when recent observations suggest that a change point has just occurred:
203 the optimal agent quickly updates the prediction by giving more weight to the new observations; the gated
204 recurrent network behaves the same but not the heuristic agents. We formally tested this dynamic updating
205 around change points by measuring the moment-by-moment effective learning rate, which normalizes the
206 amount of update in the prediction by the prediction error (i.e. the difference between the previous prediction
207 and the actual observation; see Methods, equation (2)).

208 Gated recurrent networks turned out to adjust their moment-by-moment effective learning rate as the
209 optimal agent did, showing the same characteristic peaks, at the same time and with almost the same
210 amplitude (**Fig. 3b**, top plot). By contrast, the effective learning rate of the delta-rule agents was (by
211 construction) constant, and that of the leaky agents changed only marginally.

212 When one of the mechanisms of the gated recurrence was taken out, the networks' ability to adjust their
213 effective learning rate was greatly degraded (but not entirely removed) (**Fig. 3b**, bottom plots). Without
214 gating, without lateral connections, or without recurrent weight training, the amplitude was lower (showing
215 both a lower peak value and a higher baseline value), and the peak occurred earlier.

216 This shows that gated recurrent networks can reproduce a key aspect of optimal behavior: the ability to
217 adapt the update of their prediction to change points, which is lacking in heuristic agents and alternative
218 networks.



219

220 **Figure 3. Gated recurrent but not alternative networks adjust their moment-by-moment effective learning rate around**
 221 **changes like the optimal agent. (a)** Example prediction sequence illustrating the prediction updates of different types of agents.
 222 Within each type of agent, the agent (out of 20) yielding median performance in **Fig. 2b** was selected for illustration purposes.
 223 Dots are observations, lines are predictions. **(b)** Moment-by-moment effective learning rate of each type of agent. 20 trained
 224 agents of each type were tested on 10,000 sequences whose change points were locked at the same time steps, for illustration
 225 purposes. The moment-by-moment effective learning rate was measured as the ratio of prediction update to prediction error (see
 226 Methods, equation (2)), and averaged over sequences. Lines and bands show the mean and the 95% confidence interval of the
 227 mean.

228 **Figure supplement 1. Attunement of the effective learning rate to the change point probabilities.**

229

230 Internal representation of precision and dynamic interaction with the prediction

231 Beyond behavior, we sought to determine whether a network's ability to adapt to changes relied on
 232 idiosyncratic computations or followed the more general principle of precision-weighting derived from
 233 probability theory. According to this principle, the precision of the current prediction (calculated in the optimal
 234 agent as the negative logarithm of the standard deviation of the posterior distribution over the latent
 235 probability, see equation (3) in Methods) should influence the weight of the current prediction relative to the
 236 next observation in the updating process: for a given prediction error, the lower the precision, the higher the
 237 subsequent effective learning rate. This precision-weighting principle results in an automatic adjustment of
 238 the effective learning rate in response to a change, because the precision of the prediction decreases when
 239 a change is suspected.

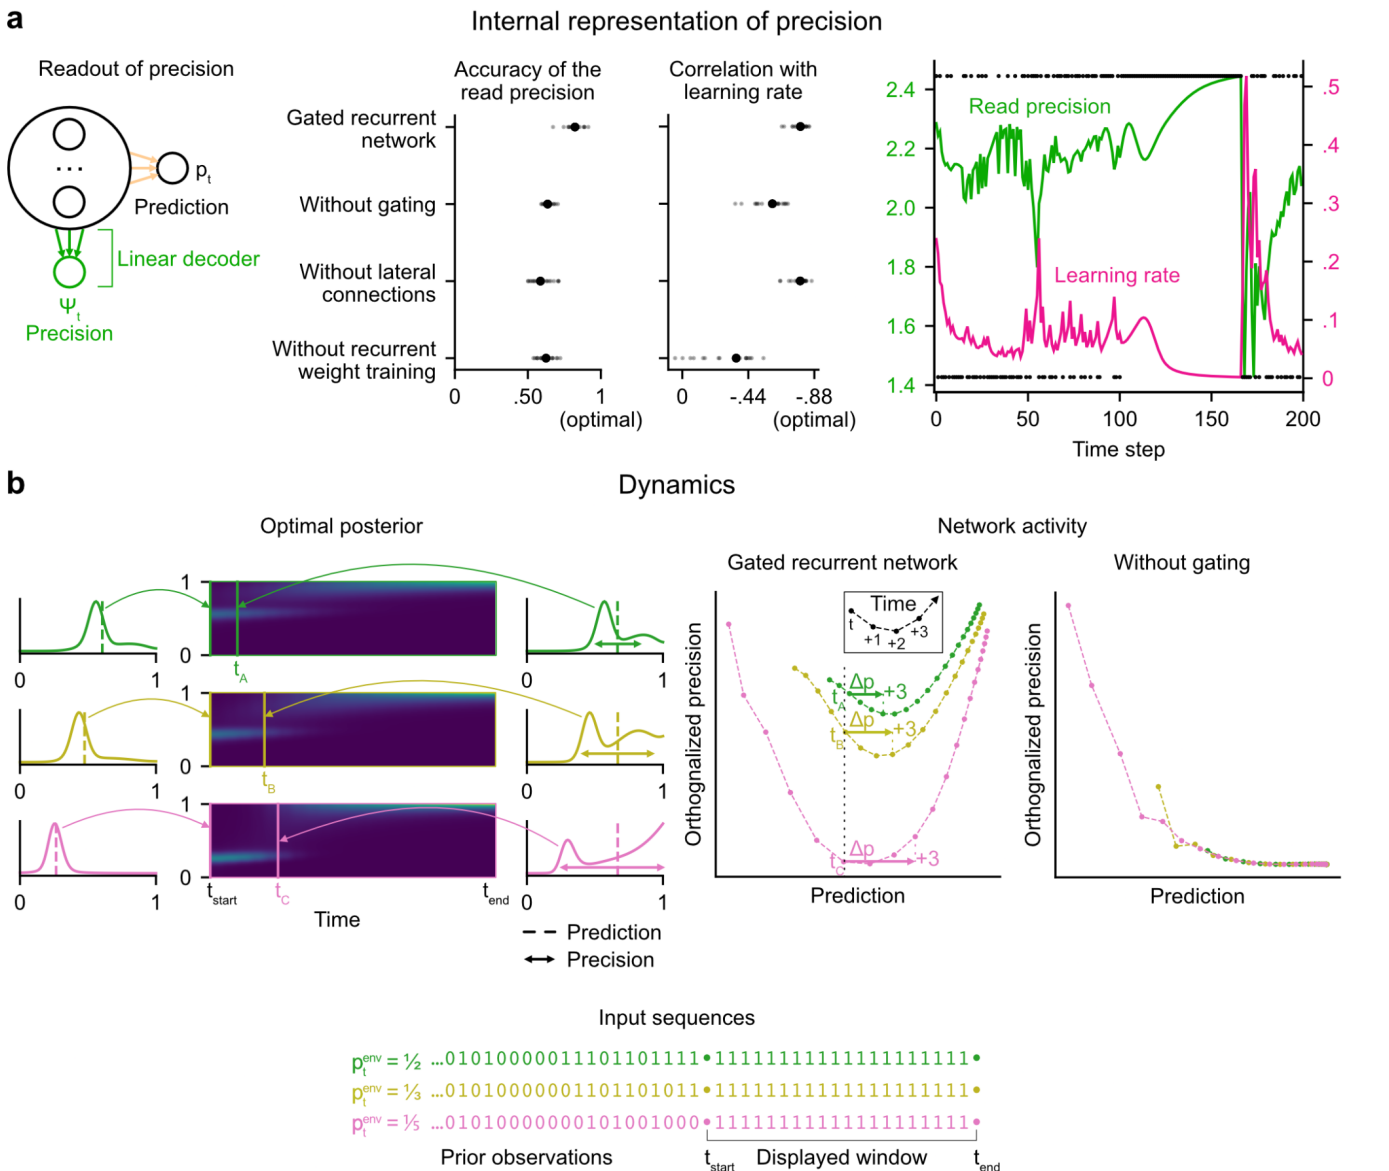
240 In line with this principle, human participants can estimate not only the prediction but also its precision as
241 estimated by the optimal agent (Boldt et al., 2019, Figure 2; Meyniel et al., 2015, Figure 4B), and this
242 precision indeed relates to the participants' effective learning rate (McGuire et al., 2014, Figure 2C and S1A;
243 Nassar et al., 2010, Figure 4C and 3B, 2012, Figure 5 and 7c).

244 We tested whether a network could represent this optimal precision too, by trying to linearly read it from
245 the network's recurrent activity (**Fig. 4a**). Note that the networks were trained only to maximize prediction
246 accuracy (not to estimate precision). Yet, in gated recurrent networks, we found that the read precision on
247 left-out data was highly accurate (**Fig. 4a**, left plot: the median Pearson correlation with the optimal precision
248 is 0.82), and correlated with their subsequent effective learning rate as in the optimal agent (**Fig. 4a**, right
249 plot: the median correlation for gated recurrent networks is -0.79; for comparison, it is -0.88 for the optimal
250 agent).

251 To better understand how precision information is represented and how it interacts with the prediction
252 dynamically in the network activity, we plotted the dynamics of the network activity in the subspace spanned
253 by the prediction and precision vectors (**Fig. 4b**). Such visualization captures both the temporal dynamics
254 and the relationships between the variables represented in the network, and has helped understand network
255 computations in other works (Mante et al., 2013; Sohn et al., 2019). Here, two observations can be made.

256 First, in the gated recurrent network (**Fig. 4b**, second plot from the right), the trajectories are well
257 separated along the precision axis (for the same prediction, the network can represent multiple precisions),
258 meaning that the representation of precision is not reducible to the prediction. By contrast, in the network
259 without gating (**Fig. 4b**, rightmost plot), these trajectories highly overlap, which indicates that the
260 representation of precision and prediction are mutually dependent. To measure this dependence, we
261 computed the mutual information between the read precision and the prediction of the network, and it turned
262 out to be very high in the network without gating (median MI=5.2) compared to the gated recurrent network
263 (median MI=0.7) and the optimal agent (median MI=0.6) (without lateral connections, median MI=1.3;
264 without recurrent weight training, median MI=1.9), confirming that gating is important to separate the
265 precision from the prediction.

266 Second, in the gated recurrent network, the precision interacts dynamically with the prediction in a
267 manner consistent with the precision-weighting principle: for a given prediction, the lower the precision, the
268 larger the subsequent updates to the prediction (**Fig. 4b**, vertical dotted line indicates the level of prediction
269 and arrows the subsequent updates).



270

271 **Figure 4. Gated recurrent networks have an internal representation of the precision of their estimate that dynamically**
 272 **interacts with the prediction following the precision-weighting principle. (a)** Left to right: Schematic of the readout of
 273 precision from the recurrent activity of a network (obtained by fitting a multiple linear regression from the recurrent activity to the
 274 log precision of the optimal posterior distribution); Accuracy of the read precision (calculated as its Pearson correlation with the
 275 optimal precision); Pearson correlation between the read precision and the network's subsequent effective learning rate (the
 276 optimal value was calculated from the optimal agent's own precision and learning rate); Example sequence illustrating their anti-
 277 correlation in the gated recurrent network. In both dot plots, large and small dots show the median and individual values,
 278 respectively. **(b)** Dynamics of the optimal posterior (left) and the network activity (right) in three sequences (green, yellow, and
 279 pink). The displayed dynamics are responses to a streak of 1s after different sequences of observations (with different generative
 280 probabilities as shown at the bottom). The optimal posterior distribution is plotted as a color map over time (dark blue and light
 281 green correspond to low and high probability densities, respectively) and as a line plot at two times: on the left, the time t_{start} just
 282 before the streak of 1s, and on the right, a time $t_A/t_B/t_C$ when the prediction (i.e. mean) is approximately equal in all three cases;
 283 note that the precision differs. The network activity was projected onto the two-dimensional subspace spanned by the prediction and
 284 precision vectors (for the visualization, the precision axis was orthogonalized with respect to the prediction axis). In the gated
 285 recurrent network, the arrow Δp shows the update to the prediction performed in the next three time steps starting at the time
 286 $t_A/t_B/t_C$ defined from the optimal posterior. Like the optimal posterior and unlike the network without gating, the gated recurrent
 287 network represents different levels of precision at an equal prediction, and the lower the precision, the higher the subsequent
 288 update to the prediction—a principle called precision-weighting. In all example plots (a-b), the displayed network is the one of the
 289 20 that yielded the median read precision accuracy.

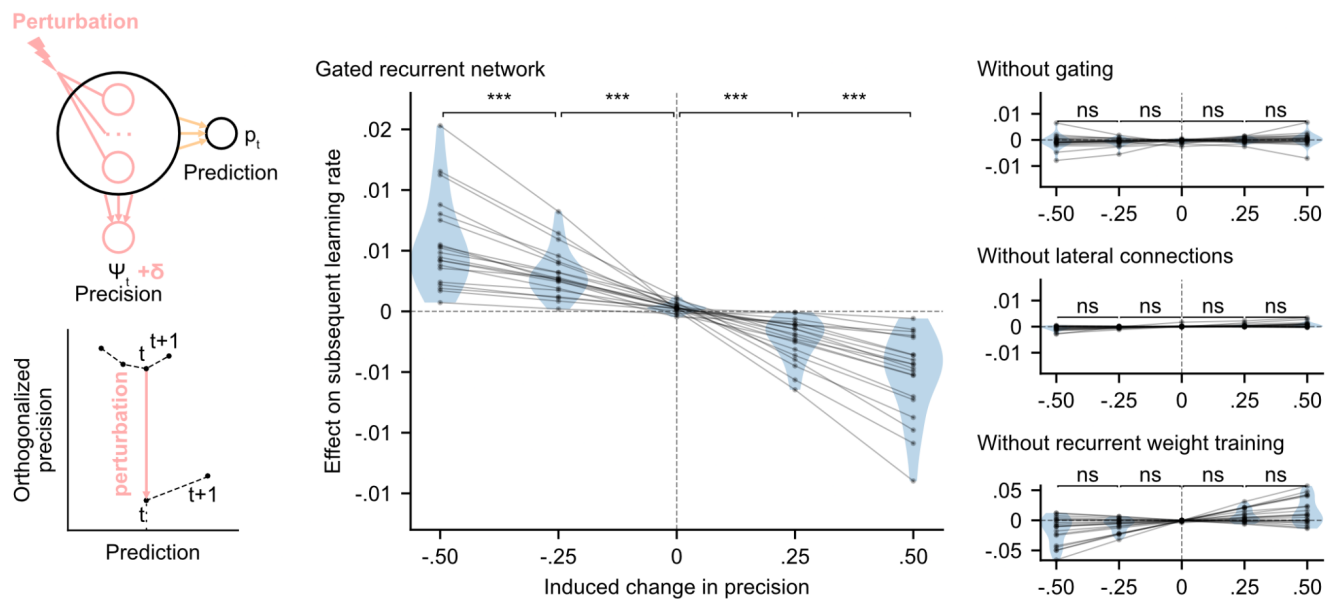
290 These results indicate that in the network without gating, precision is confounded with prediction and the
291 correlation between precision and effective learning rate is spuriously driven by the prediction itself, whereas
292 in the network with gating, there is a genuine representation of precision beyond the prediction itself, which
293 interacts with the updating of predictions. However, we have so far only provided correlational evidence; to
294 show that the precision represented in the network plays a causal role in the subsequent prediction update,
295 we need to perform an intervention that acts selectively on this precision.

296 **Causal role of precision-weighting for adaptation to changes**

297 We tested whether the internal representation of precision causally regulated the effective learning rate
298 in the networks using a perturbation experiment. We designed perturbations of the recurrent activity that
299 induced a controlled change in the read precision, while leaving the networks' current prediction unchanged
300 to control for the effect of the prediction error (for the construction of the perturbations, see **Fig. 5** bottom
301 left diagram and legend, and Methods). These perturbations caused significant changes in the networks'
302 subsequent effective learning rate, commensurate with the induced change in precision, as predicted by the
303 principle of precision-weighting (**Fig. 5**, middle plot). Importantly, this causal relationship was abolished in
304 the alternative networks that lacked one of the mechanisms of the gated recurrent architecture (**Fig. 5**, right
305 three plots; the slope of the effect was significantly different between the gated recurrent network group and
306 any of the alternative network groups, two-tailed two independent samples t-test, all $t(38) > 4.1$, all $p < 0.001$,
307 all Cohen's $d > 1.3$).

308 These results show that the gated recurrent networks' ability to adapt to changes indeed relies on their
309 precision-dependent updating and that such precision-weighting does not arise without all three
310 mechanisms of the gated recurrence.

Causal test of precision-weighting



311

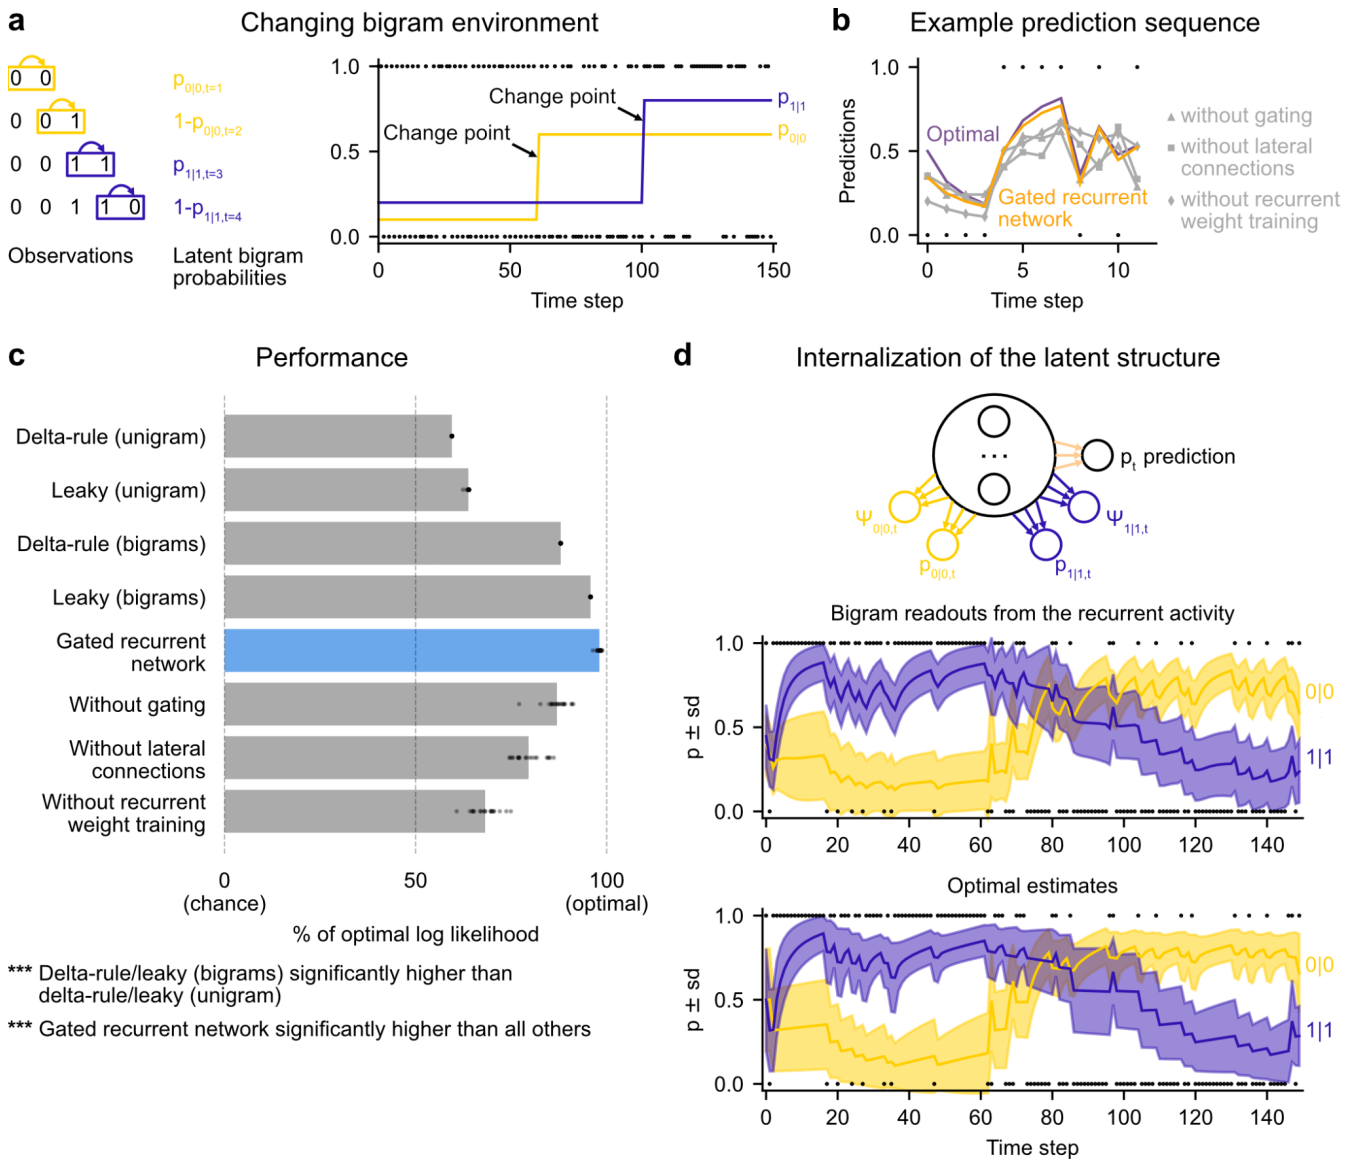
312 **Figure 5. Precision-weighting causally determines the adjustment of the effective learning rate in gated recurrent**
 313 **networks only.** Causal test of a network’s precision on its effective learning rate. The recurrent activity was perturbed to induce a
 314 controlled change δ in the read precision, while keeping the prediction at the current time step—and thus the prediction error at
 315 the next time step—constant. This was done by making the perturbation vector orthogonal to the prediction vector and making its
 316 projection onto the precision vector equal to δ (bottom left diagram). We measured the perturbation’s effect on the subsequent
 317 effective learning rate as the difference in learning rate “with perturbation” minus “without perturbation” at the next time step (four
 318 plots on the right). Each dot (and joining line) corresponds to one network. ***: $p < 0.001$, n.s.: $p > 0.05$ (one-tailed paired t-test).

319

320 Leveraging and internalizing a latent structure: bigram probabilities

321 While the changing unigram environment already covers many tasks in the behavioral and neuroscience
 322 literature, real-world sequences often exhibit more structure. To study the ability to leverage such structure,
 323 we designed a new stochastic and changing environment in which the sequence of observations is no longer
 324 generated according to a single unigram probability, $p(1)$, but two ‘bigram probabilities’ (also known as
 325 transition probabilities), $p(0|0)$ and $p(1|1)$, which denote the probability of occurrence of a 0 after a 0 and of
 326 a 1 after a 1, respectively (**Fig. 6a**; see **Fig. 1—figure supplement 1** for a graphical model). These bigram
 327 probabilities are also changing randomly, with independent change points.

328 This ‘changing bigram environment’ is well motivated because there is ample evidence that bigram
 329 probabilities play a key role in sequence knowledge in humans and other animals (Dehaene et al., 2015)
 330 even in the face of changes (Bornstein & Daw, 2013; Meyniel et al., 2015).



331

332 **Figure 6. Gated recurrent networks correctly leverage and internalize the latent bigram structure.** (a) Schematic of the
 333 changing bigram environment's latent probabilities (left) and sample generated sequence (right, dots: observations, lines: latent
 334 bigram probabilities). At each time step, a binary observation is randomly generated according to the relevant latent bigram
 335 probability, $p_{0|0}$ or $p_{1|1}$ depending on the previous observation. $p_{0|0}$ denotes the probability of occurrence of a 0 after a 0 and $p_{1|1}$
 336 that of a 1 after a 1 (note that $p_{1|0}=1-p_{0|0}$ and $p_{0|1}=1-p_{1|1}$). At any time step, each of the two bigram probabilities can suddenly
 337 change to a new value uniformly drawn in $[0,1]$, randomly with a fixed probability and independently from each other. (b) Example
 338 prediction sequence illustrating each network's ability or inability to change prediction according to the local context, compared to
 339 the optimal prediction (dots: observations, lines: predictions). (c) Prediction performance of each type of agent in the changing
 340 bigram environment. 20 new agents of each type were trained and tested as in Fig. 2b but now in the changing bigram
 341 environment (dots: agents; bars: average). The gated recurrent network significantly outperformed every other type of agent ($p <$
 342 0.001 , two-tailed two independent samples t-test with Welch's correction for unequal variances). (d) Internalization of the latent
 343 structure as shown on an out-of-sample sequence: the two bigram probabilities are simultaneously represented in the gated
 344 recurrent network (top), and closely follow the optimal estimates (bottom). The readouts were obtained through linear regression
 345 from the recurrent activity to four estimates separately: the log odds of the mean and the log precision of the optimal posterior
 346 distribution on $p_{0|0}$ and $p_{1|1}$. In (b) and (d), the networks (out of 20) yielding median performance were selected for illustration
 347 purposes.

348 **Figure supplement 1. Performance across training and test environments.**

349 We assessed how well the networks could leverage the latent bigram structure after having been trained
350 in this environment. For comparison, we tested the optimal agent for this environment as well as two groups
351 of heuristics: delta-rule and leaky estimation of unigram probabilities (as in **Fig. 2b**), and now also delta rule
352 and leaky estimation of bigram probabilities (see Methods for details).

353 The gated recurrent networks achieved 98% of optimal prediction performance (CI $\pm 0.3\%$), outperforming
354 the heuristic agents estimating bigram probabilities, and even more so those estimating a unigram
355 probability (**Fig. 6c**). To demonstrate that this was due to their internalization of the latent structure, we also
356 tested the gated recurrent networks that had been trained in the changing unigram environment: their
357 performance was much worse (**Fig. 6—figure supplement 1**).

358 At the mechanistic level, all three mechanisms of the gated recurrence are important for this ability to
359 leverage the latent bigram structure. Not only does the performance drop when one of these mechanisms
360 is removed (**Fig. 6c**), but also this drop in performance is much larger than that observed in the changing
361 unigram environment (without gating: -11.2% [CI $\pm 1.5\%$ calculated by Welch's t-interval] in the bigram
362 environment vs. -5.5% [CI $\pm 0.6\%$] in the unigram environment, without lateral connections: -18.5% [CI
363 $\pm 1.8\%$] vs. -2.9% [CI $\pm 0.2\%$]; without recurrent weight training: -29.9% [CI $\pm 1.6\%$] vs. -11.0% [CI $\pm 2.1\%$]; for
364 every mechanism, there was a significant interaction effect between the removal of the mechanism and the
365 environment on performance, all $F(1,76) > 47.9$, all $p < 0.001$).

366 **Fig. 6b** illustrates the gated recurrent networks' ability to correctly incorporate the bigram context into its
367 predictions compared to networks lacking one of the mechanisms of the gated recurrence. While a gated
368 recurrent network aptly changes its prediction from one observation to the next according to the preceding
369 observation as the optimal agent does, the other networks fail to show such context-dependent behavior,
370 sometimes even changing their prediction away from the optimal agent.

371 Altogether these results show that gated recurrent networks can leverage the latent bigram structure, but
372 this ability is impaired when one mechanism of the gated recurrence is missing.

373 Is the networks' representation of the latent bigram structure impenetrable or easily accessible? We
374 tested the latter possibility by trying to linearly read out the optimal estimate of each of the latent bigram
375 probabilities from the recurrent activity of a gated recurrent network (see Methods). Arguing in favor of an
376 explicit representation, we found that the read estimates of each of the latent bigram probabilities on left-out
377 data were highly accurate (Pearson correlation with the optimal estimates, median and CI: 0.97 [$0.97, 0.98$]
378 for each of the two bigram probabilities).

379 In addition to the point estimates of the latent bigram probabilities, we also tested whether a network
380 maintained some information about the precision of each estimate. Again, we assessed the possibility to
381 linearly read out the optimal precision of each estimate and found that the read precisions on left-out data

382 were quite accurate (Pearson correlation with the optimal precisions, median and CI: 0.77 [0.74, 0.78] for
383 one bigram probability and 0.76 [0.74, 0.78] for the other probability).

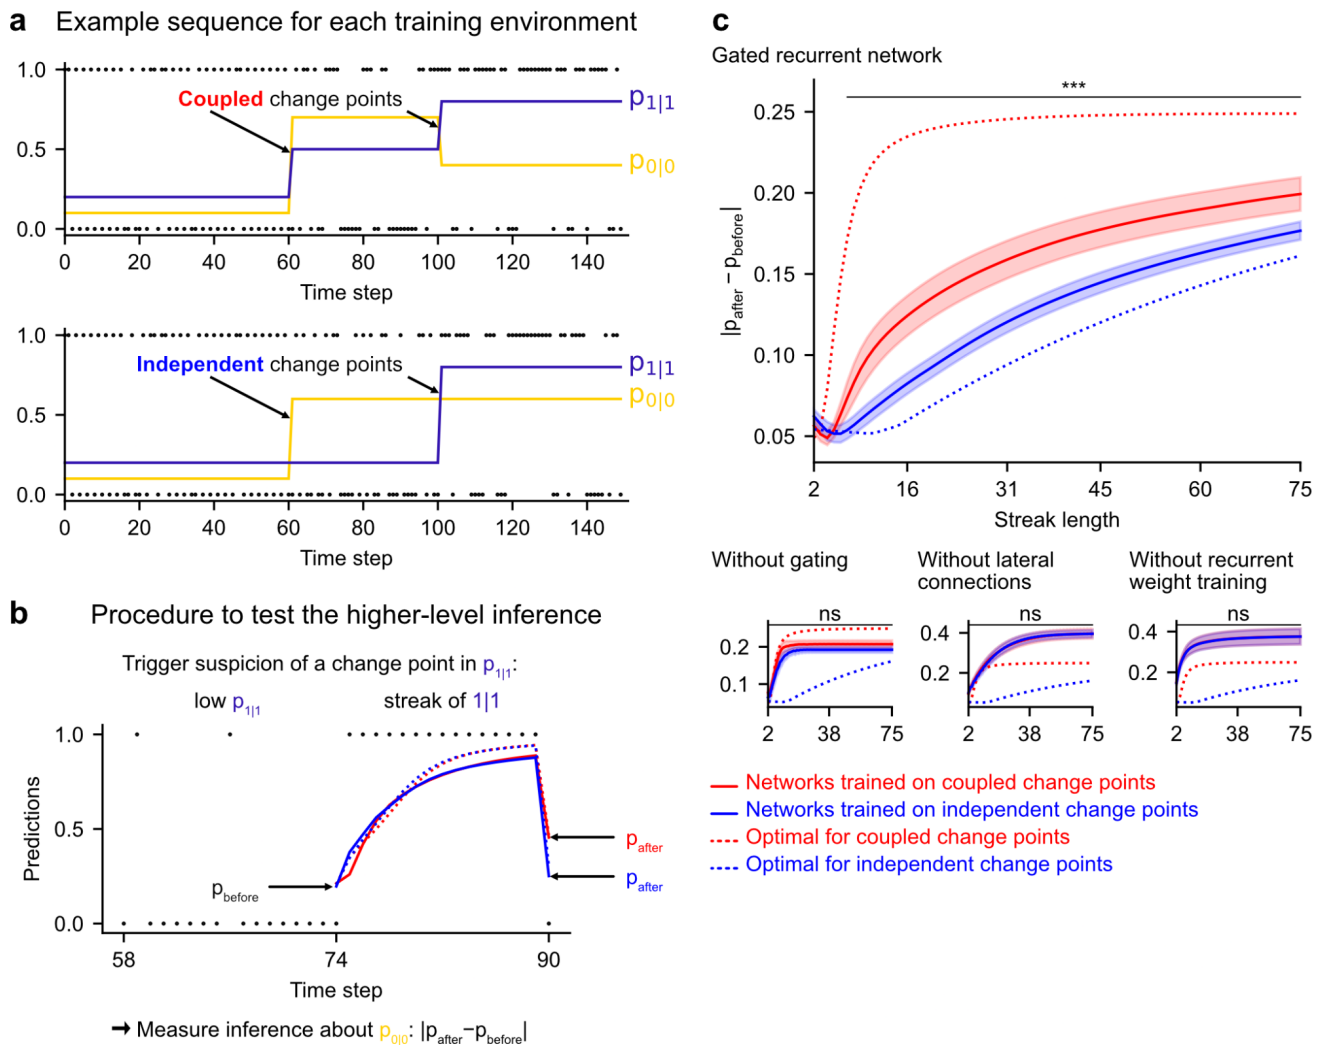
384 **Fig. 6d** illustrates the striking resemblance between the estimates read from a gated recurrent network
385 and the optimal estimates. Furthermore, it shows that the network successfully disentangles one bigram
386 probability from the other since the read estimates can evolve independently from each other (for instance
387 during the first 20 time steps, the value for 1|1 changes while the value for 0|0 does not, since only 1s are
388 observed). It is particularly interesting that both bigram probabilities are simultaneously represented, given
389 that only one of them is relevant for the moment-by-moment prediction read by the network's output unit
390 (whose weights cannot change during the sequence).

391 We conclude that gated recurrent networks internalize the latent bigram structure in such a way that both
392 bigram probabilities are available simultaneously, even though only one of the two is needed at any one
393 time for the prediction.

394 **Leveraging a higher-level structure: inference about latent changes**

395 In real life, latent structures can also exhibit different levels that are organized hierarchically (Bill et al.,
396 2020; Meyniel et al., 2015; Purcell & Kiani, 2016). To study the ability to leverage such a hierarchical
397 structure, we designed a third environment in which, in addition to bigram probabilities, we introduced a
398 higher-level factor: the change points of the two bigram probabilities are now coupled, rather than
399 independent as they were in the previous environment (**Fig. 7a**; **Fig. 1—figure supplement 1** shows the
400 hierarchical structure). Due to this coupling, from the agent's point of view, the likelihood that a change point
401 has occurred depends on the observations about both bigrams. Thus, optimal prediction requires the ability
402 to make a higher-level inference: having observed that the frequency of one of the bigrams has changed,
403 one should not only suspect that the latent probability of this bigram has changed but also transfer this
404 suspicion of a change to the latent probability of the other bigram, even without any observations about that
405 bigram.

406 Such a transfer has been reported in humans (Heilbron & Meyniel, 2019, Figure 5B). A typical situation
407 is when a streak of repetitions is encountered (**Fig. 7b**): if a long streak of 1s was deemed unlikely, it should
408 trigger the suspicion of a change point such that $p(1|1)$ is now high, and this suspicion should be transferred
409 to $p(0|0)$ by partially resetting it. This reset is reflected in the change between the prediction following the 0
410 just before the streak and that following the 0 just after the streak (**Fig. 7b**, $|p_{\text{after}} - p_{\text{before}}|$).



411

412 **Figure 7. Gated recurrent but not alternative networks leverage a higher-level structure, distinguishing the case where**
413 **change points are coupled vs. independent.** Procedure to test the higher-level inference: (a) For each network architecture, 20
414 networks were trained on sequences where the change points of the two latent bigram probabilities are coupled and 20 other
415 networks were trained on sequences where they are independent (the plots show an example training sequence for each case);
416 (b) The networks were then tested on sequences designed to trigger the suspicion of a change point in one bigram probability and
417 measure their inference about the other bigram probability: $|p_{\text{after}} - p_{\text{before}}|$ should be larger when the agent assumes change points
418 to be coupled rather than independent. The plot shows an example test sequence. Red, blue, solid, and dashed lines: as in (c),
419 except that only the gated recurrent network (out of 20) yielding median performance is shown for illustration purposes. (c)
420 Change in prediction about the unobserved bigram probability of the networks trained on coupled change points (red) and
421 independent change points (blue) for each network architecture, averaged over sequences. Solid lines and bands show the mean
422 and the 95% confidence interval of the mean over networks. Dotted lines show the corresponding values of the optimal agent for
423 the two cases. Only the gated recurrent architecture yields a significant difference between networks trained on coupled vs.
424 independent change points (one-tailed two independent samples t-test, ***: $p < 0.001$, n.s.: $p > 0.05$).

425

426 We tested the networks' ability for higher-level inference in the same way, by exposing them to such
427 streaks of repetitions and measuring their change in prediction about the unobserved bigram before and
428 after the streak. More accurately, we compared the change in prediction of the networks trained in the

429 environment with coupled change points to that of the networks trained in the environment with independent
430 change points, since the higher-level inference should only be made in the coupled case.

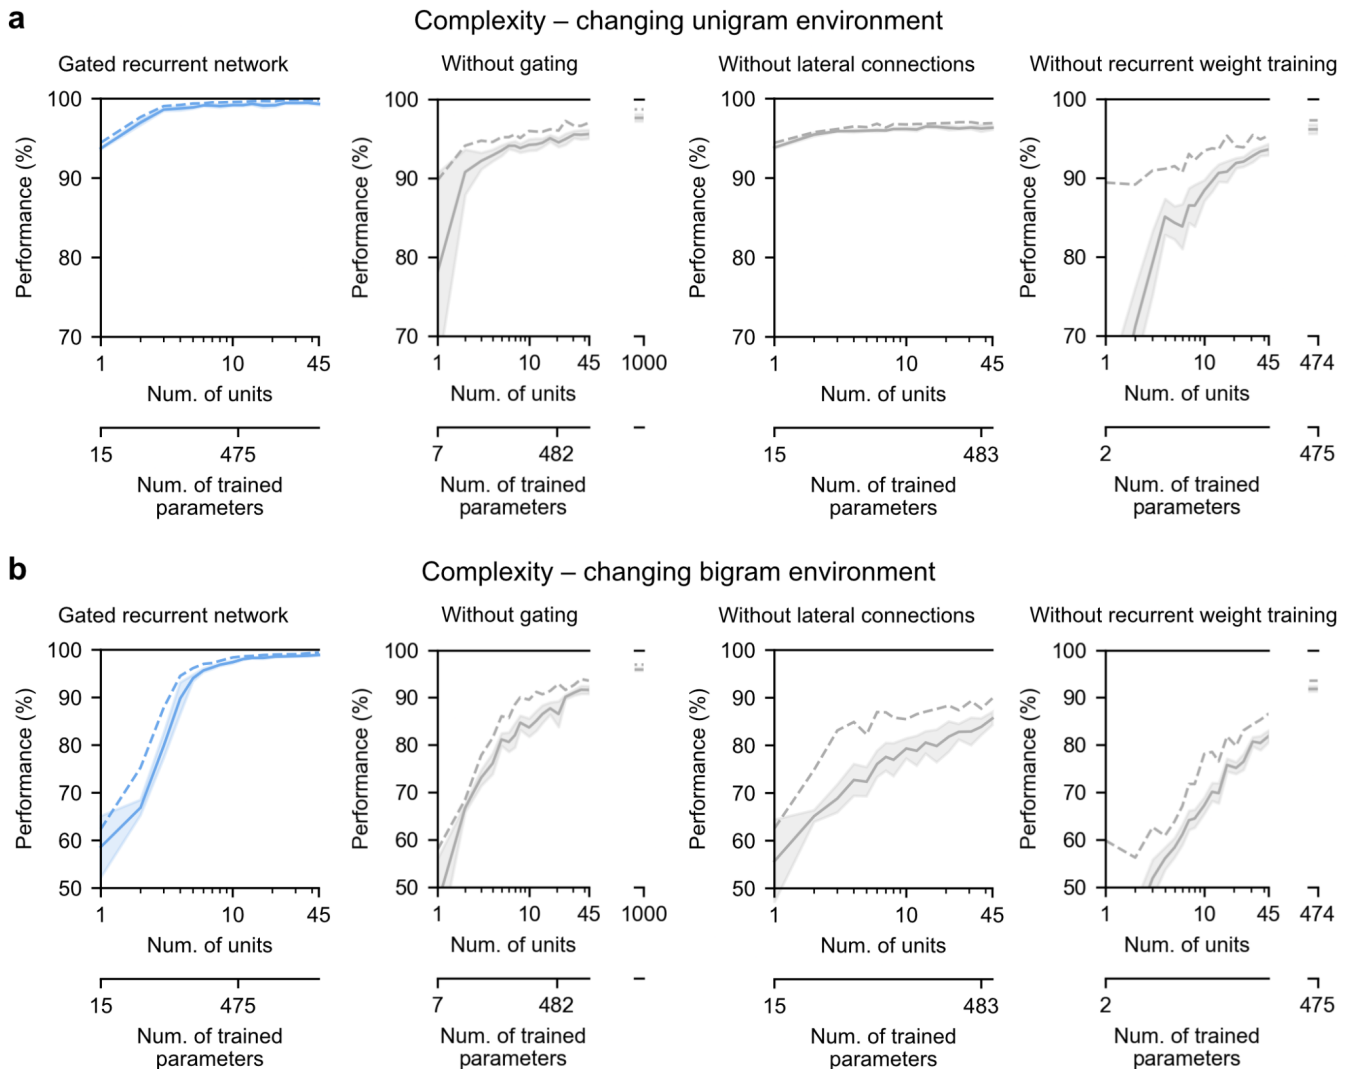
431 We found that gated recurrent networks trained in the coupled environment changed their prediction
432 about the unobserved bigram significantly more than networks trained in the independent environment, and
433 this was true across a large range of streak lengths (**Fig. 7c**, top plot). The mere presence of this effect is
434 particularly impressive given that the coupling makes very little difference in terms of raw performance (**Fig.**
435 **6—figure supplement 1**, the networks trained in either the coupled or the independent environment perform
436 very similarly when tested in either environment). All mechanisms of the gated recurrence are important to
437 achieve this higher-level inference since the networks deprived of either gating, lateral connections, or
438 recurrent weight training did not show any effect, no matter the streak length (**Fig. 7c**, bottom three plots;
439 for every mechanism, there was a significant interaction effect between the removal of the mechanism and
440 the training environment on the change in prediction over networks and streak lengths, all $F(1,6076) > 43.2$,
441 all $p < 0.001$).

442 These results show that gated recurrent networks but not alternative networks leverage the higher level
443 of structure where the change points of the latent probabilities are coupled.

444 **Gated recurrence enables simple solutions**

445 Finally, we highlight the small number of units sufficient to perform quasi-optimally in the increasingly
446 structured environments that we tested: the above-mentioned results were obtained with 11 recurrent units.
447 It turns out that gated recurrent networks can reach a similar performance with even fewer units, especially
448 in simpler environments (**Fig. 8a** and **8b**, left plot). For instance, in the unigram environment, gated recurrent
449 networks reach 99% of their asymptotic performance with no more than 3 units.

450 By contrast, without either gating, lateral connections, or recurrent weight training, even when the
451 networks are provided with more units to match the number of trained parameters in the 11-unit gated
452 recurrent networks, they are unable to achieve similar performance (**Fig. 8a** and **8b**, right three plots, the
453 twin x-axes indicate the number of units and trained parameters).



454

455 **Figure 8. Low-complexity solutions are uniquely enabled by the combination of gating, lateral connections, and recurrent**
 456 **weight training.** (a) and (b) Prediction performance of each network architecture in the changing unigram environment and the
 457 changing bigram environment, respectively, as a function of the number of recurrent units (i.e. space complexity) of the network.
 458 For each network architecture and each number of units, 20 networks were trained using hyperparameters that had been
 459 optimized prior to training, and prediction performance was measured as the % of optimal log likelihood on new test sequences.
 460 Solid lines, bands, and dotted lines show the mean, 95% confidence interval of the mean, and maximum performance,
 461 respectively. At the maximum displayed number of units, all of the alternative architectures have exceeded the complexity of the
 462 11-unit gated recurrent network shown on the left and in previous Figures, both in terms of the number of units and the number of
 463 trained parameters (indicated on the twin x-axes), but none of them have yet reached its performance.

464 **Figure supplement 1. Training speed of the gated recurrent networks in the changing unigram and bigram environments.**

465

466 With an unlimited number of units, at least in the case without gating (i.e. a vanilla RNN, short for recurrent
 467 neural network), the networks will be able to achieve such performance since they are universal
 468 approximators of dynamical systems (Cybenko, 1989; Schäfer & Zimmermann, 2006). However, our results
 469 indicate this could require a very large number of units even in the simplest environment tested here (see
 470 **Fig. 8a** and **8b**, without gating at 1000 units). Indeed, the slow growth of the vanilla RNNs' performance with

471 the number of units is well described by a power law function, of the form: $(100-p) = c(1/N)^{\alpha}$, where p is the
472 % of optimal performance and N is the number of units. We fitted this law in the unigram environment using
473 the obtained performance from 2 to 45 units and it yielded a goodness-of-fit $R^2=92.4\%$ (fitting was done by
474 linear regression on the logarithm of N and $(100-p)$). To further confirm the validity of the power law, we then
475 extrapolated to 1000 units and found that the predicted performance was within 0.2% of the obtained
476 performance for networks of this size (predicted: 97.8%, obtained: 97.6%). Based on this power law, more
477 than 10^4 units would be needed for the vanilla RNN to reach the performance exhibited by the GRU with
478 only 11 units.

479 Note that, in terms of computational complexity, the number of units is a fair measure of space complexity
480 (i.e. the amount of memory) across the architectures we considered, since in all of them it is equal to the
481 number of state variables (having one state variable h_i per unit, see Methods). What varies across
482 architectures is the number of trained parameters, i.e. the degrees of freedom that can be used during
483 training to achieve different dynamics. Still, the conclusion remains the same when an alternative network
484 exceeds the complexity of an 11-unit gated recurrent network in both its number of units and its number of
485 trained parameters.

486 Therefore, it is the specific computational properties provided by the combination of the three
487 mechanisms that afford effective low-complexity solutions.

488 Discussion

489 We have shown that the gated recurrent architecture enables simple and effective solutions: with only 11
490 units, the networks perform quasi-optimally in environments fraught with randomness, changes, and
491 different levels of latent structure. Moreover, these solutions reproduce several aspects of optimality
492 observed in organisms, including the adaptation of their effective learning rate, the ability to represent the
493 precision of their estimation and to use it to weight their updates, and the ability to represent and leverage
494 the latent structure of the environment. By depriving the architecture of one of its mechanisms, we have
495 shown that three of them are important to achieve such solutions: gating, lateral connections, and the
496 training of recurrent weights.

497 Can small neural networks behave like Bayesian agents?

498 A central and much-debated question in the scientific community is whether the brain can perform
499 Bayesian inference (Knill & Pouget, 2004; Bowers & Davis, 2012; Griffiths et al., 2012; Rahnev & Denison,
500 2018; Lee & Mumford, 2003; Rao & Ballard, 1999; Sanborn & Chater, 2016; Chater et al., 2006; Findling et
501 al., 2019; Wyart & Koechlin, 2016; Soltani & Izquierdo, 2019; Findling et al., 2021). From a computational
502 viewpoint, there exists no tractable solution (even approximate) for Bayesian inference in an arbitrary
503 environment, since it is NP-hard (Cooper, 1990; Dagum & Luby, 1993). Being a bounded agent (Simon,

504 1955, 1972), the brain cannot solve Bayesian inference in its most general form. The interesting question is
505 whether the brain can perform Bayesian inference in some environments that occur in real life. More
506 precisely, by “perform Bayesian inference” one usually means that it performs computations that satisfy
507 certain desirable properties of Bayesian inference, such as taking into account a certain type of uncertainty
508 and a certain type of latent structure (Courville et al., 2006; Deroy et al., 2016; Griffiths et al., 2012; Knill &
509 Pouget, 2004; Ma, 2010; Ma & Jazayeri, 2014; Tauber et al., 2017). In this study, we selected specific
510 properties and showed that they can indeed be satisfied when using specific (not all) neural architectures.

511 In the changing unigram and changing bigram environments, our results provide an existence proof: there
512 exist plausible solutions that are almost indistinguishable from Bayesian inference (i.e. the optimal solution).
513 They exhibit qualitative properties of Bayesian inference that have been demonstrated in humans but are
514 lacking in heuristic solutions, such as the dynamic adjustment of the effective learning rate (Behrens et al.,
515 2007; Nassar et al., 2010, 2012), the internal representation of latent variables and the precision of their
516 estimates (Boldt et al., 2019; Meyniel et al., 2015), the precision-weighting of updates (McGuire et al., 2014;
517 Nassar et al., 2010, 2012), and the ability for higher-level inference (Bill et al., 2020; Heilbron & Meyniel,
518 2019; Purcell & Kiani, 2016).

519 The performance we obtained with the gated recurrent architecture is consistent with the numerous other
520 successes it produced in other cognitive neuroscience tasks (J. X. Wang et al., 2018; Yang et al., 2019;
521 Zhang et al., 2020). Our detailed study reveals that it offers quasi-optimal low-complexity solutions to new
522 and difficult challenges, including those posed by bigram and higher-level structures and latent probabilities
523 that change unpredictably anywhere in the unit interval. We acknowledge that further generalization to
524 additional challenges remains to be investigated, including the use of more than two categories of
525 observations or continuous observations, and latent structures with longer range dependencies (beyond
526 bigram probabilities).

527 **Minimal set of mechanisms**

528 What are the essential mechanistic elements that enable such solutions? We show that it suffices to have
529 recurrent units of computation equipped with three mechanisms: 1) input, self, and lateral connections which
530 enable each unit to sum up the input with their own and other units' prior value before a non-linear
531 transformation is applied; 2) gating, which enables multiplicative interactions between activities at the
532 summation step; 3) the training of connection weights.

533 One of the advantages of such mechanisms is their generic character: they do not include any
534 components specifically designed to perform certain probabilistic operations or estimate certain types of
535 latent variables, as often done in neuroscience (Echeveste et al., 2020; Fusi et al., 2007; Jazayeri &
536 Movshon, 2006; Ma et al., 2006; Pecevski et al., 2011; Soltani & Wang, 2010). In addition, they allow
537 adaptive behavior only through recurrent activity dynamics, without involving synaptic plasticity as in other

538 models (Farashahi et al., 2017; Fusi et al., 2005; Iigaya, 2016; Schultz et al., 1997). This distinction has
539 implications for the timescale of adaptation: in the brain, recurrent dynamics and synaptic plasticity often
540 involve short and long timescales, respectively. Our study supports this view: recurrent dynamics allow the
541 networks to quickly adapt to a given change in the environment (**Fig. 3**), while synaptic plasticity allows the
542 training process to tune the speed of this adaptation to the frequency of change of the environment (**Fig.**
543 **3—figure supplement 1**).

544 Our findings suggest that these mechanisms are particularly advantageous to enable solutions with low
545 computational complexity. Without one of them, it seems that a very large number of units (i.e. a large
546 amount of memory) would be needed to achieve comparable performance (**Fig. 8**) (note that universal
547 approximation bounds in vanilla RNNs can be very large in terms of number of units (Barron, 1993;
548 Cybenko, 1989; Schäfer & Zimmermann, 2006)). These mechanisms thus seem to be key computational
549 building blocks to build simple and effective solutions. This efficiency can be formalized as the minimum
550 number of units sufficient for near-optimal performance (as in (Orhan & Ma, 2017) who made a similar
551 argument), and it is important for the brain since the brain has limited computational resources (often
552 quantified by the Shannon capacity, i.e. the number of bits that can be transmitted per unit of time, which
553 here amounts to the number of units) (Bhui et al., 2021; Lieder & Griffiths, 2020). Moreover, simplicity
554 promotes our understanding, and it is with the same goal of understanding that others have used model
555 reduction in large networks (Dubreuil et al., 2020; Jazayeri & Ostojic, 2021; Schaeffer et al., 2020).

556 Since we cannot exhaustively test all possible parameter values, it might be possible that better solutions
557 exist that were not discovered during training. However, to maximize the chances that the best possible
558 performance is achieved after training, we conducted an extensive hyperparameter optimization, repeated
559 for each environment, architecture, and several number of units, until there is no more improvement
560 according to the Bayesian optimization (see Methods).

561 **Biological implementations of the mechanisms**

562 What biological elements could implement the mechanisms of the gated recurrence? Recurrent
563 connections are ubiquitous in the brain (Douglas & Martin, 2007; Hunt & Hayden, 2017); the lesser-known
564 aspect is that of gating. In the next paragraph, we speculate on the possible biological implementations of
565 gating, broadly defined as a mechanism that modulates the effective weight of a connection as a function of
566 the network state (and not limited to the very specific form of gating of the GRU).

567 In neuroscience, many forms of gating have been observed, and they can generally be grouped into three
568 categories according to the neural process that supports them: neural circuits, neural oscillations, and
569 neuromodulation. In neural circuits, a specific pathway can be gated through inhibition/disinhibition by
570 inhibitory (GABAergic) neurons. This has been observed in microscopic circuits, e.g. in pyramidal neurons
571 a dendritic pathway can be gated by interneurons (Costa et al., 2017; Yang et al., 2016), or macroscopic

572 circuits, e.g. in basal ganglia-thalamo-cortical circuits a cortico-cortical pathway can be gated by the basal
573 ganglia and the mediodorsal nucleus of thalamus (O'Reilly, 2006; O'Reilly & Frank, 2006; Rikhye et al.,
574 2018; M. B. Wang & Halassa, 2021; Yamakawa, 2020). In addition to inhibition/disinhibition, an effective
575 gating can also be achieved by a large population of interacting neurons taking advantage of their
576 nonlinearity (Beiran et al., 2021; Dubreuil et al., 2020). Regarding neural oscillations, experiments have
577 shown that activity in certain frequency bands (typically, alpha and beta) can gate behavioral and neuronal
578 responses to the same stimulus (Baumgarten et al., 2016; Busch et al., 2009; Hipp et al., 2011; lemi et al.,
579 2019; Klimesch, 1999; Mathewson et al., 2009). One of the most influential accounts is known as “pulsed
580 inhibition” (Hahn et al., 2019; Jensen & Mazaheri, 2010; Klimesch et al., 2007): a low-frequency signal
581 periodically inhibits a high-frequency signal, effectively silencing the high-frequency signal when the low-
582 frequency signal exceeds a certain threshold. Finally, the binding of certain neuromodulators to the certain
583 receptors of a synapse changes the gain of its input-output transfer function, thus changing its effective
584 weight. This has been demonstrated in neurophysiological studies implicating noradrenaline (Aston-Jones
585 & Cohen, 2005; Salgado et al., 2016; Servan-Schreiber et al., 1990), dopamine (Moyer et al., 2007; Servan-
586 Schreiber et al., 1990; Stalter et al., 2020; Thurley et al., 2008), and acetylcholine (Gil et al., 1997; Herrero
587 et al., 2008) (see review in (Thiele & Bellgrove, 2018)).

588 We claim that gated recurrence provides plausible solutions for the brain because its mechanisms can
589 all be biologically implemented and lead to efficient solutions. However, given their multiple biological
590 realizability, the mapping between artificial units and biological neurons is not straightforward: one unit may
591 map to a large population of neurons (e.g. a brain area), or even to a microscopic, subneuronal component
592 (e.g. the dendritic level).

593 **Training: its role and possible biological counterpart**

594 Regarding the training, our results highlight that it is important to adjust the recurrent weights and thus
595 the network dynamics to the environment (and not fix them as in reservoir computing (Tanaka et al., 2019)),
596 but we make no claims about the biological process that leads to such adjustment in brains. It could occur
597 during development (Sherman et al., 2020), the life span (Lillicrap et al., 2020), or the evolution process
598 (Zador, 2019) (these possibilities are not mutually exclusive). Although our training procedure may not be
599 accurate for biology as a whole, two aspects of it may be informative for future research. First, it relies only
600 on the observation sequence (no supervision or reinforcement), leveraging prediction error signals, which
601 have been found in the brain in many studies (Den Ouden et al., 2012; Eshel et al., 2013; Maheu et al.,
602 2019). Importantly, in predictive coding (Rao & Ballard, 1999), the computation of prediction errors is part of
603 the prediction process; here we are suggesting that it may also be part of the training process (as argued in
604 (O'Reilly et al., 2021)). Second, relatively few iterations of training suffice (**Fig. 8—figure supplement 1**, in
605 the order of 10–100; for comparison, (J. X. Wang et al., 2018) reported training for 40,000 episodes in an
606 environment similar to ours).

607 **Suboptimalities in human behavior**

608 In this study we have focused on some aspects of optimality that humans exhibit in the three
609 environments we explored, but several aspects of their behavior are also suboptimal. In the laboratory, their
610 behavior is often at best qualitatively Bayesian but quantitatively suboptimal. For example, although they
611 adjust their effective learning rate to changes, the base value of their learning rate and their dynamic
612 adjustments may depart from the optimal values (Nassar et al., 2010, 2012; Prat-Carrabin et al., 2021). They
613 may also not update their prediction on every trial, unlike the optimal solution (Gallistel et al., 2014; Khaw et
614 al., 2017). Finally, there is substantial interindividual variability which does not exist in the optimal solution
615 (Khaw et al., 2021; Nassar et al., 2010, 2012; Prat-Carrabin et al., 2021). In the future, these suboptimalities
616 could be explored using our networks by making them suboptimal in three ways (among others): by stopping
617 training before quasi-optimal performance is reached (Caucheteux & King, 2021; Orhan & Ma, 2017), by
618 constraining the size of the network or its weights (with hard constraints or with regularization penalties)
619 (Mastrogiuseppe & Ostojic, 2017; Sussillo et al., 2015), or by altering the network in a certain way, such as
620 pruning some of the units or some of the connections (Blalock et al., 2020; Chechik et al., 1999; LeCun et
621 al., 1990; Srivastava et al., 2014), or introducing random noise into the activity (Findling et al., 2021; Findling
622 & Wyart, 2020; Legenstein & Maass, 2014). In this way, one could perhaps reproduce the quantitative
623 deviations from optimality while preserving the qualitative aspects of optimality observed in the laboratory.

624 **Implications for experimentalists**

625 If already trained gated recurrent networks exist in the brain, then one can be used in a new but similar
626 enough environment without further training. This is an interesting possibility because, in laboratory
627 experiments mirroring our study, humans perform reasonably well with almost no training but explicit task
628 instructions given in natural language, along with a baggage of prior experience (Gallistel et al., 2014;
629 Heilbron & Meyniel, 2019; Khaw et al., 2021; Meyniel et al., 2015; Peterson & Beach, 1967). In favor of the
630 possibility to reuse an existing solution, we found that a gated recurrent network can still perform well in
631 conditions different from those it was trained in: across probabilities of change points (**Fig. 3—figure**
632 **supplement 1**) and latent structures (**Fig. 6—figure supplement 1**, from bigram to unigram).

633 In this study, we adopted a self-supervised training paradigm to see if the networks could in principle
634 discover the latent structure from the sequences of observations alone. However, in laboratory experiments,
635 humans often do not have to discover the structure since they are explicitly told what structure they will face
636 and the experiment starts only after ensuring that they have understood it, which makes the comparison to
637 our networks impossible in this setting in terms of training (see similar argument in (Orhan & Ma, 2017)). In
638 the future, it could be interesting to study the ability of gated recurrent networks to switch from one structure
639 to another after having been informed of the current structure as humans do in these experiments. One
640 possible way would be to give a label that indicates the current structure as additional input to our networks,
641 as in (Yang et al., 2019).

642 One of our findings may be particularly interesting to experimentalists: in a gated recurrent network, the
643 representations of latent probabilities and the precision of these probability estimates (sometimes referred
644 to as confidence (Boldt et al., 2019; Meyniel et al., 2015), estimation uncertainty (McGuire et al., 2014;
645 Payzan-LeNestour et al., 2013), or epistemic uncertainty (Amini et al., 2020; Friston et al., 2015; Pezzulo et
646 al., 2015)) are linearly readable from recurrent activity, the form of decoding most frequently used in
647 neuroscience (Haxby et al., 2014; Kriegeskorte & Diedrichsen, 2019). These representations arise
648 spontaneously, and their emergence seems to come from the computational properties of gated recurrence
649 together with the need to perform well in a stochastic and changing environment. This yields an empirical
650 prediction: if such networks can be found in the brain, then latent probability estimates and their precision
651 should also be decodable in brain signals, as already found in some studies (Bach et al., 2011; McGuire et
652 al., 2014; Meyniel, 2020; Meyniel & Dehaene, 2017; Payzan-LeNestour et al., 2013; Tomov et al., 2020).

653 **Materials and methods**

654 **Sequence prediction problem**

655 The sequence prediction problem to be solved is the following. At each time step, an agent receives as
656 input a binary-valued 'observation', $x_t \in \{0, 1\}$, and gives as output a real-valued 'prediction', $p_t \in [0, 1]$
657 which is an estimate of the probability that the value of the next observation is equal to 1, $p(x_{t+1} = 1)$. Coding
658 the prediction in terms of the observation being 1 rather than 0 is inconsequential since one can be deduced
659 from the other: $p(x_{t+1} = 1) = 1 - p(x_{t+1} = 0)$. The agent's objective is to make predictions that maximize
660 the (log) likelihood of observations in the sequence, which technically corresponds to the negative binary
661 cross-entropy cost function:

$$662 \quad L(p; x) = \sum_{t=0}^{T-1} \log [x_{t+1}p_t + (1 - x_{t+1})(1 - p_t)] \quad (1)$$

663 **Network architectures**

664 All network architectures consist of a binary input unit, which codes for the current observation, one
665 recurrent layer (sometimes called hidden layer) with a number N of recurrent units, and an output unit, which
666 represents the network's prediction. Unless otherwise stated, $N=11$. At every time step, the recurrent unit i
667 receives as input the value of the observation x_t and the previous activation values of the recurrent units j
668 that connect to i $h_{j,t-1}$. It produces as output a new activation value $h_{i,t}$, which is a real number. The output
669 unit receives as input the activations of all of the recurrent units and produces as output the prediction p_t .

670 The parameterized function of the output unit is the same for all network architectures:

$$671 \quad p_t = \sigma \left(\sum_{i=1}^N w_{hp,i} h_{i,t} + b_{hp} \right)$$

672 where σ is the logistic sigmoid, $w_{hp,i}$ is the weight parameter of the connection from the i -th recurrent
 673 unit to the output unit, and b_{hp} is the bias parameter of the output unit.

674 The updating of h_i takes a different form depending on whether gating or lateral connections are included,
 675 as described below.

676 **Gated recurrent network.** A gated recurrent network includes both gating and lateral connections. This
 677 enables multiplicative interactions between the input and recurrent activity as well as the activities of different
 678 recurrent units during the updating of h_i . The variant of gating used here is GRU (Cho et al., 2014; Chung
 679 et al., 2014) For convenience of exposition, we introduce, for each recurrent unit i , two intermediate variables
 680 in the calculation of the update: the reset gate r_i and the update gate z_i , both of which have their own set
 681 of weights and bias. The update gate corresponds to the extent to which a unit can change its values from
 682 one time step to the next, and the reset gate corresponds to the balance between recurrent activity and input
 683 activity in case of update. Note that r_i and z_i do not count as state variables since the system would be
 684 equivalently characterized without them by injecting their expression into the update equation of h_i below.
 685 The update is calculated as follows:

$$r_{i,t+1} = \sigma \left(w_{xr,i}x_{t+1} + b_{xr,i} + w_{hr,ii}h_{i,t} + \sum_{j \neq i} w_{hr,ji}h_{j,t} + b_{hr,i} \right)$$

$$z_{i,t+1} = \sigma \left(w_{xz,i}x_{t+1} + b_{xz,i} + w_{hz,ii}h_{i,t} + \sum_{j \neq i} w_{hz,ji}h_{j,t} + b_{hz,i} \right)$$

$$h_{i,t+1} = z_{i,t+1}h_{i,t} + (1 - z_{i,t+1}) \tanh \left[w_{xh,i}x_{t+1} + b_{xh,i} + r_{i,t+1}(w_{hh,ii}h_{i,t} + \sum_{j \neq i} w_{hh,ji}h_{j,t}) + b_{hh,i} \right]$$

686 $h_{i,t=-1} = 0$

687 where $(w_{xr,i}, b_{xr,i}, w_{hr,ji}, b_{hr,i}), (w_{xz,i}, b_{xz,i}, w_{hz,ji}, b_{hz,i}), (w_{xh,i}, b_{xh,i}, w_{hh,ji}, b_{hh,i})$ are the connection weights
 688 and biases from the input unit and the recurrent units to unit i corresponding to the reset gate, the update
 689 gate, and the ungated new activity, respectively.

690 Another variant of gating is the LSTM (Hochreiter & Schmidhuber, 1997). It incorporates similar gating
 691 mechanisms as that of the GRU and can achieve the same performance in our task. We chose the GRU
 692 because it is simpler than the LSTM and it turned out sufficient.

693 **Without gating.** Removing the gating mechanism from the gated recurrent network is equivalent to
 694 setting the above variables r_i equal to 1 and z_i equal to 0. This simplifies the calculation of the activations
 695 to a single equation, which boils down to a weighted sum of the input and the recurrent units' activity before
 696 applying a non-linearity, as follows:

$$h_{i,t+1} = \tanh \left[w_{xh,i}x_{t+1} + b_{xh,i} + w_{hh,ii}h_{i,t} + \sum_{j \neq i} w_{hh,ji}h_{j,t} + b_{hh,i} \right]$$

697

698 Another possibility (not considered here) would be to set the value of z_i to a constant other than 1 and
699 treat this value (which amounts to a time constant) as a hyperparameter.

700 **Without lateral connections.** Removing lateral connections from the gated recurrent network is
701 equivalent to setting the weights $w_{hr,ji}$, $w_{hz,ji}$, and $w_{hh,ji}$ to 0 for all $j \neq i$. This abolishes the possibility of
702 interaction between recurrent units, which simplifies the calculation of the activations as follows:

$$r_{i,t+1} = \sigma(w_{xr,i}x_{t+1} + b_{xr,i} + w_{hr,ii}h_{i,t} + b_{hr,i})$$

$$z_{i,t+1} = \sigma(w_{xz,i}x_{t+1} + b_{xz,i} + w_{hz,ii}h_{i,t} + b_{hz,i})$$

$$h_{i,t+1} = z_{i,t+1}h_{i,t} + (1 - z_{i,t+1}) \tanh[w_{xh,i}x_{t+1} + b_{xh,i} + r_{i,t+1}w_{hh,ii}h_{i,t} + b_{hh,i}]$$

703

704 Note that this architecture still contains gating. We could have tested a simpler architecture without lateral
705 connection and without gating; however, our point is to demonstrate the specific importance of lateral
706 connections to solve the problem we are interested in with few units, and the result is all the more convincing
707 if the network lacking lateral connections has gating (without gating, it would fail even more dramatically).

708 **Without recurrent weight training.** The networks referred to as “without recurrent weight training” have
709 the same architecture as the gated recurrent networks and differ from them only in the way they are trained.
710 While in the other networks, all of the weights and bias parameters are trained, for those networks, only the
711 weights and bias of the output unit, w_{hp} and b_{hp} , are trained; other weights and biases are fixed to the value
712 drawn at initialization.

713 Environments

714 An environment is characterized by its data generating process, i.e. the stochastic process used to
715 generate a sequence of observations in that environment. Each of the generative processes is described by
716 a graphical model in **Fig. 1—figure supplement 1** and further detailed below.

717 **Changing unigram environment.** In the changing unigram environment, at each time step, one
718 observation is drawn from a Bernoulli distribution whose probability parameter is the latent variable p_t^{env} .
719 The evolution of this latent variable is described by the following stochastic process.

- 720 • Initially, $p_{t=0}^{env}$ is drawn from a uniform distribution on $[0, 1]$.
- 721 • At the next time step, with probability p_c , p_{t+1}^{env} is drawn anew from a uniform distribution on $[0, 1]$ (this
722 event is called a 'change point'), otherwise, p_{t+1}^{env} remains equal to p_t^{env} . The change point probability
723 p_c is fixed in a given environment.

724 **Changing bigram environments.** In the changing bigram environments, at each time step, one
725 observation is drawn from a Bernoulli distribution whose probability parameter is either equal to the latent
726 variable $p_{1|1,t}^{env}$, if the previous observation was equal to 1, or to the latent variable $1 - p_{0|0,t}^{env}$ otherwise (at
727 $t=0$, the previous observation is considered to be equal to 0). The evolution of those latent variables is
728 described by a stochastic process which differs depending on whether the change points are independent
729 or coupled.

- 730 • In both cases, initially, $p_{0|0,t=0}^{env}$ and $p_{1|1,t=0}^{env}$ are both drawn independently from a uniform distribution
731 on $[0,1]$.
- 732 • In the case of **independent change points**, at the next time step, with probability p_c , $p_{0|0,t+1}^{env}$ is
733 drawn anew from a uniform distribution on $[0,1]$, otherwise, $p_{0|0,t+1}^{env}$ remains equal to $p_{0|0,t}^{env}$. Similarly,
734 $p_{1|1,t+1}^{env}$ is either drawn anew with probability p_c or remains equal to $p_{1|1,t}^{env}$ otherwise, and critically,
735 the occurrence of a change point in $p_{1|1}^{env}$ is independent from the occurrence of a change point in
736 $p_{0|0}^{env}$.
- 737 • In the case of **coupled change points**, at the next time step, with probability p_c , $p_{0|0,t+1}^{env}$ and $p_{1|1,t+1}^{env}$
738 are both drawn anew and independently from a uniform distribution on $[0,1]$, otherwise, both remain
739 equal to $p_{0|0,t}^{env}$ and $p_{1|1,t}^{env}$ respectively.

740 The changing bigram environment with independent change points and that with coupled change points
741 constitute two distinct environments. When the type of change points is not explicitly mentioned, the default
742 case is independent change points. For conciseness, we sometimes refer to the changing unigram and
743 changing bigram environments simply as “unigram” and “bigram” environments.

744 In all environments, unless otherwise stated, the length of a sequence is $T = 380$ observations, and the
745 change point probability is $p_c = \frac{1}{75}$, as in previous experiments done with human participants (Heilbron &
746 Meyniel, 2019; Meyniel et al., 2015).

747 **Optimal solution**

748 For a given environment among the three possibilities defined above, the optimal solution to the
749 prediction problem can be determined as detailed in (Heilbron & Meyniel, 2019). This solution consists in
750 inverting the data-generating process of the environment using Bayesian inference, i.e. computing the
751 posterior probability distribution over the values of the latent variables given the history of observation
752 values, and then marginalizing over that distribution to compute the prediction (which is the probability of
753 the next observation given the history of observations). This can be done using a hidden Markov model
754 formulation of the data-generating process where the hidden state includes the values of the latent variables
755 as well as the previous observation in the bigram case, and using the forward algorithm to compute the

756 posterior distribution over the hidden state. Because it would be impossible to compute the probabilities for
 757 the infinitely many possible values of the latent variables in the continuous interval $[0,1]$, we discretized the
 758 interval into 20 equal-width bins for each of the latent variables. For a more exhaustive treatment, see
 759 (Heilbron & Meyniel, 2019) and the online code (<https://github.com/florentmeyniel/TransitionProbModel>).

760 Heuristic solutions

761 The four heuristic solutions used here can be classified into 2×2 groups depending on:

- 762 • which kind of variables are estimated: a unigram probability or two bigram probabilities.
- 763 • which heuristic rule is used in the calculation of the estimates: the delta-rule or the leaky rule.

764 The equations used to calculate the estimates are provided below.

765 Unigram, delta-rule:

$$\widehat{p}_{t+1} = \widehat{p}_t + \alpha (x_{t+1} - \widehat{p}_t)$$

$$766 \widehat{p}_{t=-1} = 0.5$$

767 Unigram, leaky rule:

$$n_{0,t+1} = \alpha n_{0,t} + (1 - x_{t+1})$$

$$n_{1,t+1} = \alpha n_{1,t} + x_{t+1}$$

$$n_{0,t=-1} = n_{1,t=-1} = 0$$

$$768 \widehat{p}_t = \frac{n_{1,t} + 1}{n_{1,t} + n_{0,t} + 2}$$

769 Bigrams, delta-rule:

$$\widehat{p}_{0|0,t+1} = \widehat{p}_{0|0,t} + \alpha (1 - x_t) (1 - x_{t+1} - \widehat{p}_{0|0,t})$$

$$\widehat{p}_{1|1,t+1} = \widehat{p}_{1|1,t} + \alpha x_t (x_{t+1} - \widehat{p}_{1|1,t})$$

$$770 \widehat{p}_{0|0,t=-1} = \widehat{p}_{1|1,t=-1} = 0.5$$

771 Bigrams, leaky rule:

$$n_{0|0,t+1} = \alpha n_{0|0,t} + (1 - x_t) (1 - x_{t+1})$$

$$n_{01,t+1} = \alpha n_{01,t} + (1 - x_t) x_{t+1}$$

$$n_{10,t+1} = \alpha n_{10,t} + x_t (1 - x_{t+1})$$

$$n_{1|1,t+1} = \alpha n_{1|1,t} + x_t x_{t+1}$$

$$n_{0|0,t=-1} = n_{01,t=-1} = n_{10,t=-1} = n_{1|1,t=-1} = 0$$

$$\widehat{p}_{0|0,t} = \frac{n_{0|0,t} + 1}{n_{0|0,t} + n_{01,t} + 2}$$

$$772 \widehat{p}_{1|1,t} = \frac{n_{1|1,t} + 1}{n_{1|1,t} + n_{10,t} + 2}$$

773 The delta-rule corresponds to the update rule of the Rescorla-Wagner model (Rescorla & Wagner, 1972).
774 The leaky rule corresponds to the mean of an approximate posterior which is a Beta distribution whose
775 parameters depend on the leaky counts of observations: $n_1 + 1$ and $n_0 + 1$ (see (Meyniel et al., 2016) for
776 more details).

777 The output prediction value is equal to \hat{p}_t in the unigram case, and in the bigram case, to $\hat{p}_{1|1,t}$ if $x_t = 1$
778 and $1 - \hat{p}_{0|0,t}$ otherwise. The parameter α is a free parameter which is trained (using the same training data
779 as the networks) and thus adjusted to the training environment.

780 Training

781 For a given environment and a given type of agent among the network types and heuristic types, all the
782 reported results are based on 20 agents, each sharing the same set of hyperparameters and initialized with
783 a different random seed. During training, the parameters of a given agent were adjusted to minimize the
784 binary cross-entropy cost function (see equation (1)). During one iteration of training, the gradients of the
785 cost function with respect to the parameters are computed on a subset of the training data (called a
786 minibatch) using backpropagation through time and are used to update the parameters according to the
787 selected training algorithm. The training algorithm was Adam (Kingma & Ba, 2015) for the network types
788 and stochastic gradient descent for the heuristic types.

789 For the unigram environment, the analyses reported in **Fig. 2** to **5** were conducted after training on a
790 common training dataset of 160 minibatches of 20 sequences. For each of the two bigram environments,
791 the analyses reported in **Fig. 6** to **7** were conducted after training on a common training dataset (one per
792 environment) of 400 minibatches of 20 sequences. These sizes were sufficient for the validation
793 performance to converge before the end of training for all types of agents.

794 **Parameters initialization.** For all of the networks, the bias parameters are randomly initialized from a
795 uniform distribution on $[-1/\sqrt{N}, +1/\sqrt{N}]$, and the weights w_{hp} are randomly initialized from a normal
796 distribution with standard deviation $1/\sqrt{N}$ and mean 0. For all of the networks, the weights w_{xr} , w_{xz} , w_{xh}
797 are randomly initialized from a normal distribution with standard deviation $\sigma_{0,x}$ and mean 0, and the weights
798 w_{hr,j_i} , w_{hz,j_i} , w_{hh,j_i} are randomly initialized from a normal distribution with standard deviation $\sigma_{0,h}$ and
799 mean 0 for all $j \neq i$ and $\mu_{0,h,ii}$ for $j = i$. $\sigma_{0,x}$, $\sigma_{0,h}$, $\mu_{0,h,ii}$ are hyperparameters that were optimized for
800 a given environment, type of network, and number of units as detailed in the hyperparameter optimization
801 section (the values resulting from this optimization are listed in **Table 1**).

802 For the initialization of the parameter α in the heuristic solutions, a random value r is drawn from a log-
803 uniform distribution on the interval $[10^{-2.5}, 10^{-0.5}]$, and the initial value of α is set to r in the delta-rule case or
804 $\exp(-r)$ in the leaky rule case.

805 **Hyperparameter optimization**

806 Each type of agent had a specific set of hyperparameters to be optimized. For all network types, it
807 included the initial learning rate of Adam η_0 and the initialization hyperparameters $\sigma_{0,x.}$, $\sigma_{0,h.}$. For the
808 networks without lateral connections specifically, it also included $\mu_{0,h.,ii}$ (for those networks, setting it close
809 to 1 can help avoid the vanishing gradient problem during training (Bengio et al., 1994; Sutskever et al.,
810 2013)); for the other networks, this was set to 0. For the heuristic types, it included only the learning rate of
811 the stochastic gradient descent. A unique set of hyperparameter values was determined for each type of
812 agent, each environment, and, for the network types, each number of units, through the optimization
813 described next.

814 We used Bayesian optimization (Agnihotri & Batra, 2020) with Gaussian processes and the upper
815 confidence bound acquisition function to identify the best hyperparameters for each network architecture,
816 environment, and number of units. During the optimization, combinations of hyperparameter values were
817 iteratively sampled, each evaluated over 10 trials with different random seeds, for a total of 60 iterations
818 (hence, 600 trials) for a given architecture, environment, and number of units. In each trial, one network was
819 created, trained, and its cross-entropy was measured on independent test data. The training and test
820 datasets used for the hyperparameter optimization procedure were not used in any other analyses. The
821 training datasets contained respectively 160 and 400 minibatches of 20 sequences for the unigram and the
822 bigram environment; the test datasets contained 200 sequences for each environment. We selected the
823 combination of hyperparameter values corresponding to the iteration that led to the lowest mean test cross-
824 entropy over the 10 trials. The selected values are listed in **Table 1**.

825 **Table 1. Selected hyperparameter values after optimization. (*: fixed value.)**

environment	network type	N	η_0	σ_x	σ_h	$\mu_{h,ii}$
unigram	gated recurrent network	3	8.00E-02	0.02	0.02	0*
unigram	gated recurrent network	11	6.60E-02	0.43	0.21	0*
unigram	gated recurrent network	45	4.20E-02	1	0.02	0*
unigram	without gating	3	2.50E-02	1	0.07	0*
unigram	without gating	11	1.70E-02	1	0.07	0*
unigram	without gating	45	7.60E-03	1	0.08	0*
unigram	without gating	1000	1.34E-04	1	0.04	0*
unigram	without lateral connections	3	5.30E-02	0.02	0.02	1
unigram	without lateral connections	11	2.70E-02	1	0.02	1
unigram	without lateral connections	45	1.30E-02	1	1	1
unigram	without recurrent weight training	3	1.00E-01	1.07	0.55	0*
unigram	without recurrent weight training	11	1.00E-01	2	0.41	0*
unigram	without recurrent weight training	45	1.00E-01	2	0.26	0*
unigram	without recurrent weight training	474	9.60E-03	1	0.1	0*
bigram	gated recurrent network	3	6.30E-02	0.02	1	0*
bigram	gated recurrent network	11	4.40E-02	1	0.02	0*
bigram	gated recurrent network	45	1.60E-02	1	0.02	0*
bigram	without gating	3	5.50E-02	0.02	0.13	0*
bigram	without gating	11	3.20E-02	1	0.05	0*
bigram	without gating	45	8.90E-03	1	0.06	0*
bigram	without gating	1000	5.97E-05	1	0.03	0*
bigram	without lateral connections	3	4.30E-02	1	0.02	0
bigram	without lateral connections	11	4.30E-02	1	1	0
bigram	without lateral connections	45	2.80E-02	1	1	0
bigram	without recurrent weight training	3	6.60E-02	0.73	0.55	0*
bigram	without recurrent weight training	11	1.00E-01	2	0.45	0*

826

827 For the heuristic types, we used random search from a log uniform distribution in the $[10^{-6}, 10^{-1}]$ range
 828 over 80 trials to determine the optimal learning rate of the stochastic gradient descent. This led to selecting
 829 the value $3 \cdot 10^{-3}$ for all heuristic types and all three environments.

830 Performance analyses

831 All agents were tested in the environment they were trained in (except for **Fig. 6—figure supplement 1**
832 which tests cross-environment performance). We used a single test dataset per environment of 1,000
833 sequences independent of the training dataset. The log likelihood L of a given agent was measured from
834 its predictions according to equation (1). The optimal log likelihood $L_{optimal}$ was measured from the
835 predictions of the optimal solution for the given environment. The chance log likelihood L_{chance} was
836 measured using a constant prediction of 0.5. To facilitate the interpretation of the results, the prediction
837 performance of the agent was expressed as the % of optimal log likelihood, defined as

$$838 \frac{L - L_{chance}}{L_{optimal} - L_{chance}} \times 100 .$$

839 To test the statistical significance of a comparison of performance between two types of agents, we used
840 a two-tailed two independent samples t-test with Welch's correction for unequal variances.

841 Analysis of the effective learning rate

842 The instantaneous effective learning rate of an agent that updates its prediction from p_t to p_{t+1} upon
843 receiving when given as observation x_{t+1} is calculated as:

$$844 \begin{aligned} \alpha_{t+1} &= \frac{p_{t+1} - p_t}{x_{t+1} - p_t} \\ \alpha_{t=0} &= \frac{p_0 - 0.5}{x_0 - 0.5} \quad (2) \end{aligned}$$

845 We call it “effective learning rate” because, had the agent been using a delta-rule algorithm, it would be
846 equivalent to the learning rate of the delta-rule (as can be seen by rearranging the above formula into an
847 update equation), and because it can be measured even if the agent uses another algorithm.

848 Readout analyses

849 The readout of a given quantity from the recurrent units of a network consists of a weighted sum of the
850 activation values of each unit. To determine the weights of the readout for a given network, we ran a multiple
851 linear regression using, as input variables, the activation of each recurrent unit at a given time step, $h_{i,t}$,
852 and as target variable, the desired quantity calculated at the same time step. The regression was run on a
853 training dataset of 900 sequences of 380 observations each (hence, 342,000 samples).

854 In the unigram environment, the precision readout was obtained using as desired quantity the log
855 precision of the posterior distribution over the unigram variable calculated by the optimal solution as
856 previously described, i.e. $\psi_t = -\log \sigma_t$, where σ_t is the standard deviation of the posterior distribution over
857 p_{t+1}^{env} : $\sigma_t = \text{SD}[p_{t+1}^{env} | x_0, \dots, x_t]$ (3).

858 In the bigram environment, the readout of the estimate of a given bigram variable was obtained using as
859 desired quantity the log odds of the mean of the posterior distribution over that bigram variable calculated
860 by the optimal solution, and the readout of the precision of that estimate was obtained using the log precision
861 of that same posterior under the above definition of precision.

862 In **Fig. 4a**, to measure the accuracy of the readout from a given network, we calculated the Pearson
863 correlation between the quantity read from the network and the optimal quantity on a test dataset of 100
864 sequences (hence, 38,000 samples), independent from any training dataset. To measure the Pearson
865 correlation between the read precision and the subsequent effective learning rate, we used 300 out-of-
866 sample sequences (hence, 114,000 samples). To measure the mutual information between the read
867 precision and the prediction of the network, we also used 300 out-of-sample sequences (114,000 samples).

868 In **Fig. 6d**, the log odds and log precision were transformed back into mean and standard deviation for
869 visualization purposes.

870 **Dynamics of network activity in the prediction-precision subspace**

871 In **Fig. 4b**, the network activity (i.e. the population activity of the recurrent units in the network) was
872 projected onto the two-dimensional subspace spanned by the prediction vector and the precision vector.
873 The prediction vector is the vector of the weights from the recurrent units to the output unit of the network,
874 w_{hp} . The precision vector is the vector of the weights of the precision readout described above, $w_{h\psi}$. For
875 the visualization, we orthogonalized the precision vector against the prediction vector using the Gram-
876 Schmidt process (i.e. by subtracting from the precision vector its projection onto the prediction vector), and
877 used the orthogonalized precision vector to define the y-axis shown in **Fig. 4b**.

878 **Perturbation experiment to test precision-weighting**

879 The perturbation experiment reported in **Fig. 5** is designed to test the causal role of the precision read
880 from a given network on its weighting of the next observation, measured through its effective learning rate.
881 We performed this perturbation experiment on each of the 20 networks that were trained within each of the
882 4 architectures we considered. The causal instrument is a perturbation vector q that is added to the network's
883 recurrent unit activations. The perturbation vector was randomly generated subject to the following
884 constraints:

- 885 • $q \cdot w_{h\psi} = \delta\psi$ is the desired change in precision (we used 5 levels) that is read from the units'
886 activities; it is computed by projecting the perturbation onto the weight vector of the precision readout
887 ($w_{h\psi}$, \cdot is the dot product);
- 888 • the perturbation q induces no change in the prediction of the network: $q \cdot w_{hp} = 0$, where w_{hp} is the
889 weight vector of the output unit of the network;

- the perturbation has a constant intensity c across simulations, which we formalize as the norm of the perturbation: $\|q\| = c$.

We describe below the algorithm that we used to generate random perturbations q that satisfy these constraints. The idea is to decompose q into two components: both components leave the prediction unaffected, the first (q_ψ) is used to induce a controlled change in precision, the second (q_r) does not change the precision but is added to ensure a constant intensity of the perturbation across simulations.

- To ensure no change in precision, we compute Q , the subspace of the activation space spanned by all vectors q that are orthogonal to the prediction weight vector w_{hp} , as the null space of w_{hp} (i.e. the orthogonal complement of the subspace spanned by w_{hp} , dimension N-1).
- We compute q_ψ , the vector component of Q that affects precision, as the orthogonal projection of w_{hp} onto Q (q_ψ is thus collinear to the orthogonalized precision axis shown in **Fig. 4b** and described above).
- We compute β_ψ , the coefficient to assign to q_ψ in the perturbation vector to produce the desired change in precision $\delta\psi$, as $\beta_\psi = \frac{\delta\psi}{\|q_\psi \cdot w_{hp}\|}$.
- We compute R , the subspace spanned by all vector components of Q that do not affect precision, as the null space of q_ψ (dimension N-2). A perturbation vector in R therefore leaves both the prediction and the precision unchanged.
- We draw a random unit vector q_r within R (by drawing from all N-2 components).
- We compute β_r , the coefficient to assign to q_r in the perturbation vector so as to ensure that the final perturbation's norm equals c , as $\beta_r = \sqrt{c^2 - \beta_\psi^2 \|q_\psi\|^2}$.
- We combine q_ψ and q_r into the final perturbation vector as $q = \beta_\psi q_\psi + \beta_r q_r$.

The experiment was run on a set of 1,000 sample time points randomly drawn from 300 sequences. First, the unperturbed learning rate was measured by running the network on all of the sequences. Second, for each sample time point, the network was run unperturbed up until that point, a perturbation vector was randomly generated for the desired change of precision and applied to the network at that point, then the perturbed network was run on the next time point and its perturbed learning rate was measured. This was repeated for each level of change in precision. Finally, for a given change in precision, the change in learning rate was calculated as the difference between the perturbed and the unperturbed learning rate.

For statistical analysis, we ran a one-tailed paired t-test to test whether the population's mean change in learning rate was higher at one level of precision change than at the next level of precision change. This was done for each of the four consecutive pairs of levels of change in precision.

921 **Test of higher-level inference about changes**

922 For a given network architecture, higher-level inference about changes was assessed by comparing the
923 population of 20 networks trained in the environment with coupled change points to the population of 20
924 networks trained in the environment with independent change points.

925 In **Fig. 7c**, the change in unobserved bigram prediction for a given streak length m was computed as
926 follows. First, prior sequences were generated and each network was run on each of the sequences. We
927 generated initial sequences of 74 observations each with a probability of 0.2 for the 'observed' bigram (which
928 will render its repetition surprising) and a probability p for the 'unobserved' bigram equal to 0.2 or 0.8 (such
929 probabilities, symmetric and substantially different from the default prior 0.5, should render a change in their
930 inferred value detectable). We crossed all possibilities (0|0 or 1|1 as observed bigram, 0.2 or 0.8 for p) and
931 generated 100 sequences for each (hence 400 sequences total). Second, at the end of each of these initial
932 sequences, the prediction for the unobserved bigram, p_{before} , was queried by retrieving the output of the
933 network after giving it as input '0' if the unobserved bigram was 0|0 or '1' otherwise. Third, the network was
934 further presented with m repeated observations of the same value: '1' if the observed bigram was 1|1 or '0'
935 otherwise. Finally, after this streak of repetition, the new prediction for the unobserved bigram, p_{after} , was
936 queried (as before) and we measured its change with respect to the previous query, $|p_{\text{after}} - p_{\text{before}}|$. This
937 procedure was repeated for m ranging from 2 and 75.

938 For statistics, we ran a one-tailed two independent samples t-test to test whether the mean change in
939 unobserved bigram prediction of the population trained on coupled change points was higher than that of
940 the population trained on independent change points.

941 **Complexity analyses**

942 The complexity analysis reported in **Fig. 8** consisted in measuring, for each network architecture and
943 each environment, the performance of optimally trained networks as a function of the number of units N . For
944 optimal training, hyperparameter optimization was repeated at several values of N , for each type of network
945 and each environment (the resulting values are listed in **Table 1**). For the complexity analysis, a grid of
946 equally spaced N values in logarithmic space between 1 and 45 was generated, and an additional value of
947 474 was included specifically for the networks without recurrent weight training so as to match their number
948 of trained parameters to that of an 11-unit gated recurrent network. For every value on this grid, 20 networks
949 of a given architecture in a given environment were randomly initialized with the set of hyperparameter
950 values that was determined to be optimal for the nearest neighboring N value in logarithmic space. The
951 performance of these networks after training was evaluated using a new couple of training and test datasets
952 per environment, each consisting of 400 minibatches of 20 sequences for training and 1,000 sequences for
953 testing.

954 **Statistics**

955 To assess the variability between different agent solutions, we trained 20 agents for each type of agent
956 and each environment. These agents have different random seeds (which changes their parameter
957 initialization and how their training data is shuffled). Throughout the article, we report mean or median over
958 these agents, and individual data points when possible or 95% confidence intervals (abbreviated as "CI")
959 otherwise, as fully described in the text and figure legends. No statistical methods were used to pre-
960 determine sample sizes but our sample sizes are similar to those reported in previous publications (Masse
961 et al., 2019; Yang et al., 2019). Data analysis was not performed blind to the conditions of the experiments.
962 No data were excluded from the analyses. All statistical tests were two-tailed unless otherwise noted. The
963 data distribution was assumed to be normal, but this was not formally tested. The specific details of each
964 statistical analysis are reported directly in the text.

965 **Code availability**

966 The code to reproduce exhaustively the analyses of this paper is available at
967 https://github.com/cedricfoucault/networks_for_sequence_prediction and archived on Zenodo with DOI:
968 [10.5281/zenodo.5707498](https://doi.org/10.5281/zenodo.5707498). This code also enables to train new networks equipped with any number of units
969 and generate Figures 2 to 7 with those networks.

970 **Data availability**

971 This paper presents no experimental data. All synthetic data are available in the code repository.

972 **Acknowledgements**

973 We thank Yair Lakretz for useful feedback, advice, and discussions throughout the project, Alexandre
974 Pouget for his input when starting this project, and Charles Findling for comments on a previous version of
975 the manuscript.

976 **Author contributions**

977 C.F. and F.M. designed the research. C.F. performed the model simulations and analyses. C.F. and F.M.
978 interpreted the results. C.F. and F.M. wrote the manuscript.

979 **Competing interests**

980 The authors declare no competing interests.

References

- 982 Agnihotri, A., & Batra, N. (2020). Exploring Bayesian Optimization. *Distill*, 5(5), e26. <https://doi.org/10.23915/distill.00026>
- 983 Amini, A., Schwarting, W., Soleimany, A., & Rus, D. (2020). Deep Evidential Regression. *Advances in Neural Information Processing Systems*, 33, 14927–14937.
- 984
- 985 Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: Adaptive gain and optimal performance. *Annual Review of Neuroscience*, 28(1), 403–450. <https://doi.org/10.1146/annurev.neuro.28.061604.135709>
- 986
- 987 Aston-Jones, G., Rajkowski, J., & Kubiak, P. (1997). Conditioned responses of monkey locus coeruleus neurons anticipate acquisition of discriminative behavior in a vigilance task. *Neuroscience*, 80(3), 697–715. [https://doi.org/10.1016/S0306-4522\(97\)00060-2](https://doi.org/10.1016/S0306-4522(97)00060-2)
- 988
- 989 Bach, D. R., Hulme, O., Penny, W. D., & Dolan, R. J. (2011). The Known Unknowns: Neural Representation of Second-Order Uncertainty, and Ambiguity. *Journal of Neuroscience*, 31(13), 4811–4820. <https://doi.org/10.1523/JNEUROSCI.1452-10.2011>
- 990
- 991 Barron, A. R. (1993). Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3), 930–945. <https://doi.org/10.1109/18.256500>
- 992
- 993 Baumgarten, T. J., Schnitzler, A., & Lange, J. (2016). Prestimulus Alpha Power Influences Tactile Temporal Perceptual Discrimination and Confidence in Decisions. *Cerebral Cortex*, 26(3), 891–903. <https://doi.org/10.1093/cercor/bhu247>
- 994
- 995 Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9), 1214–1221. <https://doi.org/10.1038/nn1954>
- 996
- 997 Beiran, M., Dubreuil, A., Valente, A., Mastrogiuseppe, F., & Ostojic, S. (2021). Shaping Dynamics With Multiple Populations in Low-Rank Recurrent Networks. *Neural Computation*, 33(6), 1572–1615. https://doi.org/10.1162/neco_a_01381
- 998
- 999 Bengio, Y., Simard, P., & Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5(2), 157–166.
- 1000
- 1001 Berniker, M., & Kording, K. (2008). Estimating the sources of motor errors for adaptation and generalization. *Nature Neuroscience*, 11(12), 1454–1461. <https://doi.org/10.1038/nn.2229>
- 1002
- 1003 Bhui, R., Lai, L., & Gershman, S. J. (2021). Resource-rational decision making. *Current Opinion in Behavioral Sciences*, 41, 15–21. <https://doi.org/10.1016/j.cobeha.2021.02.015>
- 1004
- 1005 Bill, J., Pailian, H., Gershman, S. J., & Drugowitsch, J. (2020). Hierarchical structure is employed by humans during visual motion perception. *Proceedings of the National Academy of Sciences*, 117(39), 24581–24589. <https://doi.org/10.1073/pnas.2008961117>
- 1006
- 1007 Blalock, D., Ortiz, J. J. G., Frankle, J., & Guttag, J. (2020). What is the State of Neural Network Pruning? *ArXiv:2003.03033 [Cs, Stat]*. <http://arxiv.org/abs/2003.03033>
- 1008
- 1009 Boldt, A., Blundell, C., & De Martino, B. (2019). Confidence modulates exploration and exploitation in value-based learning. *Neuroscience of Consciousness*, 2019(niz004). <https://doi.org/10.1093/nc/niz004>
- 1010
- 1011 Bornstein, A. M., & Daw, N. D. (2013). Cortical and Hippocampal Correlates of Deliberation during Model-Based Decisions for Rewards in Humans. *PLOS Computational Biology*, 9(12), e1003387. <https://doi.org/10.1371/journal.pcbi.1003387>
- 1012
- 1013 Bowers, J. S., & Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, 138(3), 389–414. <https://doi.org/10.1037/a0026450>
- 1014
- 1015 Busch, N. A., Dubois, J., & VanRullen, R. (2009). The Phase of Ongoing EEG Oscillations Predicts Visual Perception. *Journal of Neuroscience*, 29(24), 7869–7876. <https://doi.org/10.1523/JNEUROSCI.0113-09.2009>
- 1016
- 1017 Caucheteux, C., & King, J.-R. (2021). *Language processing in brains and deep neural networks: Computational convergence and its limits* (p. 2020.07.03.186288). <https://doi.org/10.1101/2020.07.03.186288>
- 1018
- 1019 Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences*, 10(7), 287–291. <https://doi.org/10.1016/j.tics.2006.05.007>
- 1020
- 1021 Chechik, G., Meilijson, I., & Ruppin, E. (1999). Neuronal Regulation: A Mechanism for Synaptic Pruning During Brain Maturation. *Neural Computation*, 11(8), 2061–2080. <https://doi.org/10.1162/089976699300016089>
- 1022
- 1023 Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1724–1734. <https://doi.org/10.3115/v1/D14-1179>
- 1024
- 1025 Chung, J., Gulcehre, C., Cho, K., & Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. *NIPS 2014 Workshop on Deep Learning, December 2014*. <https://nyuscholars.nyu.edu/en/publications/empirical-evaluation-of-gated-recurrent-neural-networks-on-sequen>
- 1026
- 1027
- 1028
- 1029 Cooper, G. F. (1990). The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, 42(2–3), 393–405.
- 1030
- 1031 Costa, R., Assael, I. A., Shillingford, B., de Freitas, N., & Vogels, Ti. (2017). Cortical microcircuits as gated-recurrent neural networks. *Advances in Neural Information Processing Systems*, 30. <https://papers.nips.cc/paper/2017/hash/45fbc6d3e05ebd93369ce542e8f2322d-Abstract.html>
- 1032
- 1033
- 1034 Courville, A. C., Daw, N. D., & Touretzky, D. S. (2006). Bayesian theories of conditioning in a changing world. *Trends in Cognitive Sciences*, 10(7), 294–300. <https://doi.org/10.1016/j.tics.2006.05.004>
- 1035

- 1036 Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4), 303–314.
1037 <https://doi.org/10.1007/BF02551274>
- 1038 Dagum, P., & Luby, M. (1993). Approximating probabilistic inference in Bayesian belief networks is NP-hard. *Artificial Intelligence*, 60(1),
1039 141–153. [https://doi.org/10.1016/0004-3702\(93\)90036-B](https://doi.org/10.1016/0004-3702(93)90036-B)
- 1040 De Lange, F. P., Heilbron, M., & Kok, P. (2018). How do expectations shape perception? *Trends in Cognitive Sciences*, 22(9), 764–779.
- 1041 Dehaene, S., Meyniel, F., Wacongne, C., Wang, L., & Pallier, C. (2015). The Neural Representation of Sequences: From Transition
1042 Probabilities to Algebraic Patterns and Linguistic Trees. *Neuron*, 88(1), 2–19. <https://doi.org/10.1016/j.neuron.2015.09.019>
- 1043 Den Ouden, H. E., Kok, P., & De Lange, F. P. (2012). How Prediction Errors Shape Perception, Attention, and Motivation. *Frontiers in*
1044 *Psychology*, 3. <https://doi.org/10.3389/fpsyg.2012.00548>
- 1045 Deroy, O., Spence, C., & Noppeney, U. (2016). Metacognition in Multisensory Perception. *Trends in Cognitive Sciences*, 20(10), 736–747.
1046 <https://doi.org/10.1016/j.tics.2016.08.006>
- 1047 Dolan, R. J., & Dayan, P. (2013). Goals and Habits in the Brain. *Neuron*, 80(2), 312–325. <https://doi.org/10.1016/j.neuron.2013.09.007>
- 1048 Douglas, R. J., & Martin, K. A. C. (2007). Recurrent neuronal circuits in the neocortex. *Current Biology*, 17(13), R496–R500.
1049 <https://doi.org/10.1016/j.cub.2007.04.024>
- 1050 Dubreuil, A., Valente, A., Beiran, M., Mastrogiuseppe, F., & Ostojic, S. (2020). *Complementary roles of dimensionality and population*
1051 *structure in neural computations* (p. 2020.07.03.185942). <https://doi.org/10.1101/2020.07.03.185942>
- 1052 Echeveste, R., Aitchison, L., Hennequin, G., & Lengyel, M. (2020). Cortical-like dynamics in recurrent circuits optimized for sampling-based
1053 probabilistic inference. *Nature Neuroscience*, 23(9), 1138–1149. <https://doi.org/10.1038/s41593-020-0671-1>
- 1054 Elman, J. L. (1990). Finding Structure in Time. *Cognitive Science*, 14(2), 179–211. https://doi.org/10.1207/s15516709cog1402_1
- 1055 Elman, J. L. (1991). Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*, 7(2), 195–225.
1056 <https://doi.org/10.1007/BF00114844>
- 1057 Eshel, N., Tian, J., & Uchida, N. (2013). Opening the black box: Dopamine, predictions, and learning. *Trends in Cognitive Sciences*, 17(9),
1058 430–431. <https://doi.org/10.1016/j.tics.2013.06.010>
- 1059 Fairhall, A. L., Lewen, G. D., Bialek, W., & de Ruyter van Steveninck, R. R. (2001). Efficiency and ambiguity in an adaptive neural code.
1060 *Nature*, 412(6849), 787–792. <https://doi.org/10.1038/35090500>
- 1061 Farashahi, S., Donahue, C. H., Khorsand, P., Seo, H., Lee, D., & Soltani, A. (2017). Metaplasticity as a Neural Substrate for Adaptive Learning
1062 and Choice under Uncertainty. *Neuron*, 94(2), 401–414.e6. <https://doi.org/10.1016/j.neuron.2017.03.044>
- 1063 Findling, C., Chopin, N., & Koehlin, E. (2021). Imprecise neural computations as a source of adaptive behaviour in volatile environments.
1064 *Nature Human Behaviour*, 5(1), 99–112. <https://doi.org/10.1038/s41562-020-00971-z>
- 1065 Findling, C., Skvortsova, V., Dromnelle, R., Palminteri, S., & Wyart, V. (2019). Computational noise in reward-guided learning drives
1066 behavioral variability in volatile environments. *Nature Neuroscience*, 22(12), 2066–2077. <https://doi.org/10.1038/s41593-019-0518-9>
- 1067 Findling, C., & Wyart, V. (2020). *Computation noise promotes cognitive resilience to adverse conditions during decision-making* (p.
1068 2020.06.10.145300). <https://doi.org/10.1101/2020.06.10.145300>
- 1069 Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: From behavior to neural representations.
1070 *Trends in Cognitive Sciences*, 14(3), 119–130. <https://doi.org/10.1016/j.tics.2010.01.003>
- 1071 Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive*
1072 *Neuroscience*, 6(4), 187–214. <https://doi.org/10.1080/17588928.2015.1020053>
- 1073 Fusi, S., Asaad, W. F., Miller, E. K., & Wang, X.-J. (2007). A neural circuit model of flexible sensorimotor mapping: Learning and forgetting
1074 on multiple timescales. *Neuron*, 54(2), 319–333. <https://doi.org/10.1016/j.neuron.2007.03.017>
- 1075 Fusi, S., Drew, P. J., & Abbott, L. F. (2005). Cascade Models of Synaptically Stored Memories. *Neuron*, 45(4), 599–611.
1076 <https://doi.org/10.1016/j.neuron.2005.02.001>
- 1077 Gallistel, C. R., Krishan, M., Liu, Y., Miller, R., & Latham, P. E. (2014). The perception of probability. *Psychological Review*, 121(1), 96–123.
1078 <https://doi.org/10.1037/a0035232>
- 1079 Gijssen, S., Grundei, M., Lange, R. T., Ostwald, D., & Blankenburg, F. (2021). Neural surprise in somatosensory Bayesian learning. *PLOS*
1080 *Computational Biology*, 17(2), e1008068. <https://doi.org/10.1371/journal.pcbi.1008068>
- 1081 Gil, Z., Connors, B. W., & Amitai, Y. (1997). Differential Regulation of Neocortical Synapses by Neuromodulators and Activity. *Neuron*,
1082 19(3), 679–686. [https://doi.org/10.1016/S0896-6273\(00\)80380-3](https://doi.org/10.1016/S0896-6273(00)80380-3)
- 1083 Griffiths, T. L., Chater, N., Norris, D., & Pouget, A. (2012). How the Bayesians got their beliefs (and what those beliefs actually are):
1084 Comment on Bowers and Davis (2012). *Psychological Bulletin*, 138(3), 415–422. <https://doi.org/10.1037/a0026884>
- 1085 Hahn, G., Ponce-Alvarez, A., Deco, G., Aertsen, A., & Kumar, A. (2019). Portraits of communication in neuronal networks. *Nature Reviews*
1086 *Neuroscience*, 20(2), 117–127. <https://doi.org/10.1038/s41583-018-0094-0>
- 1087 Hauser, M. D., Newport, E. L., & Aslin, R. N. (2001). Segmentation of the speech stream in a non-human primate: Statistical learning in
1088 cotton-top tamarins. *Cognition*, 78(3), B53–B64. [https://doi.org/10.1016/S0010-0277\(00\)00132-3](https://doi.org/10.1016/S0010-0277(00)00132-3)
- 1089 Haxby, J. V., Connolly, A. C., & Guntupalli, J. S. (2014). Decoding Neural Representational Spaces Using Multivariate Pattern Analysis.
1090 *Annual Review of Neuroscience*, 37(1), 435–456. <https://doi.org/10.1146/annurev-neuro-062012-170325>
- 1091 Heilbron, M., & Meyniel, F. (2019). Confidence resets reveal hierarchical adaptive learning in humans. *PLOS Computational Biology*, 15(4),
1092 e1006972. <https://doi.org/10.1371/journal.pcbi.1006972>

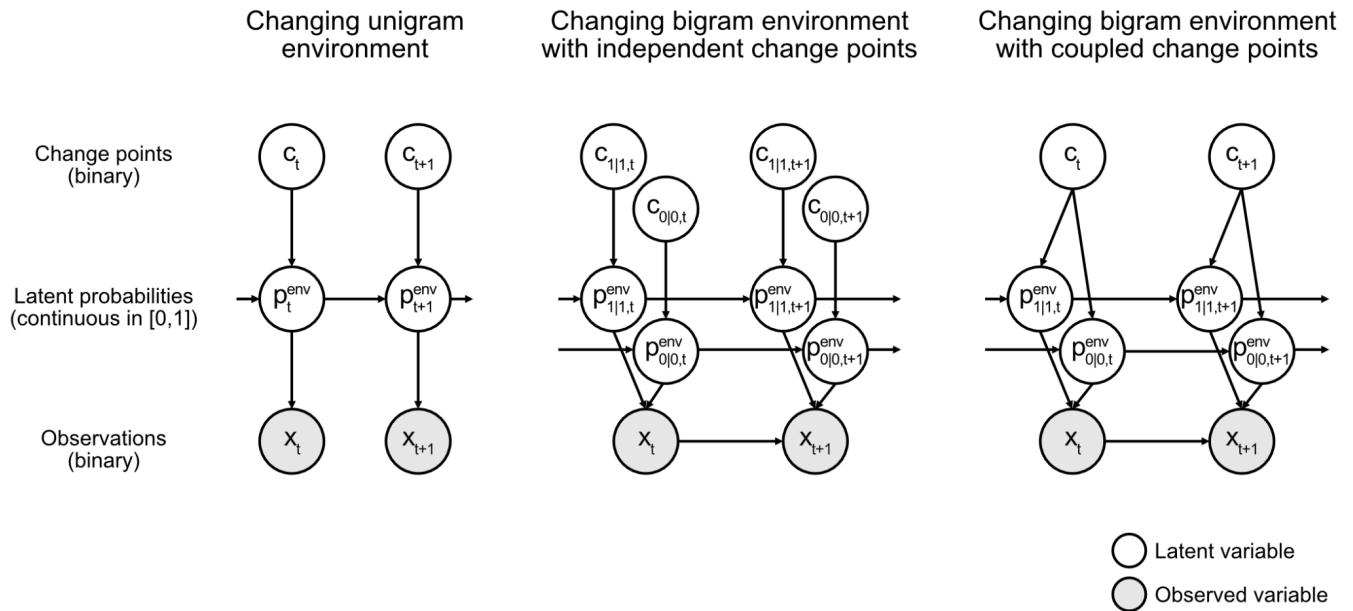
- 1093 Herrero, J. L., Roberts, M. J., Delicato, L. S., Gieselmann, M. A., Dayan, P., & Thiele, A. (2008). Acetylcholine contributes through
1094 muscarinic receptors to attentional modulation in V1. *Nature*, *454*(7208), 1110–1114. <https://doi.org/10.1038/nature07141>
- 1095 Hipp, J. F., Engel, A. K., & Siegel, M. (2011). Oscillatory Synchronization in Large-Scale Cortical Networks Predicts Perception. *Neuron*,
1096 *69*(2), 387–396. <https://doi.org/10.1016/j.neuron.2010.12.027>
- 1097 Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Computation*, *9*(8), 1735–1780.
1098 <https://doi.org/10.1162/neco.1997.9.8.1735>
- 1099 Hunt, L. T., & Hayden, B. Y. (2017). A distributed, hierarchical and recurrent framework for reward-based choice. *Nature Reviews*
1100 *Neuroscience*, *18*(3), 172–182. <https://doi.org/10.1038/nrn.2017.7>
- 1101 Jemi, L., Busch, N. A., Laudini, A., Haegens, S., Samaha, J., Villringer, A., & Nikulin, V. V. (2019). Multiple mechanisms link prestimulus
1102 neural oscillations to sensory responses. *ELife*, *8*, e43620. <https://doi.org/10.7554/eLife.43620>
- 1103 Iglesias, S., Mathys, C., Brodersen, K. H., Kasper, L., Piccirelli, M., den Ouden, H. E., & Stephan, K. E. (2013). Hierarchical prediction errors
1104 in midbrain and basal forebrain during sensory learning. *Neuron*, *80*(2), 519–530.
- 1105 Iigaya, K. (2016). Adaptive learning and decision-making under uncertainty by metaplastic synapses guided by a surprise detection system.
1106 *ELife*, *5*. <https://doi.org/10.7554/eLife.18073>
- 1107 Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience*, *9*(5), 690–
1108 696. <https://doi.org/10.1038/nn1691>
- 1109 Jazayeri, M., & Ostojic, S. (2021). Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity.
1110 *Current Opinion in Neurobiology*, *70*, 113–120. <https://doi.org/10.1016/j.conb.2021.08.002>
- 1111 Jensen, O., & Mazaheri, A. (2010). Shaping Functional Architecture by Oscillatory Alpha Activity: Gating by Inhibition. *Frontiers in Human*
1112 *Neuroscience*, *4*. <https://doi.org/10.3389/fnhum.2010.00186>
- 1113 Kaliukhovich, D. A., & Vogels, R. (2014). Neurons in Macaque Inferior Temporal Cortex Show No Surprise Response to Deviants in Visual
1114 Oddball Sequences. *The Journal of Neuroscience*, *34*(38), 12801–12815. <https://doi.org/10.1523/JNEUROSCI.2154-14.2014>
- 1115 Khaw, M. W., Stevens, L., & Woodford, M. (2017). Discrete adjustment to a changing environment: Experimental evidence. *Journal of*
1116 *Monetary Economics*, *91*, 88–103. <https://doi.org/10.1016/j.jmoneco.2017.09.001>
- 1117 Khaw, M. W., Stevens, L., & Woodford, M. (2021). Individual differences in the perception of probability. *PLOS Computational Biology*,
1118 *17*(4), e1008871. <https://doi.org/10.1371/journal.pcbi.1008871>
- 1119 Kingma, D. P., & Ba, J. (2015). Adam: A Method for Stochastic Optimization. In Y. Bengio & Y. LeCun (Eds.), *3rd International Conference*
1120 *on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
1121 <http://arxiv.org/abs/1412.6980>
- 1122 Klimesch, W. (1999). EEG alpha and theta oscillations reflect cognitive and memory performance: A review and analysis. *Brain Research*
1123 *Reviews*, *29*(2), 169–195. [https://doi.org/10.1016/S0165-0173\(98\)00056-3](https://doi.org/10.1016/S0165-0173(98)00056-3)
- 1124 Klimesch, W., Sauseng, P., & Hanslmayr, S. (2007). EEG alpha oscillations: The inhibition–timing hypothesis. *Brain Research Reviews*, *53*(1),
1125 63–88. <https://doi.org/10.1016/j.brainresrev.2006.06.003>
- 1126 Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*,
1127 *27*(12), 712–719. <https://doi.org/10.1016/j.tins.2004.10.007>
- 1128 Kriegeskorte, N., & Diedrichsen, J. (2019). Peeling the Onion of Brain Representations. *Annual Review of Neuroscience*, *42*(1), 407–432.
1129 <https://doi.org/10.1146/annurev-neuro-080317-061906>
- 1130 LeCun, Y. (2016). Predictive learning. *Proc. Speech NIPS*.
- 1131 LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444. <https://doi.org/10.1038/nature14539>
- 1132 LeCun, Y., Denker, J., & Solla, S. (1990). Optimal Brain Damage. *Advances in Neural Information Processing Systems*, *2*.
1133 <https://proceedings.neurips.cc/paper/1989/hash/6c9882bbac1c7093bd25041881277658-Abstract.html>
- 1134 Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America. A, Optics,*
1135 *Image Science, and Vision*, *20*(7), 1434–1448.
- 1136 Legenstein, R., & Maass, W. (2014). Ensembles of Spiking Neurons with Noise Support Optimal Probabilistic Inference in a Dynamically
1137 Changing Environment. *PLoS Comput Biol*, *10*(10), e1003859. <https://doi.org/10.1371/journal.pcbi.1003859>
- 1138 Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational
1139 resources. *Behavioral and Brain Sciences*, *43*. <https://doi.org/10.1017/S0140525X1900061X>
- 1140 Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J., & Hinton, G. (2020). Backpropagation and the brain. *Nature Reviews Neuroscience*,
1141 *21*(6), 335–346. <https://doi.org/10.1038/s41583-020-0277-3>
- 1142 Ma, W. J. (2010). Signal detection theory, uncertainty, and Poisson-like population codes. *Vision Research*, *50*(22), 2308–2319.
1143 <https://doi.org/10.1016/j.visres.2010.08.035>
- 1144 Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*,
1145 *9*(11), 1432–1438. <https://doi.org/10.1038/nn1790>
- 1146 Ma, W. J., & Jazayeri, M. (2014). Neural coding of uncertainty and probability. *Annual Review of Neuroscience*, *37*, 205–220.
1147 <https://doi.org/10.1146/annurev-neuro-071013-014017>
- 1148 Maheu, M., Dehaene, S., & Meyniel, F. (2019). Brain signatures of a multiscale process of sequence learning in humans. *ELife*, *8*, e41541.
1149 <https://doi.org/10.7554/eLife.41541>

- 1150 Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex.
1151 *Nature*, 503(7474), 78–84. <https://doi.org/10.1038/nature12742>
- 1152 Masse, N. Y., Yang, G. R., Song, H. F., Wang, X.-J., & Freedman, D. J. (2019). Circuit mechanisms for the maintenance and manipulation of
1153 information in working memory. *Nature Neuroscience*, 22(7), 1159–1167. <https://doi.org/10.1038/s41593-019-0414-3>
- 1154 Mastrogiuseppe, F., & Ostojic, S. (2017). Intrinsically-generated fluctuating activity in excitatory-inhibitory networks. *PLOS Computational
1155 Biology*, 13(4), e1005498. <https://doi.org/10.1371/journal.pcbi.1005498>
- 1156 Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., & Ro, T. (2009). To See or Not to See: Prestimulus α Phase Predicts Visual
1157 Awareness. *Journal of Neuroscience*, 29(9), 2725–2732. <https://doi.org/10.1523/JNEUROSCI.3963-08.2009>
- 1158 McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally Dissociable Influences on Learning Rate in a Dynamic
1159 Environment. *Neuron*, 84(4), 870–881. <https://doi.org/10.1016/j.neuron.2014.10.013>
- 1160 Meyniel, F. (2020). Brain dynamics for confidence-weighted learning. *PLOS Computational Biology*, 16(6), e1007935.
1161 <https://doi.org/10.1371/journal.pcbi.1007935>
- 1162 Meyniel, F., & Dehaene, S. (2017). Brain networks for confidence weighting and hierarchical inference during probabilistic learning.
1163 *Proceedings of the National Academy of Sciences*, 114(19), E3859–E3868. <https://doi.org/10.1073/pnas.1615773114>
- 1164 Meyniel, F., Maheu, M., & Dehaene, S. (2016). Human Inferences about Sequences: A Minimal Transition Probability Model. *PLOS
1165 Computational Biology*, 12(12), e1005260. <https://doi.org/10.1371/journal.pcbi.1005260>
- 1166 Meyniel, F., Schlunegger, D., & Dehaene, S. (2015). The Sense of Confidence during Probabilistic Learning: A Normative Account. *PLOS
1167 Computational Biology*, 11(6), e1004305. <https://doi.org/10.1371/journal.pcbi.1004305>
- 1168 Moyer, J. T., Wolf, J. A., & Finkel, L. H. (2007). Effects of Dopaminergic Modulation on the Integrative Properties of the Ventral Striatal
1169 Medium Spiny Neuron. *Journal of Neurophysiology*, 98(6), 3731–3748. <https://doi.org/10.1152/jn.00335.2007>
- 1170 Nassar, M. R., Rumsey, K. M., Wilson, R. C., Parikh, K., Heasly, B., & Gold, J. I. (2012). Rational regulation of learning dynamics by pupil-
1171 linked arousal systems. *Nature Neuroscience*, 15(7), 1040–1046. <https://doi.org/10.1038/nn.3130>
- 1172 Nassar, M. R., Wilson, R. C., Heasly, B., & Gold, J. I. (2010). An Approximately Bayesian Delta-Rule Model Explains the Dynamics of Belief
1173 Updating in a Changing Environment. *The Journal of Neuroscience*, 30(37), 12366–12378.
1174 <https://doi.org/10.1523/JNEUROSCI.0822-10.2010>
- 1175 O'Reilly, R. C. (2006). Biologically Based Computational Models of High-Level Cognition. *Science*, 314(5796), 91–94.
1176 <https://doi.org/10.1126/science.1127242>
- 1177 O'Reilly, R. C., & Frank, M. J. (2006). Making Working Memory Work: A Computational Model of Learning in the Prefrontal Cortex and
1178 Basal Ganglia. *Neural Computation*, 18(2), 283–328. <https://doi.org/10.1162/089976606775093909>
- 1179 O'Reilly, R. C., Russin, J. L., Zolfaghar, M., & Rohrlich, J. (2021). Deep Predictive Learning in Neocortex and Pulvinar. *Journal of Cognitive
1180 Neuroscience*, 1–39. https://doi.org/10.1162/jocn_a_01708
- 1181 Orhan, A. E., & Ma, W. J. (2017). Efficient probabilistic inference in generic neural networks trained with non-probabilistic feedback. *Nature
1182 Communications*, 8(1), 1–14.
- 1183 Payzan-LeNestour, E., Dunne, S., Bossaerts, P., & O'Doherty, J. P. (2013). The Neural Representation of Unexpected Uncertainty during
1184 Value-Based Decision Making. *Neuron*, 79(1), 191–201. <https://doi.org/10.1016/j.neuron.2013.04.037>
- 1185 Pecevski, D., Buesing, L., & Maass, W. (2011). Probabilistic Inference in General Graphical Models through Sampling in Stochastic Networks
1186 of Spiking Neurons. *PLOS Computational Biology*, 7(12), e1002294. <https://doi.org/10.1371/journal.pcbi.1002294>
- 1187 Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychological Bulletin*, 68(1), 29–46. <https://doi.org/10.1037/h0024722>
- 1188 Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active Inference, homeostatic regulation and adaptive behavioural control. *Progress in
1189 Neurobiology*, 134, 17–35. <https://doi.org/10.1016/j.pneurobio.2015.09.001>
- 1190 Prat-Carrabin, A., Wilson, R. C., Cohen, J. D., & da Silveira, R. A. (2021). Human Inference in Changing Environments With Temporal
1191 Structure. *Psychological Review*, 128(5), 879–912. <https://doi.org/10.1037/rev0000276>
- 1192 Purcell, B. A., & Kiani, R. (2016). Hierarchical decision processes that operate over distinct timescales underlie choice and changes in strategy.
1193 *Proceedings of the National Academy of Sciences*, 113(31), E4531–E4540. <https://doi.org/10.1073/pnas.1524685113>
- 1194 Rahnev, D., & Denison, R. N. (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences*, 41.
1195 <https://doi.org/10.1017/S0140525X18000936>
- 1196 Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field
1197 effects. *Nature Neuroscience*, 2(1), 79–87. <https://doi.org/10.1038/4580>
- 1198 Rescorla, Robert A., & Wagner, Allan R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and
1199 nonreinforcement. In A.H., Black, & W.F., W. F. Prokasy, & A. H. Back (Eds.), *Classical conditioning II: Current research and
1200 theory* (pp. 64–99). New York: Appleton-Century Crofts.
- 1201 Rikhye, R. V., Gilra, A., & Halassa, M. M. (2018). Thalamic regulation of switching between cortical representations enables cognitive
1202 flexibility. *Nature Neuroscience*, 21(12), 1753–1763. <https://doi.org/10.1038/s41593-018-0269-z>
- 1203 Robinson, J. G. (1979). An Analysis of the Organization of Vocal Communication in the Titi Monkey *Callicebus moloch*. *Zeitschrift Für
1204 Tierpsychologie*, 49(4), 381–405. <https://doi.org/10.1111/j.1439-0310.1979.tb00300.x>
- 1205 Rose, G. J., Goller, F., Gritton, H. J., Plamondon, S. L., Baugh, A. T., & Cooper, B. G. (2004). Species-typical songs in white-crowned
1206 sparrows tutored with only phrase pairs. *Nature*, 432(7018), 753–758. <https://doi.org/10.1038/nature02992>

- 1207 Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926–1928.
- 1208 Sahani, M., & Dayan, P. (2003). Doubly Distributional Population Codes: Simultaneous Representation of Uncertainty and Multiplicity. *Neural Computation*, 15(10), 2255–2279. <https://doi.org/10.1162/089976603322362356>
- 1209 Salgado, H., Treviño, M., & Atzori, M. (2016). Layer- and area-specific actions of norepinephrine on cortical synaptic transmission. *Brain Research*, 1641, 163–176. <https://doi.org/10.1016/j.brainres.2016.01.033>
- 1211 Sanborn, A. N., & Chater, N. (2016). Bayesian Brains without Probabilities. *Trends in Cognitive Sciences*, 20(12), 883–893. <https://doi.org/10.1016/j.tics.2016.10.003>
- 1212 Saxe, A., Nelli, S., & Summerfield, C. (2021). If deep learning is the answer, what is the question? *Nature Reviews Neuroscience*, 22(1), 55–67. <https://doi.org/10.1038/s41583-020-00395-8>
- 1215 Schaeffer, R., Khona, M., Meshulam, L., Laboratory, I. B., & Fiete, I. R. (2020). *Reverse-engineering Recurrent Neural Network solutions to a hierarchical inference task for mice* (p. 2020.06.09.142745). <https://doi.org/10.1101/2020.06.09.142745>
- 1217 Schäfer, A. M., & Zimmermann, H. G. (2006). Recurrent Neural Networks Are Universal Approximators. In S. D. Kollias, A. Stafylopatis, W. Duch, & E. Oja (Eds.), *Artificial Neural Networks – ICANN 2006* (pp. 632–640). Springer Berlin Heidelberg.
- 1218 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>
- 1220 Schultz, W., Dayan, P., & Montague, P. R. (1997). A Neural Substrate of Prediction and Reward. *Science*, 275(5306), 1593–1599. <https://doi.org/10.1126/science.275.5306.1593>
- 1222 Servan-Schreiber, D., Printz, H., & Cohen, J. D. (1990). A network model of catecholamine effects: Gain, signal-to-noise ratio, and behavior. *Science*, 249(4971), 892–895. <https://doi.org/10.1126/science.2392679>
- 1224 Sherman, B. E., Graves, K. N., & Turk-Browne, N. B. (2020). The prevalence and importance of statistical learning in human cognition and behavior. *Current Opinion in Behavioral Sciences*, 32, 15–20. <https://doi.org/10.1016/j.cobeha.2020.01.015>
- 1226 Simon, H. A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1), 99–118. <https://doi.org/10.2307/1884852>
- 1228 Simon, H. A. (1972). Theories of bounded rationality. *Decision and Organization*, 1(1), 161–176.
- 1230 Sohn, H., Narain, D., Meirhaeghe, N., & Jazayeri, M. (2019). Bayesian Computation through Cortical Latent Dynamics. *Neuron*, 103(5), 934–947.e5. <https://doi.org/10.1016/j.neuron.2019.06.012>
- 1231 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>
- 1233 Soltani, A., & Wang, X.-J. (2010). Synaptic computation underlying probabilistic inference. *Nature Neuroscience*, 13(1), 112–119. <https://doi.org/10.1038/nn.2450>
- 1235 Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1), 1929–1958.
- 1237 Stalter, M., Westendorff, S., & Nieder, A. (2020). Dopamine Gates Visual Signals in Monkey Prefrontal Cortex Neurons. *Cell Reports*, 30(1), 164–172.e4. <https://doi.org/10.1016/j.celrep.2019.11.082>
- 1239 Sterling, P. (2004). Principles of allostasis: Optimal design, predictive regulation, pathophysiology, and rational therapeutics. *Allostasis, Homeostasis, and the Costs of Physiological Adaptation*, 17–64.
- 1242 Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: Neural and computational mechanisms. *Nature Reviews Neuroscience*, 15(11), 745–756. <https://doi.org/10.1038/nrn3838>
- 1243 Sussillo, D., Churchland, M. M., Kaufman, M. T., & Shenoy, K. V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Neuroscience*, 18(7), 1025–1033. <https://doi.org/10.1038/nn.4042>
- 1245 Sutskever, I., Martens, J., Dahl, G., & Hinton, G. (2013). On the importance of initialization and momentum in deep learning. *International Conference on Machine Learning*, 1139–1147. <http://proceedings.mlr.press/v28/sutskever13.html>
- 1247 Sutton, R. (1992). Gain Adaptation Beats Least Squares? In *Proceedings of the 7th Yale Workshop on Adaptive and Learning Systems*, 161–166.
- 1249 Sutton, R. S., & Barto, A. G. (1998). *Introduction to Reinforcement Learning* (1st ed.). MIT Press.
- 1251 Tanaka, G., Yamane, T., Héroux, J. B., Nakane, R., Kanazawa, N., Takeda, S., Numata, H., Nakano, D., & Hirose, A. (2019). Recent advances in physical reservoir computing: A review. *Neural Networks*, 115, 100–123. <https://doi.org/10.1016/j.neunet.2019.03.005>
- 1253 Tauber, S., Navarro, D. J., Perfors, A., & Steyvers, M. (2017). Bayesian models of cognition revisited: Setting optimality aside and letting data drive psychological theory. *Psychological Review*, 124(4), 410–441. <https://doi.org/10.1037/rev0000052>
- 1255 Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science (New York, N.Y.)*, 331(6022), 1279–1285. <https://doi.org/10.1126/science.1192788>
- 1257 Thiele, A., & Bellgrove, M. A. (2018). Neuromodulation of Attention. *Neuron*, 97(4), 769–785. <https://doi.org/10.1016/j.neuron.2018.01.008>
- 1259 Thurley, K., Senn, W., & Lüscher, H.-R. (2008). Dopamine Increases the Gain of the Input-Output Response of Rat Prefrontal Pyramidal Neurons. *Journal of Neurophysiology*, 99(6), 2985–2997. <https://doi.org/10.1152/jn.01098.2007>
- 1260 Todd, P. M., & Gigerenzer, G. (2000). Précis of Simple heuristics that make us smart. *Behavioral and Brain Sciences*, 23(5), 727–741. <https://doi.org/10.1017/S0140525X00003447>
- 1262 Tomov, M. S., Truong, V. Q., Hundia, R. A., & Gershman, S. J. (2020). Dissociable neural correlates of uncertainty underlie different

- 1264 exploration strategies. *Nature Communications*, 11(1), 2371. <https://doi.org/10.1038/s41467-020-15766-z>
- 1265 Ulanovsky, N., Las, L., Farkas, D., & Nelken, I. (2004). Multiple time scales of adaptation in auditory cortex neurons. *The Journal of*
- 1266 *Neuroscience: The Official Journal of the Society for Neuroscience*, 24(46), 10440–10453.
- 1267 <https://doi.org/10.1523/JNEUROSCI.1905-04.2004>
- 1268 Vinckier, F., Gaillard, R., Palminteri, S., Rigoux, L., Salvador, A., Fornito, A., Adapa, R., Krebs, M. O., Pessiglione, M., & Fletcher, P. C.
- 1269 (2016). Confidence and psychosis: A neuro-computational account of contingency learning disruption by NMDA blockade.
- 1270 *Molecular Psychiatry*, 21(7), 946–955. <https://doi.org/10.1038/mp.2015.73>
- 1271 Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., Hassabis, D., & Botvinick, M. (2018). Prefrontal cortex as a
- 1272 meta-reinforcement learning system. *Nature Neuroscience*, 21(6), 860–868. <https://doi.org/10.1038/s41593-018-0147-8>
- 1273 Wang, M. B., & Halassa, M. M. (2021). Thalamocortical contribution to solving credit assignment in neural systems. *ArXiv:2104.01474 [q-*
- 1274 *Bio]*. <http://arxiv.org/abs/2104.01474>
- 1275 Wark, B., Fairhall, A., & Rieke, F. (2009). Timescales of Inference in Visual Adaptation. *Neuron*, 61(5), 750–761.
- 1276 <https://doi.org/10.1016/j.neuron.2009.01.019>
- 1277 Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, 269(5232), 1880–1882.
- 1278 Wyart, V., & Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments. *Current Opinion in Behavioral Sciences*,
- 1279 11, 109–115. <https://doi.org/10.1016/j.cobeha.2016.07.003>
- 1280 Yamakawa, H. (2020). Attentional Reinforcement Learning in the Brain. *New Generation Computing*, 38(1), 49–64.
- 1281 <https://doi.org/10.1007/s00354-019-00081-z>
- 1282 Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T., & Wang, X.-J. (2019). Task representations in neural networks trained to perform
- 1283 many cognitive tasks. *Nature Neuroscience*, 22(2), 297–306. <https://doi.org/10.1038/s41593-018-0310-2>
- 1284 Yang, G. R., Murray, J. D., & Wang, X.-J. (2016). A dendritic disinhibitory circuit mechanism for pathway-specific gating. *Nature*
- 1285 *Communications*, 7(1), 12815. <https://doi.org/10.1038/ncomms12815>
- 1286 Yu, A. J., & Cohen, J. D. (2008). Sequential effects: Superstition or rational behavior? *Advances in Neural Information Processing Systems*,
- 1287 21, 1873–1880.
- 1288 Zador, A. M. (2019). A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications*,
- 1289 10(1), 3770. <https://doi.org/10.1038/s41467-019-11786-6>
- 1290 Zhang, Z., Cheng, H., & Yang, T. (2020). A recurrent neural network framework for flexible and adaptive decision making based on sequence
- 1291 learning. *PLOS Computational Biology*, 16(11), e1008342. <https://doi.org/10.1371/journal.pcbi.1008342>

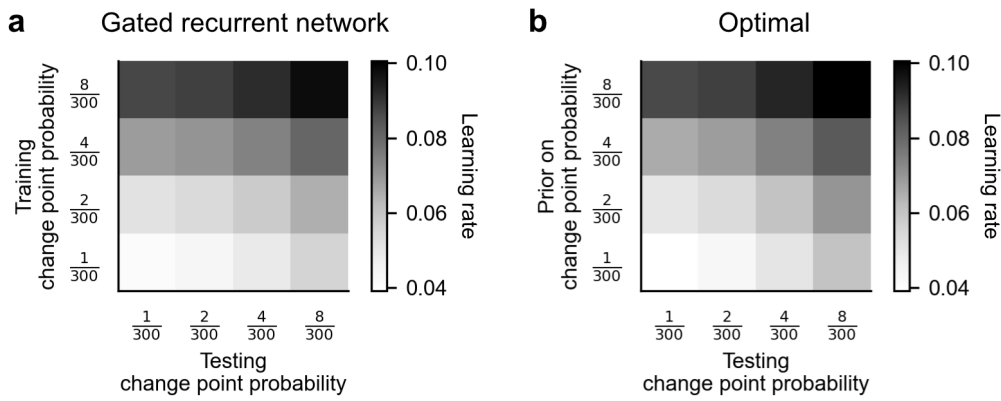
1292 **Figure supplements**



1293

1294 **Figure 1—figure supplement 1. Graphical model of the generative process of each environment.** Nodes encode the
 1295 variables and edges the conditional dependencies between variables. Each graph represents a factorization of the joint probability
 1296 distribution of all variables in the generative process: this joint distribution is the product of the conditional probability distributions
 1297 of each variable given its parents in the graph. For further details on the generative processes, see Methods. In all environments,
 1298 inferring the next observation from previous observations using such a graph is computationally difficult because it requires
 1299 computing and marginalizing over the continuous probability distribution of the latent probabilities. This distribution is not easy to
 1300 compute because it incorporates the likelihoods of the observations (for any latent probability value) and the change point
 1301 probabilities from all previous time steps, and requires normalization. Notice also the increasingly complex conditional structures
 1302 of the graphs from left to right. In the unigram environment, observations are conditionally independent given the latent
 1303 probabilities, but in the bigram environments, they interact. In the bigram environment with coupled change points, the hierarchical
 1304 structure implies that the two latent bigram probabilities are no longer conditionally independent of each other given their values at
 1305 the previous time step, since they are connected by a common parent (the change point).

1306



1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

Figure 3—figure supplement 1. Attunement of the effective learning rate to the change point probabilities. (a) Average effective learning rate of the gated recurrent networks as a function of the change point probability used during testing (columns) and during training (rows). Each row corresponds to a different set of 20 networks trained in the changing unigram environment with the indicated change point probability. Each column corresponds to a different test set with the indicated change point probability, each of 1,000 out-of-sample sequences. The networks' effective learning rate was measured and averaged over time, sequences, and networks. (b) Average effective learning rate of the optimal agent as a function of the change point probability used during testing (columns) and the prior on the change point probability assumed by the model (rows). The optimal agent was tested on the same sets of sequences as the gated recurrent networks and its effective learning rate was averaged over time and sequences.

Performance across training and test environments


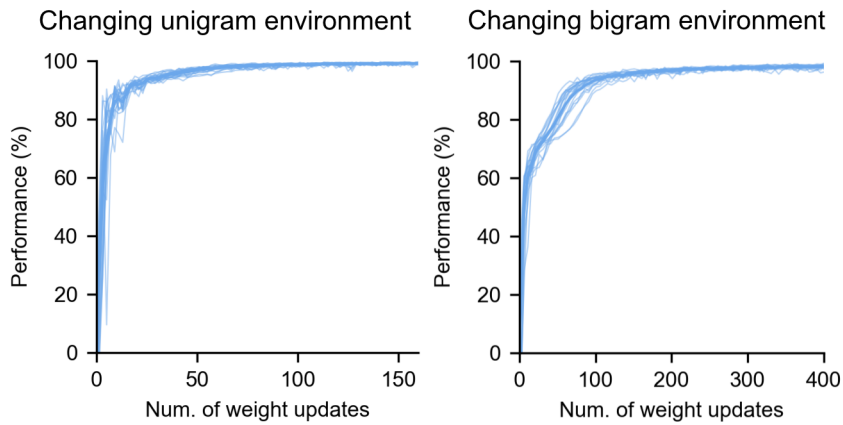


Figure 6—figure supplement 1. Performance across training and test environments. For each type of agent and each environment, a set of 20 agents was trained in the given environment as in Fig. 2, 5, and 6. The performance of each set of trained agents was then evaluated in each test environment, using 1,000 new sequences per environment and the same performance measure as in Fig. 2 and 5. ch.: changing; ind.: independent change points; coup: coupled change points.

Training speed



1322

1323

1324

1325

1326

1327

Figure 8—figure supplement 1. Training speed of the gated recurrent networks in the changing unigram and bigram environments. During training, the networks' weights were iteratively updated, with each update based on the evaluation of the cost function on 20 sequences. Prediction performance was repeatedly measured after each iteration as the % of optimal log likelihood on an out-of-sample validation set of 200 sequences. The thin lines and the thick line show the mean and the individual performances of the 20 networks, respectively.