

Self-Healing Neural Codes

M. E. Rule^{a,1} and T. O’Leary^a

^aUniversity of Cambridge, Engineering Department, Trumpington Street, Cambridge CB2 1PZ, United Kingdom

This manuscript was compiled on March 22, 2021

Neural representations change, even in the absence of overt learning. To preserve stable behavior and memories, the brain must track these changes. Here, we explore homeostatic mechanisms that could allow neural populations to track drift in continuous representations without external error feedback. We build on existing models of Hebbian homeostasis, which have been shown to stabilize representations against synaptic turnover and allow discrete neuronal assemblies to track representational drift. We show that a downstream readout can use its own activity to detect and correct drift, and that such a self-healing code could be implemented by plausible synaptic rules. Population response normalization and recurrent dynamics could stabilize codes further. Our model reproduces aspects of drift observed in experiments, and posits neurally plausible mechanisms for long-term stable readouts from drifting population codes.

Representational Drift | Hebbian Plasticity | Homeostasis

The cellular and molecular components of the brain change over time. In addition to synaptic turnover (1), ongoing reconfiguration of the tuning properties of single neurons has been seen in hippocampus (2, 3) and neocortex, including parietal (4), frontal (5), prefrontal (6), visual (7, 8), and olfactory (9) cortices. Remarkably, the reconfiguration observed in these studies occurs in the absence of any obvious change in behavior, task performance, or perception. How can we reconcile this stability with widespread ongoing changes in how the brain encodes experiences?

These recent and widespread observations seem to be at odds with well established evidence of homeostasis in neural circuit properties. Homeostasis is a feature of all biological systems, and examples of homeostatic plasticity in the nervous system are pervasive (e.g. (10) for review). Broadly speaking, homeostatic plasticity is a negative feedback process that maintains physiological properties such as average firing rates (e.g. 11), neuronal variability (e.g. 12), distributions of synaptic strengths (e.g. 13, 14), and population-level statistics (e.g. (15)). This maintains collective properties, such as the total synaptic drive to a neuron or an average firing rate in a population. Regulation of collective properties is consistent with substantial variability in internal components (16). This suggests that known homeostatic mechanisms may be capable of maintaining a consistent readout from a continually reconfiguring code (17, 18).

In our model of representational drift, homeostatic processes maintain selectivity and function in neural population codes, while allowing individual neurons to reconfigure. We also develop a second sense of homeostasis that allows consolidated representations to maintain stable relationships with unstable neural population codes. This form of homeostasis arises from the interaction between single-cell homeostatic processes, and Hebbian learning in a predictive coding framework. When combined with recurrent network dynamics, such “Hebbian homeostasis” stabilizes consolidated neural representations in the presence of drift.

In this paper, we show that two kinds of homeostatic plasticity can stabilize a population code despite drift. We first argue that single-cell processes can stabilize the information-coding capacity of populations. We then describe a novel form of homeostatic plasticity that allows consolidated representations to interoperate with unstable neural populations. The implication of this finding is that long-term storage of memories and percepts is possible dynamically, with relatively simple, known mechanisms. This potentially reconciles stable behavior with representational drift. The mechanisms we propose here are theoretical, but they are grounded in well-established principles of neuronal function. Our model therefore yields testable predictions about how Hebbian plasticity and homeostasis should interact to stabilize neural representations.

Background. We briefly review representational drift and the broader context of the ideas used in this manuscript. Representational drift refers to seemingly random changes in neural responses during a learned task that are not associated with learning (17). For example, in Driscoll et al. (4) mice navigated to one of two endpoints in a T-shaped maze (Figure 1a), based on a visual cue. Population activity in Posterior Parietal Cortex (PPC) was recorded over several weeks using fluorescence calcium imaging. Neurons in PPC were tuned to the animal’s past, current, and planned behavior. Gradually, the tuning of individual cells changed: neurons that might initially fire at the start of the maze, could start to fire more toward the end—or become disengaged from the task entirely (Figure 1b). The neural population code eventually reconfigured completely (Figure 1c). However, neural tunings continued to tile the task, indicating stable task information at the population level. These features of drift have been observed throughout the brain (3, 7, 8).

Gradual drift would be relatively easy for a downstream readout to track using external error feedback, e.g. from

Significance Statement

The brain reconfigures itself continuously while maintaining stable long-term memories and learned skills. This work examines how stable and unstable neurons can interoperate, despite complete reconfiguration of neural codes, and in the absence of external error signals. We suggest that homeostasis in single neurons can allow the brain to continuously re-interpret shifting neural codes. This could allow the brain to reconfigure how single neurons are used without forgetting by continuously reconsolidating previously learned representations.

TO: Conceptualization, Funding acquisition, Project administration, Resources, Supervision, Writing - review & editing MR: Formal analysis, Investigation, Methodology, Software, Visualization, Writing - original draft, Writing - review & editing

The authors declare no competing interests here.

¹To whom correspondence should be addressed. E-mail: mer49@eng.cam.ac.uk

ongoing rehearsal (18). Indeed, recent simulation studies confirm that learning in the presence of noise can lead to a steady state, in which drift is balanced by error feedback (19, 20). Here, we will show that it is possible to track drift without an external learning objective.

Previous studies have shown how stable functional connectivity can be maintained despite synaptic turnover (21, 22). However, we are interested in the scenario where functional connectivity itself is unstable, allowing the roles of single neurons to change. Additionally, recent work has shown that discrete representations can be stabilized despite drift using neural assemblies (23, 24). Since assembly activation is all-or-nothing, no fidelity is lost if a few neurons enter or leave the assembly. A readout can detect this, and update how it interprets neural population activity (24).

Self-correcting assemblies provide a compelling model for the longevity of discrete information, such as semantic knowledge. However, the brain must contend with continuous sensorimotor variables. Recent experiments suggest that neural representations of these variables are also continuous (25). Even if internal representations are discrete (26, 27), the external world is not. Some states will always lie at ambiguous boundaries between different assemblies. Here, small amounts of drift can introduce large changes.

Despite this, neural representations of continuous tasks are stable. Neural activity is typically confined to a low-dimensional manifold that reflects sensory, motor, and cognitive variables (28). The geometry of these low-dimensional representations is consistent over time, although the way it is reflected in neuronal firing changes (29, 30). Engineers have applied online recalibration and transfer learning and to track drift in brain-machine interface decoders (31). Could neurons in the brain do something similar? We argue that neuronal homeostasis and Hebbian plasticity driven by internally-generated prediction errors allows neural networks to, in effect, "self-heal".

Results

Here, we explore how neural networks could track drift in sensorimotor representations. There are two important general principles to keep in mind throughout. First, distributed neural representations are redundant. To create ambiguity at the macroscopic level, many smaller disruptive changes must occur in a coordinated way. Neurons can exploit this to improve their robustness to drift. Second, learning creates recurrent connections that allow neural populations to predict their own inputs and activity. Even if learning has ceased, these connections continue to constrain activity. This allows a downstream readout to repair inputs corrupted by drift, and use these error-corrected readouts as a training signal.

In the first half of the manuscript, we discuss how homeostasis achieves stable population-level representations, despite instability in single-neuron tunings. We then explore how a single neuron might stabilize its own readout in the presence of drift using homeostasis, and updating its synaptic weights. In the latter half of the manuscript, we show that these rules imply a form of Hebbian learning that achieves homeostasis. We extend these ideas to neural populations, and show that recurrent dynamics can stabilize a readout of an unstable neural code.

A model for representational drift. To understand how neurons cope with unstable population codes, we must first build a model of representational drift. We focus on continuous representations, like those studied in Ziv et al. (2) and Driscoll et al. (4), and simplify our model as much as possible.

Figure 1b illustrates average neuronal fluorescence intensities as a function of progress through the task, mapped to a pseudo-location variable $\theta \in [0, 1]$ (Methods: *Data and analysis*). Neurons fired preferentially in specific parts of the maze. Preferred tunings were typically stable, but occasionally changed abruptly. Figure 1c shows a population of forty neurons tracked over thirty-nine days. Neurons could be sorted according to their preferred location on the first day, and tiled the task space. Preferred tunings gradually switched over time to new locations, leaving little trace of the original code after a month. To model this, we consider a population of N neurons that encodes states θ . We assume that the encoded states θ lie on a continuous low-dimensional manifold. We neglect noise, and assume that θ is encoded in the vector of instantaneous firing rates in a neural population, with tuning curves $\mathbf{x}(\theta) = \{x_1(\theta), \dots, x_N(\theta)\}^\top$.

The population statistics (2, 4, 9), and low-dimensional geometry (29, 30) of drifting population codes remains stable. The properties of single-neuron tuning curves are also preserved: place cells may change their preferred location, but always look like place cells (2). We incorporate these constraints by viewing tuning curves as random samples from the space of possible tuning curves, constrained by the statistics of the encoded variables.

To define this random process, we assume that a task is associated with a set of K features, $\mathbf{s}(\theta) = \{s_1(\theta), \dots, s_K(\theta)\}^\top$. These features have a fixed relationship to the external world, for example visual input or the space of joint configurations, and capture the statistics of the encoded variables θ . To model this, we take $s(\theta)$ to be fixed samples from a Gaussian process on θ :

$$s(\theta) \sim \mathcal{GP}[0, \Sigma(\theta, \theta')] \quad [1]$$

These features are combined linearly through an encoding weight matrix $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_N]$, to yield the synaptic activations $\mathbf{a}(\theta) = \{a_1(\theta), \dots, a_N(\theta)\}^\top$ of the encoding population. Each column \mathbf{u}_i is the encoding weights for a single unit x_i . The firing rates $\mathbf{x}(\theta)$ are then given as a nonlinear function of these activation functions:

$$\begin{aligned} \mathbf{a}(\theta) &= \mathbf{U}^\top \mathbf{s}(\theta) \\ \mathbf{x}(\theta) &= \phi[\mathbf{a}(\theta)] \end{aligned} \quad [2]$$

The nonlinearity $\phi[\cdot]$ can be any function that is rectifying and monotonically increasing; We use the exponential here.

If the encoding weights are taken as i.i.d. samples from a standard normal distribution, $\mathbf{u} \sim \mathcal{N}(0, I_N)$, then the activation functions will follow a zero-mean Gaussian process on θ with covariance inherited from $\mathbf{s}(\theta)$. This converts the problem of defining drift as a random walk through the space of possible activation curves $\mathbf{a}(\theta)$, to a simpler random walk in the space of encoding weights, \mathbf{U} . (See Methods: *Simulated drift* for details of how these weights evolve, and why this preserves information about θ in the population.)

At this point we should pause to address two caveats of this model. First, the fixed features $\mathbf{s}(\theta)$ do not exist in a literal sense. It is true that primary sensory and motor connections are fixed, but these do not provide a sufficiently rich basis to

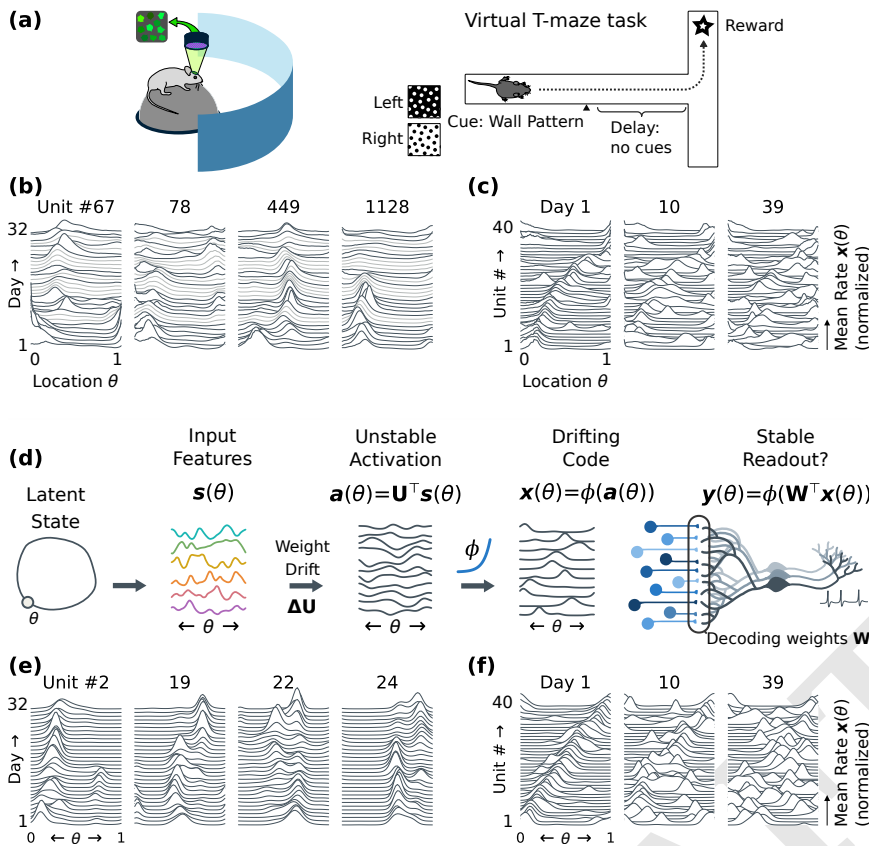


Fig. 1. A model for representational drift. (a) Driscoll et al. (4) imaged population activity in PPC for several weeks, after mice had learned to navigate a virtual T-maze. Neuronal responses continued to change even without overt learning. (b) Tunings were often similar between days, but could change unexpectedly. Plots show average firing rates as a function of task pseudotime (0=beginning, 1=complete) for select cells from (4). Tuning curves from subsequent days are stacked vertically, from day 1 up to day 32. Missing days (light gray) are interpolated. Peaks indicate that a cell fired preferentially at a specific location (Methods: ??). (c) Neuronal tunings tiled the task. Within a day, one can decode the mouse's behavior from population activity (4, 18). Plots show normalized tuning curves for 40 random cells, stacked vertically. Cells are sorted by their preferred location on day 1. By day 10, many cells have changed tuning. Day 39 shows little trace of the original code. (d) We model drift in a simulated rate network (Methods: ??). An encoding population $\mathbf{x}(\theta)$ receives input $\mathbf{s}(\theta)$ with low-dimensional structure, in this case a circular track with location θ . The encoding weights \mathbf{U} driving the activations $\mathbf{a}(\theta)$ of this population drift, leading to unstable tuning. Homeostasis preserves bump-like tuning curves. (e) As in the data (a-c), this model shows stable tuning punctuated by large changes. (f) The neural code reorganizes, while continuing to tile the task. We will examine strategies that a downstream readout could use to update how it decodes $\mathbf{x}(\theta)$ to keep its own representation $\mathbf{y}(\theta)$ stable. This readout is also modeled as linear-nonlinear rate neurons, with decoding weights \mathbf{W} .

190 describe all possible sensorimotor transformations. Richer representations are constructed through transformations within the brain (e.g. 32). The synapses involved in these transformations are also subject to drift. The decomposition of fixed $\mathbf{s}(\theta)$ and drifting $\mathbf{a}(\theta)$ captures the abstract principles that 191 192 193 194 195 196 197 198

199 The second caveat we should address is that this model is not, on its own, especially stable. We have assumed that inputs $\mathbf{s}(\theta)$ and encoding weights \mathbf{U} follow particular distributions, which yield synaptic activations $\mathbf{a}(\theta)$ that produce sensible firing rates when passed through nonlinearity $\phi[\cdot]$. These constraints are easily enforced in a computer, but biological systems must achieve them through homeostatic tuning or regulation of the network activity. 200 201 202 203 204 205 206

207 To model these homeostatic processes, we impose an additional constraint on the mean and the variance of the firing rate for each encoding neuron $x_n(\theta)$: 208 209

$$\begin{aligned} \langle \mathbf{x}_n \rangle &= \mu_0 \\ \text{var}[\mathbf{x}_n] &= \sigma_0^2 \end{aligned} \quad [3]$$

210 211 These moments are fixed by homeostatically adapting a bias β and gain γ of each neuron separately: 212

$$x(\theta) = \phi[\gamma \mathbf{a}(\theta) + \beta]. \quad [4]$$

213 214 The bias can be viewed as threshold adaptation, and the gain as synaptic scaling. These processes control the excitability and variability of the encoding neuron, respectively. They 215 216

217 occur over hours to days, through homeostatic regulation in single neurons (12). For a fixed average firing rate, larger variability invariably corresponds to higher selectivity. Homeostatic regulation of these statistics ensures that (I) encoding neurons retain a reasonable range of firing rates and (II) the tuning curves of these encoding neurons remain selective for a particular preferred stimulus θ_0 (or set of stimuli that are similar in some way). 218 219 220 221 222 223 224


225 For encoding neuron $x(\theta)$, we adjust the gain and bias based on the error between the neuron's firing rate statistics, and the homeostatic targets Eq. (3). 226 227

$$\begin{aligned} \Delta\gamma &\propto \varepsilon_\sigma = (\sigma_0^2 - \text{var}[x])/\sigma_0^2 \\ \Delta\beta &\propto \varepsilon_\mu = \mu_0 - \langle x \rangle \end{aligned} \quad [5]$$

228 229 Multiple homeostatic processes acting in parallel can interact, potentially leading to instability (12). One solution is to allow threshold adaptation to be much faster than synaptic scaling. Another is for the synaptic scaling process to also adapt the threshold, canceling out any influence on excitability. 230 231 232 233

234 Figure 1 shows examples of tuning curve drift from Driscoll et al (4), compared to the Gaussian-process model of drift described above. Figure 1d-f illustrates simulated tuning curve drift in the model. We define a circular environment with location $\theta \in [0, 2\pi)$. This location drives fixed input features $\mathbf{s}(\theta)$, which then drive activity in the encoding population $\mathbf{x}(\theta)$ via encoding weights \mathbf{U} . Drift is simulated as a random walk on these encoding weights, and the encoding cells' tuning curves are homeostatically maintained according to Eq. (3) and Eq. (4) (Methods: *Simulated drift*). Notably, the model mimics changes in tuning curves seen *in vivo*. In Figure 1e, we see that individual encoding neurons show a punctuated 235 236 237 238 239 240 241 242 243 244 245

(a) Fixed Weights in the Presence of Drift



Readout Tuning $y(\theta)$






Fig. 2. Homeostatic Hebbian plasticity enables stable readout from unstable populations. (a) Simulated linear-nonlinear units that are driven by redundant population activity show a loss of excitability, not a change in tuning, when their inputs drift. Since the cell is selective to a conjunction of features, it loses excitatory drive when some of its inputs change. Since most drift is orthogonal to this readout, however, the preferred tuning θ_0 does not change. The right-most plot shows that the excitability gradually diminishes as a larger fraction of inputs change. (b) Homeostatic adjustments to neuron sensitivity stabilizes readouts for small amounts of drift. As more inputs reconfigure, the cell compensates for loss of excitatory drive by increasing an effective gain parameter γ . However, the readout changes to a new, random location once a substantial fraction of inputs have reconfigured (right). This phenomenon is the same as the model for tuning curve drift in the encoding population (c.f. Fig. 1e). (c) Hebbian homeostasis increases neuronal variability by potentiating synaptic inputs that are correlated with post-synaptic activity, or depressing those same synapses when neuronal variability is too high. This results in the neuron re-learning how to decode its own tuning curve from the shifting population code, supporting a stable readout despite complete reconfiguration (right).

(b) Sensitivity Homeostasis



(c) Hebbian Homeostasis



246 stability in their tuning, similar to Figure 1b. Likewise, Figure
 247 1f shows that the tuning curves of the encoding population tile
 248 the state space, but gradually reconfigure over several weeks.
 249 Overall, this illustrates that neural population codes displaying
 250 drift similar to that seen in the brain arise under very generic
 251 circumstances. The only constraints are (I) that inputs to the
 252 population reflect the similarity space of the encoded variables
 253 θ , and (II) that neuronal excitability and selectivity are
 254 homeostatically maintained.

Hebbian homeostasis stabilizes readouts without error feedback.

255 Neural population codes are massively redundant. For
 256 example, most of the neural variability in (4) is explained by
 257 progress through the maze, conditioned on the current and
 258 planned turn direction. Nonlinear dimensionality reduction
 259 algorithms recover the latent T-shaped structure of the task
 260 (17). Because of redundancy, there are many valid ways to
 261 decode information from the population. We propose that,
 262 in the absence of external error feedback or sensorimotor
 263 rehearsal, a readout could use this to generate a surrogate error
 264 signal. The error signal supports a plasticity rule that could
 265 allow unstable neural codes to be continuously reconsolidated.

266 This self training re-encodes a learned readout function $y(\theta)$
 267 in terms of the new neural code $x(\theta)$, allowing the network
 268 to track an unstable representation. Surprisingly, this "self-
 269 healing" plasticity stabilizes the readout of unstable population
 270 codes even in single neurons. We first sketch an example of
 271 this plasticity, and then explore why this works.

272 Using our drifting population code as input, we model
 273 a readout population of M neurons with tuning curves
 274 $y(\theta) = \{y_1(\theta), \dots, y_M(\theta)\}^T$ (Figure 1d). If this readout is
 275 stable, then the responses $y(\theta)$ should remain fixed, even as the
 276 encoding population $x(\theta)$ reconfigures completely. We model
 277 this decoder as a linear-nonlinear function, using decoding
 278 weights \mathbf{W} and biases (thresholds) \mathbf{b} :

$$280 \quad y(\theta) = \phi[\mathbf{W}^T \mathbf{x}(\theta) + \mathbf{b}]. \quad [6]$$

281 On each simulated day, we re-train the decoding weights using

a Hebbian rule. This potentiates decoding weights whose input
 $x_n(\theta)$ correlates with the post-synaptic firing rate $y_m(\theta)$. We
 also adapt the threshold \mathbf{b} to maintain the average firing rate,
 and include some weight decay:

$$282 \quad \Delta \mathbf{W} \propto \varepsilon_\sigma [\langle \mathbf{x}(\theta) \mathbf{y}(\theta)^T \rangle_\theta - \mathbf{W}] \quad 283$$

$$284 \quad \Delta \mathbf{b} \propto \varepsilon_\mu (\langle \mathbf{x}(\theta) \rangle_\theta - \mathbf{b}). \quad 285$$

286 In some ways, Eq. (7) resembles the homeostatic rules ex-
 287 plored earlier (Eq. (3)). Firing rate statistics are controlled
 288 through negative feedback, driven by measurements of the
 289 deviations from the target set-points ε_μ and ε_σ . However,
 290 rather than scale all weights uniformly, this rule adjusts the
 291 component of the weights that is most correlated with the post-
 292 synaptic output, $y(\theta)$. Traditionally, "homeostatic Hebbian
 293 plasticity" refers to processes that stabilize synaptic weights
 294 and responses under ongoing rehearsal and learning. The
 295 role of "Hebbian homeostasis" here is more specific: the neu-
 296 rons adjust their activity toward homeostatic set-points using
 297 Hebbian (or anti-Hebbian) learning. 298

299 Figure 2 simulates a single neuron driven by the unstable
 300 population code. With fixed weights (Figure 2a), drift reduces
 301 the excitability without changing its tuning. This is because
 302 the readout requires a conjunction of specific inputs to fire.
 303 Drift gradually destroys this conjunction, and is unlikely to
 304 spontaneously create a similar conjunction at a different part
 305 of the coding space. A similar phenomena may underlie forms
 306 of drift that consist of changes in excitability, but stable
 307 preferred tuning (5, 7, 33). For small amounts of drift, firing-
 308 rate homeostasis Eq. (5) can temporarily stabilize the readout
 309 (Figure 2b). Eventually, however, the encoding population
 310 reconfigures so drastically that no trace of the original code
 311 remains, and the cell acquires a new preferred stimulus.

312 In contrast, Figure 2c illustrates the consequences of Heb-
 313 bian homeostasis. As the encoding population $x(\theta)$ drifts, the
 314 excitatory drive to the neuron decreases. This activates home-
 315 ostatic plasticity to restore neuronal excitability. However,
 316 instead of scaling up all synapses uniformly, the neuron selec-

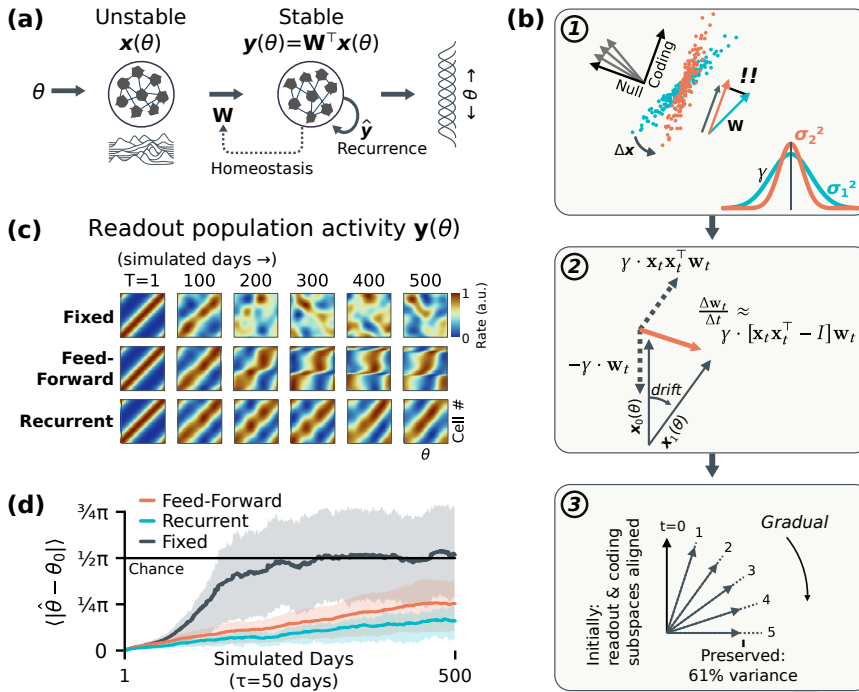


Fig. 3. Self-healing codes in a linear model. (a) Network schema: An unstable population $\mathbf{x}(\theta)$ encodes variables θ (c.f. Fig. 1). A linear readout $\mathbf{y}(\theta)$ seeks to homeostatically preserve its representation, and can use recurrent activity as a training signal $\hat{\mathbf{y}}$. (b-1) Drift changes the low-dimensional structure of population activity. Most drift occurs in non-coding directions, and readouts can detect when low-dimensional activity no longer aligns with their synaptic weights. In linear models, this corresponds to reduced firing-rate variability. (b-2) Hebbian homeostasis restores a target variability by re-aligning the decoding weights with low-dimensional activity. This is the sum of a Hebbian and weight-decay term, scaled by the homeostatic error γ . (b-3) For small amounts of drift, this self-repair has low (but nonzero) error. Large amounts of drift can be tracked if changes are gradual. (c) Readout stability $\mathbf{y}(\theta)$ with bump-like tuning curves tiling a circular space. Encoding cells $\mathbf{x}(\theta)$ drift with time-constant $\tau = 50$ days ("one epoch"). We simulate ten epochs, applying continuous-time Hebbian homeostatic learning rules (Eq. 13). Fixed weights degrade rapidly. Single-cell homeostasis provides some stability for ≈ 3 epochs, but preferred directions shift. Recurrent dynamics better preserve population correlation structure. (d) Hebbian homeostasis reduces the drift of the readout, and recurrence stabilizes it further. The ability of the linear network to error-correct is limited, so the readout still drifts in the long-term (but see Fig. 4). Shaded regions reflect the interquartile range over twenty realizations.

317 tively potentiates the component of $\mathbf{x}(\theta)$ that correlates with
 318 its own output. This leverages the fact that small amounts
 319 of drift change neuronal excitability, but not tuning. The
 320 neuron's own output provides a teaching signal to re-learn
 321 decoding weights for inputs that have changed.

322 If Hebbian homeostasis is applied continuously, a readout
 323 can track drift despite complete reconfiguration in the encoding
 324 population $\mathbf{x}(\theta)$. In effect, the readout's initial tuning curve
 325 is transported to a new set of weights that estimate the same
 326 function from an entirely different input (Methods: *Weight*
 327 *filtering*). This homeostatic rule might seem ad-hoc. However,
 328 we will show that such a rule arises naturally as a plausible
 329 consequence of the interaction between prevailing models of
 330 learning and homeostasis.

331 **Internal models track drift.** Since most neurons are not coupled
 332 directly to the external world, learning must incorporate con-
 333 straints on perception and behavior into local networks (34).
 334 Neural populations learn internal models that recapitulate
 335 and predict the statistics of the external world (35–39). We
 336 propose that these internal models provide the error signals
 337 needed to integrate stable and volatile neural representations.
 338 In essence, the brain generates a teaching signal that trains
 339 neurons how to re-interpret the meaning of neurons whose
 340 function have changed. By computing this teaching from re-
 341 current dynamics, the brain continually re-trains itself. This
 342 implies that a strategy for tracking drift in a neural population
 343 should contain three components.

- 344 I The readout should leverage redundancy to minimize the
 345 error caused by drift.
- 346 II The readout should use its own activity as a training
 347 signal to update its decoding weights.
- 348 III The correlation structure of the readout population should
 349 be homeostatically preserved.

To show how these principles imply Hebbian homeostasis, we
 350 unpack them in a linear network. We then illustrate that
 351 these principles lead to long-term stability, despite drift, in a
 352 nonlinear network. 353

A self-healing linear readout. In a linear network (Fig. 3a),
 354 the readouts $\mathbf{y}(\theta)$ can be viewed as the output of ordinary
 355 least-squares linear regression. Although this network is not
 356 particularly good at correcting errors, it does provide useful
 357 intuition. We incorporate the three components of self-healing
 358 codes (robustness, self-training, and correlation homeostasis)
 359 as follows: (I) We regularize decoding weights to improve
 360 robustness; (II) We use the readout's own activity as a
 361 training signal; (III) We use homeostasis to stabilise firing-
 362 rate variability, and recurrent dynamics to stabilize correlations.
 363

We assume that the readout is initially trained from an
 364 external error signal, and consider a drifting population code
 365 $\mathbf{x}_d(\theta)$ that evolves randomly over several days ' d '. Given a
 366 training signal \mathbf{y}_0 , the regularized least-squares solution for
 367 the ideal decoding weights for the following day $d + 1$ is:
 368

$$\mathbf{W}_{d+1} = [\Sigma_d + \Sigma_\Delta]^{-1} \langle \mathbf{x}_d \mathbf{y}_0^\top \rangle, \quad [8] \quad 369$$

where $\Sigma_d = \langle \mathbf{x}_d \mathbf{x}_d^\top \rangle$ is the covariance of the encoding popu-
 370 lation on day d , and Σ_Δ is a regularizing term reflecting the
 371 expected covariance of day-to-day drift. 372

To incorporate self-training, we generate the training signal
 373 for the weights on day $d + 1$ from the network's own output on
 374 day d . For a linear readout, the readout is the linear projection
 375 $\hat{\mathbf{y}} = \mathbf{W}_d^\top \mathbf{x}_d$. The expectation $\langle \mathbf{x}_d \mathbf{y}_0^\top \rangle$ therefore equals $\Sigma_d \mathbf{W}_d$
 376 and on may write:
 377

$$\mathbf{W}_{d+1} = [\Sigma_d + \Sigma_\Delta]^{-1} \Sigma_d \mathbf{W}_d \quad [9] \quad 378$$

This update applies recursive filtering to the weights (Methods:
 379 *Weight filtering*). However, filtering alone is unhelpful (Fig.
 380 3e), since it allows activity to decay as predictions become
 381

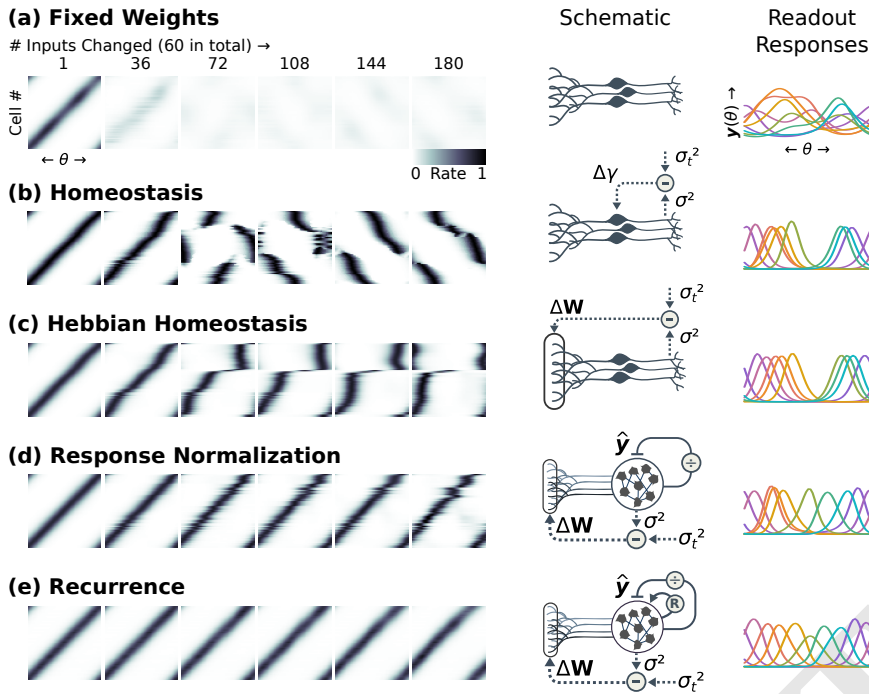


Fig. 4. Self-healing readout in a nonlinear rate network. Each plot shows (left) the stability of a population readout $\mathbf{y}(\theta)$ from a drifting code $\mathbf{x}(\theta)$ over time, (middle) a schematic of the readout dynamics, and (right) a plot of select readout unit's tuning to θ if 55 out of 60 (92%) of the encoding cells were to abruptly switch to a new, random tuning. **(a)** For fixed readouts, representational drift in the encoding population gradually destroys the feature conjunctions used to define selective activity in the readout. **(b)** Homeostatic processes could stabilize the mean firing rate and variability in readout cells. For small amounts of drift, homeostasis can compensate for loss of drive. However, drift eventually disrupts the readout's tuning curve. **(c)** Hebbian homeostasis can preserve the statistics of tuning curves in single cells, by using a neuron's own output as a training signal to update decoding weights. However, this process is not lossless, and the population code in the readout degrades over time. **(d)** Response normalization controls the population firing rate, causing neurons to compete for activation. This stabilizes the statistics of the population code, but readout neurons can still swap preferred tunings, degrading the readout. **(e)** Recurrent activity, in which the network predicts its own activity, can enforce population correlations. This limits the structure of the readout to the ring-like encoding in which it was first trained. Here, the only drift that is permitted is along the symmetry of the circular state θ .

uncertain. To stabilize the firing-rate variability, we rescale the training signal to compensate for any loss of variability σ_y^2 , from its homeostatic target σ_0^2 . For a single readout neuron with weights \mathbf{w} , this gives the homeostatic update:

$$\mathbf{w}_{d+1} = [\Sigma_d + \Sigma_\Delta]^{-1} \Sigma_d \mathbf{w}_d \frac{\sigma_0}{\sigma_y} \quad [10]$$

This update can be solved by online stochastic gradient descent using a Hebbian learning rule (Methods: *Synaptic learning rules*).

$$\Delta \mathbf{w}_t \propto \gamma \mathbf{x}_t \mathbf{y}_t^\top - \Sigma_\Delta \mathbf{w}_t \quad [11]$$

We can use loss of excitatory drive as an indicator of the current drift rate, setting $\Sigma_\Delta \approx \gamma I$ (Methods: *Estimating the rate of drift*). This gives a Hebbian rule:

$$\Delta \mathbf{w}_t \propto \gamma \cdot [\mathbf{x}_t \mathbf{x}_t^\top - I] \mathbf{w}_t \quad [12]$$

This learning rule is the same as the Hebbian homeostasis rule proposed earlier Eq. (7). It acts as follows: In redundant, low-dimensional codes, most drift occurs in directions that are not used for coding (Fig. 3b-1). Drift does, however, reduce input drive to a readout. Neurons can detect this, and apply Hebbian homeostasis to re-align their decoding weights with the encoding subspace (Fig. 3b-2). This process allows synaptic weight to track drift as it occurs. If drift is gradual, a stable readout can survive multiple complete reconfigurations of the input code (Fig. 3b-3). This update resembles classic linear approximations to Hebbian learning (40) with weight decay. Such learning rules extract the leading principle component(s) of their input. This can cause different cells to regress to encoding the same salient inputs. Population interactions can counter this, as we will explore later.

Recurrence in a linear model Hebbian homeostasis improves stability, but does not stabilize the population code in the long-term, since the tuning of each neuron can diffuse slowly. Recurrent dynamics address this by deleting changes in $\mathbf{y}(\theta)$ that are inconsistent with the learned structure of θ . We

define recurrent weights \mathbf{R} that transform the feed-forward activations $\mathbf{y}_f = \mathbf{W}^\top \mathbf{x}$ into an error-corrected training signal $\mathbf{y}_r = \mathbf{R}^\top \mathbf{y}_f$. This gives a new Hebbian learning term that cancels the difference between feed-forward and recurrent activity:

$$\Delta \mathbf{W} \propto \gamma [\langle \mathbf{x} \mathbf{y}_r^\top \rangle - \mathbf{w}_t] + \rho \langle \mathbf{x} (\mathbf{y}_r - \mathbf{y}_f)^\top \rangle, \quad [13]$$

where ρ sets the influence of recurrent dynamics on the decoding weights. The error signal $\mathbf{y}_r - \mathbf{y}_f$ can be computed using recurrent negative feedback in a predictive coding framework (Methods: *Linear network with recurrence*). The benefits of recurrence in a linear network are limited (Fig. 3d), but more substantial in a nonlinear network (Fig. 4).

Overall, the linear model provides important intuition: Hebbian homeostasis is an inevitable consequence of the interaction between Hebbian learning and homeostatic processes in single cells. This stabilizes neural function in the presence of drift; Recurrent dynamics can provide further stability (Fig. 3cd). As we discuss next, further constraints, such as nonlinear recurrent dynamics and response normalization, can confer marked stability.

Nonlinearity and response normalization. Much of the intuition from the linear network extends to the nonlinear case. We assume that neuronal responses are (approximately) locally linear, so the same Hebbian learning rules apply. However, a nonlinear network has key advantages: It is better at correcting errors, and it lets us examine the effect of response normalization on readout stability.

Response normalization controls the average firing rate in a local population of neurons, causing neurons to compete to remain active. It is supported experimentally, and implicated in diverse sensory computations (for review, see (41)). Competition can encourage neurons to acquire diverse tunings, forming a population of localized receptive fields that tile the encoded latent variable space (42, 43).

Nonlinear recurrent networks require specific architectural details to ensure stable dynamics. To avoid this complexity, we model recurrent dynamics and response normalization as discrete transformations. For response normalization, we divide the rates by the average firing rate across the population $\langle \mathbf{y}_f(\theta) \rangle$:

$$\mathbf{y}_d(\theta) = \mathbf{y}_f(\theta) / \langle \mathbf{y}_f(\theta) \rangle \cdot \mu_p, \quad [14]$$

where μ_p is the target average firing rate across the population. For recurrent connections, we train the readout to predict its own activity using fixed set of recurrent weights \mathbf{R} :

$$\mathbf{y}_r(\theta) = \phi[\mathbf{R}^\top \mathbf{y}_d(\theta)] \quad [15]$$

This signal $\mathbf{y}_r(\theta)$ can be used as a training signal to continuously update the forward encoding weights, as in Eq. (5).

Figure 4 summarizes the impact of drift on a nonlinear population readout in several scenarios (Methods: *Nonlinear simulations*). As in the linear case, fixed weights are unstable. Classical homeostasis provides only short term stability. Hebbian homeostasis stabilizes tuning curve statistics, but does not prevent collapse of the population code (Fig. 4a-c).

Surprisingly, response normalization alone improves stability substantially (Fig. 4d). It creates repulsive force between neurons' preferred tunings under the influence of Hebbian plasticity. For the one-dimensional θ explored here, this repulsion constrains the possible rearrangements. Drift must be large to cause two readout neurons to exchange their preferred tunings. Note that tuning curves would be much less constrained in higher dimensional spaces, and we should expect the stabilizing effect of crowding to diminish in higher dimensions.

With recurrent dynamics, the nonlinear readout is exceptionally stable (Fig. 4d). The recurrent weights strongly constrain the correlated activity patterns in $\mathbf{y}(\theta)$, and suppressing any activity that does not match the ring structure learned initially. Drift can only occur along directions of symmetry in the underlying encoded space θ . For the circular θ explored here, drift can rotate the readout, but no other changes are permitted. This illustrates that internal models can strongly constrain network activity, and that these constraints allow populations of neurons to tolerate complete reconfiguration of the inputs that drive them.

Discussion

In this work, we outlined homeostatic principles that could allow stable and plastic representations to coexist in the brain. We argue that self-healing codes should have of three components: (I) Neuronal responses should be tolerate small amounts of drift; (II) Neurons should use their own output as a training signal to update their decoding weights, and (III) Stable codes should homeostatically preserve internal models, which are reflected in stable population statistics.

Here, we considered two populations, one stable and one unstable. This could reflect communication between stable and plastic components of the brain, or the interaction between stable and plastic neurons within the same population. This is consistent with experiments that find consolidated stable representations (44), and with the view that neural populations contain a mixture of stable and unstable cells (45).

However, there is no requirement that a neuron that is stable at present must remain so. Over time, neurons could enter or leave this stable core. As long as some stable neurons

remain, long-term representations could persist. This implies a general principle that supports reallocation of the function of single neurons, while preserving internal models. It also raises the question of whether a stable population is even necessary: could functional stability be achieved by several plastic populations tracking each-other? This points to a potentially powerful generalization of homeostatic principles, which could explain the long-term robustness of distributed neural representations.

Here, we considered how networks might stabilize a pre-existing trained structure. How are these stable representations learned? Once learned, can they be updated? A crucial assumption in our work is that neurons generate their own internal training signals. For single cells, this amounts to error correcting across the pool of its own synaptic inputs. For networks, this corresponds to prediction errors coming from recurrent or top-down dynamics. These error signals are precisely the same ones that would be used for learning from external error feedback. During learning, recurrent and top-down prediction errors propagate high-level reinforcement signals back to local neural populations (34). These prediction errors are carried by the same mechanisms that we use here to achieve homeostasis. Hebbian homeostasis, then, can be viewed as a natural consequence of predictive learning mechanisms in the absence of external error feedback.

The brain supports both consolidated and volatile representations, respectively associated with memory and learning. Artificial neural networks have so far failed to imitate this, and suffer from catastrophic forgetting wherein new learning erases previously learned representation (46). Many strategies have been proposed to mitigate this. Broadly, all of these methods segregate stable and unstable representations into distinct subspaces of the possible synaptic weight changes (c.f. 47). These learning rules therefore amount to preventing disruptive drift in the first place.

The strategies we explore here are fundamentally different. We do not restrict changes in weights or activity: the encoding population is free to reconfigure arbitrarily. However, any change in a neural code leads to an equal and opposite change in how that code is interpreted—The brain must publish new translations of its changing internal language. This constraint preserves the functional relationships between neurons. The approach shares some similarities with approaches to attenuate forgetting using replay during sleep, or the equivalent in artificial networks (e.g. (48, 49)). The internal models must be occasionally re-activated through either rehearsal or replay, in order to detect and correct inconsistencies caused by drift. If this process occurs too infrequently, drift becomes large, and the error correction will fail.

Here, we focused on homeostatic maintenance of function despite drifting population codes. It is worth exploring whether a similar process can explain how the brain preserves learned representations despite neuronal death. In developmental pruning, the brain removes synapses and neurons without loss of function (50). Existing models of pruning require ongoing learning to prevent loss of learned representations (51, 52). Homeostatic preservation of predictive models may allow the brain to benefit from large networks during learning (53–55), and optimize these networks without extensive re-training.

To integrate stable and plastic representations, changes anywhere in the brain must be accompanied by compensatory

changes throughout the brain. The learning rules we explored here placed a particular emphasis on Hebbian homeostasis, and the role of predictive coding in generating robust representations. In the long term, these processes could support widespread reallocation or reconsolidation of neuronal function. Further exploration of these principles may clarify how the brain can be simultaneously plastic and stable, and provide clues to how to build artificial networks that share these properties.

Materials and Methods

Data and analysis. Data shown in Figure 1b,c were taken from Driscoll et al. (4), and are available online at at Dryad (56). Examples of tuning curve drift were taken from mouse four, which tracked a sub-population of cells for over a month. Normalized dF/F calcium transients were band-pass filtered between 0.3 and 3 Hz, and individual trial runs through the T maze were extracted. Calcium fluorescence traces from select cells were aligned based on task pseudotime (0: start, 1: reward). The activity of each cell was z-scored within each trial to yield a normalized log-fluorescence signal. On each day, normalized log-fluorescence was averaged over all trials and then exponentiated to generate the average tuning curves shown in Figure 1b. For Figure 1c, a sub-population of forty cells was selected at random, and sorted based on their peak firing location on the first day. For further details, see (4, 18).

Simulated drift. We sample a random walk on encoding weights \mathbf{U} as an Ornstein Uhlenbeck (OU) process with unit steady-state variance and time constant τ , measured in days. Given τ , and the constraint that the steady-state variance of an OU process is $\frac{1}{2}\tau\sigma^2 = 1$, we set the noise variance to $\sigma^2 = 2/\tau$. In discrete time this is sampled with $\alpha = \sigma^2\Delta t$:

$$u_{ij}^{t+1} = u_{ij}^t\sqrt{1-\alpha} + \sqrt{\alpha}\xi, \quad \xi \sim \mathcal{N}(0,1) \quad [16]$$

This yields an embedding of θ in the activity of the N -dimensional encoding population that changes gradually and randomly over time. The structure of θ encoded in $\mathbf{s}(\theta)$ is inherited by $\mathbf{a}(\theta) = \mathbf{U}^\top \mathbf{s}(\theta)$.

This model preserves the amount of population variability in $\mathbf{a}(\theta)$ driven by θ , in expectation:

$$\langle \|\nabla_{\theta} \mathbf{a}(\theta, t)\|^2 \rangle = N \cdot \text{tr}[\nabla_{\theta} \Sigma(\theta, \theta') \nabla_{\theta'}^\top] = N \cdot \|\nabla_{\theta} \mathbf{s}(\theta)\|^2 \quad [17]$$

In the special case of an exponential nonlinearity $\phi = \exp$, the trace of Fisher information of $\mathbf{x}(\theta)$ is proportional to the average variation in $\mathbf{a}(\theta)$ driven by θ :

$$\text{tr}[\mathcal{I}(\theta)] \propto \langle \|\nabla_{\theta} \ln[\mathbf{x}(\theta, t)]\|^2 \rangle = \langle \|\nabla_{\theta} \mathbf{a}(\theta, t)\|^2 \rangle \quad [18]$$

Formally, the Fisher information is infinite when the noise in \mathbf{x} is zero, but Eq. (18) can be viewed as the zero-variance limit of homogeneous and i.i.d. Gaussian noise with suitable normalization.

In expectation then, this random walk in the encoding weight space preserves the overall population code statistics: It preserves the geometry of θ in the correlations of $\mathbf{a}(\theta)$, and the average amount of information about θ encoded in the population activations.

Weight filtering. We consider a linear version of our encoding-decoding model (Eqs. 2-6), whose weights and activity change across days ("d")

$$\begin{aligned} \mathbf{x}_d(\theta) &= \mathbf{U}_d^\top \mathbf{s}(\theta) \\ \mathbf{y}_d(\theta) &= \mathbf{W}_d^\top \mathbf{x}_d(\theta) \end{aligned} \quad [19]$$

Drift can be viewed as a slow-timescale component of noise, and a readout that is robust to noise can also tolerate some amount of drift. Denote the drift in the code between days as $\Delta \mathbf{x}(\theta)$, and assume that it can be modeled as Gaussian:

$$\Delta \mathbf{x}(\theta) \sim \mathcal{N}(0, \Sigma_{\Delta}) \quad [20]$$

This Gaussian model captures diffusive drift like the OU process Eq. (16) introduced earlier. For training signals $(\mathbf{x}_0, \mathbf{y}_0^*)$, the least-squares optimal weights for day $d+1$ trained on activity on day d is given by regularized linear regression:

$$\mathbf{W}_{d+1} = [\Sigma_d + \Sigma_{\Delta}]^{-1} \Sigma_{0, \mathbf{y}_0^*} \quad [21]$$

where Σ_d is the covariance of $\mathbf{x}_d(\theta)$, and $\Sigma_{0, \mathbf{y}_0^*}$ is the cross covariance between the encoding population activity and the target readout tuning curves \mathbf{y}_0^* .

We needn't estimate these regularized weights from scratch. If we have already weights \mathbf{W}_d trained on day d , then we can prepare regularized weights for the subsequent day \mathbf{W}_{d+1} by updating these existing weights. This also realigns the decoding weights with the correlation structure of the current encoding, $\Sigma_d = \langle \mathbf{x}_d \mathbf{x}_d^\top \rangle$:

$$\mathbf{W}_{d+1} = [\Sigma_d + \Sigma_{\Delta}]^{-1} \Sigma_d \mathbf{W}_d. \quad [22]$$

(c.f. Eq. 9) This is equivalent to using the activity on the current day, \mathbf{x}_{d+1} , to predict the corresponding activity on the previous day \mathbf{x}_d :

$$\hat{\mathbf{x}}_d = \Sigma_d [\Sigma_d + \Sigma_{\Delta}]^{-1} \mathbf{x}_{d+1} \quad [23]$$

Applying Eq. (23) iteratively yields an estimate of the original code $\hat{\mathbf{x}}_0$, thereby translating the current representation \mathbf{x}_d back in time to when the readout was first learned:

$$\hat{\mathbf{y}}(\theta) = \mathbf{W}_0^\top \left\{ \prod_{d'=0..d-1} \Sigma_{d'} [\Sigma_{d'} + \Sigma_{\Delta}]^{-1} \right\} \mathbf{x}_d(\theta). \quad [24]$$

Since the readout activity is driven by these decoding weights, $\mathbf{y}_d = \mathbf{W}_d^\top \mathbf{x}_d$, this recursive filtering can be interpreted by the network re-training itself using its own output:

$$\begin{aligned} \mathbf{y}^* &= \mathbf{W}_d^\top \mathbf{x}_d \\ \mathbf{W}_{d+1} &= [\Sigma_d + \Sigma_{\Delta}]^{-1} \Sigma_d \mathbf{y}^* \\ \Sigma_d \mathbf{y}^* &= \langle \mathbf{x}_d \mathbf{y}^{*\top} \rangle = \langle \mathbf{x}_d \mathbf{x}_d^\top \mathbf{W}_d \rangle = \Sigma_d \mathbf{W}_d \\ \Rightarrow \mathbf{W}_{d+1} &= [\Sigma_d + \Sigma_{\Delta}]^{-1} \Sigma_d \mathbf{W}_d \quad [\text{c.f. Eqs. 9,22}] \end{aligned} \quad [25]$$

To summarize, tracking an unstable code involves filtering the current code-words \mathbf{x}_d to recover the original code \mathbf{x}_0 against which the readout was first trained. In a linear, Gaussian model, this can be computed by iteratively re-training the decoding weights using the network's own output.

The linear Bayesian model (Eq. 19-25) incorporates the assumption that the encoding \mathbf{x} changes, but not that $\Pr(\theta)$ and the primary inputs $\mathbf{s}(\theta)$ are fixed. How might neurons incorporate this? The readout population cannot access $\mathbf{s}(\theta)$, but it could measure its own statistics:

$$\Pr(\mathbf{y}) = \int \mathbf{y}(\theta) \Pr(\theta) d\theta. \quad [26]$$

For example, in the linear model (Eq. 8-12), $\mathbf{y}(\theta)$ is a zero-mean Gaussian variable, so $\Pr(\mathbf{y})$ is encoded fully in the covariance $\Sigma_{\mathbf{y}}$:

$$\Sigma_{\mathbf{y}} = \langle \mathbf{y} \mathbf{y}^\top \rangle = \int \mathbf{y}(\theta) \mathbf{y}(\theta)^\top \Pr(\theta) d\theta \quad [27]$$

Since $\Sigma_{\mathbf{y}}$ is inherited from $\Pr(\mathbf{s}(\theta))$, stable readouts must exhibit stable $\Sigma_{\mathbf{y}}$. The converse is not true, but is a useful constraint that can improve stability. This covariance is readily accessible: its diagonal is simply the firing rate variability of single neurons, and its off-diagonal terms can be encoded in recurrent connections that constrain population activity.

Synaptic learning rules. The homeostatic learning rule Eq. (10) is simple, but unrealistic: it requires tracking the covariance of the encoding population, and solving a linear system by matrix inversion. Neither of these are things that single neurons could do. However, these operations are equivalent to linear regression, which can be computed in an online manner using stochastic gradient descent.

Least Mean Squares (LMS; 57) is an online stochastic gradient descent algorithm that solves the linear regression problem $\mathbf{y} = \mathbf{W}^\top \mathbf{x}$, converging (with noise) to the solution $\mathbf{W} = \Sigma_{\mathbf{x}}^{-1} \Sigma_{\mathbf{x}\mathbf{y}}$, by minimizing the following objective via stochastic gradient descent:

$$\mathbf{w} = \underset{\mathbf{w}}{\text{argmin}} \frac{1}{2} \langle \|\mathbf{w}^\top \mathbf{x} - \mathbf{y}\|^2 \rangle \quad [28]$$

Given a single observation $(\mathbf{x}_t, \mathbf{y}_t)$ at time t , LMS computes the following online weight update:

$$\begin{aligned} \Delta \mathbf{w}_t &\propto -\nabla_{\mathbf{w}_t} \frac{1}{2} \langle \|\mathbf{w}_t^\top \mathbf{x} - \mathbf{y}\|^2 \rangle \\ &= \Sigma_{\mathbf{x}, \mathbf{y}} - \Sigma_{\mathbf{x}} \mathbf{w}_t \\ &\approx \mathbf{x}_t \mathbf{y}_t^\top - \mathbf{x}_t \mathbf{x}_t^\top \mathbf{w}_t \\ &= \mathbf{x}_t (\mathbf{y}_t - \mathbf{w}_t^\top \mathbf{x}_t)^\top. \end{aligned} \quad [29]$$

685 Recall the formula for the filtering weight update, with homeostatic
686 gain re-scaling of $g=\sigma_0/\sigma_y$:

$$687 \quad \mathbf{w}_{d+1} = g \cdot [\Sigma_d + \Sigma_\Delta]^{-1} \Sigma_d \mathbf{w}_d \quad [30]$$

688 This is a batched update, which uses activity on a given day to
689 update the weights for the following day. It minimizes the following
690 objective:

$$691 \quad \begin{aligned} \mathbf{y} &= g \cdot \mathbf{w}^\top \mathbf{x} \\ \mathbf{w} &= \underset{\mathbf{w}}{\operatorname{argmin}} \frac{1}{2} \left\{ \langle \|\mathbf{w}^\top \mathbf{x} - \mathbf{y}\|^2 \rangle + \mathbf{w}^\top \Sigma_\Delta \mathbf{w} \right\} \\ &= \underset{\mathbf{w}}{\operatorname{argmin}} \frac{1}{2} \mathbf{w}^\top \left\{ (1 - g^2) \Sigma_d + \Sigma_\Delta \right\} \mathbf{w} \end{aligned} \quad [31]$$

692 In the online model, we treat drift as occurring gradually and
693 continuously, over small intervals Δt . The incremental drift is
694 therefore $\Delta t \cdot \Sigma_\Delta$, and the homeostatic gain adjustments are small,
695 $g^2 \approx 1 + \gamma \Delta t$. The weight update Eq. (30) for a self-healing code is
696 also a linear least-squares problem. In analogy to LMS, an online
697 stochastic gradient solution for the self-healing weight update is:

$$698 \quad \begin{aligned} \frac{\Delta \mathbf{w}_t}{\Delta t} &\propto -\frac{1}{\Delta t} \left[(1 - g^2) \Sigma_d + \Delta t \cdot \Sigma_\Delta \right] \mathbf{w}_t \\ &\approx [\gamma \mathbf{x}_t \mathbf{x}_t^\top - \Sigma_\Delta] \mathbf{w}_t \\ &= \gamma \mathbf{x}_t \mathbf{y}_t^\top - \Sigma_\Delta \mathbf{w}_t \quad [\text{c.f. Eq. (11)}] \end{aligned} \quad [32]$$

699 This reduces to the Hebbian homeostatic weight update, Eq. (7),
700 with $\gamma = \varepsilon_\sigma$ providing negative feedback to stabilize the neuron's
701 firing-rate variability. Eq. (32) also contains an extra term, $-\Sigma_\Delta \mathbf{w}_t$,
702 which acts as regularizing weight decay. The drift Σ_Δ could be
703 estimated in several ways. It might simply be initialized heuristically
704 as a constant weight decay $\Sigma_\Delta \propto I$. It is also possible to use changes
705 in neuronal variability as a proxy for drift.

706 **Estimating the rate of drift.** Empirically, we observe that the popu-
707 lation statistics for the tuning curves $\mathbf{x}(\theta)$ are stable despite drift
708 (4). The tuning curve of each encoding cell $x(\theta)$ can be viewed as
709 a vector in this space of possible tuning curves. For large popula-
710 tions, the total amount of task-related variability is approximately
711 conserved. This implies that drift is, on average, mostly rotational.
712 If rotational drift rotates our code by amount ϕ away from the
713 subspace spanned by our current decoding weights, it will lead to a
714 loss of drive to the readout neurons, which is approximately $\cos^2(\phi)$.

715 The homeostatic gain adjustment acts based on the observed loss
716 of drive. Assuming our target variance is one, $\sigma_t^2=1$, a variability
717 decreases of $\cos^2(\phi)$ requires a gain adjustment of $g=1/\cos(\phi)$. For
718 small amounts of drift, a first-order Taylor expansion yields $\gamma \approx \phi^2$.
719 The current value γ is therefore also an estimate of the drift rate,
720 i.e. $\hat{\Sigma}_\Delta \approx \gamma I$, and one may write:

$$721 \quad \frac{\Delta \mathbf{w}_t}{\Delta t} \approx \gamma \cdot [\mathbf{x}_t \mathbf{x}_t^\top - I] \mathbf{w}_t \quad [33]$$

722 Another way to arrive at Eq. (33) is to assume that drift (and
723 therefore any compensatory weight changes) should be tangent to the
724 current decoding weight vector (Figure 3a). This has an intuitive
725 interpretation: if we assume that the encoding of θ is stable over
726 time at the population level, then we know that there is always
727 some linear combination of decoding weights that can read out a
728 target tuning curve $y(\theta)$ from $\mathbf{x}(\theta)$. That is, the overall statistics of
729 the weight vector should also be stable. Drift causes these decoding
730 weights to point in a slightly different direction. Tracking drift
731 therefore amounts to rotating the weight vector to point in this
732 new direction. Large reconfigurations of the encoding space can
733 therefore be tracked if drift is gradual (Figure 3b).

734 To ensure that gain homeostasis can converge in the absence of
735 drift, one might use a faster learning rate $\eta_\gamma > 1$ for gain adjustment,
736 which amounts to:

$$737 \quad \frac{\Delta \mathbf{w}_t}{\Delta t} \approx \gamma \cdot [\eta_\gamma \mathbf{x}_t \mathbf{x}_t^\top - I] \mathbf{w}_t \quad [34]$$

738 **Linear network with recurrence.** So far, we have explored self-healing
739 codes in the case of a single neuron, which uses a measurement of its
740 own variability to detect and correct for drift. One way to extend
741 this to populations is to assume that the activity in the readout,
742 \mathcal{Y} , is constrained by local recurrent connections. This recurrent
743 activity provides additional error correction (27). In this scenario,

744 the decoding weights and recurrent connections incorporate the
745 prior knowledge that Σ_y should remain stable over time.

746 A simple version of this mechanism might use feed-forward
747 activity $\mathbf{y}_f = \mathbf{W}^\top \mathbf{x}$ to generate regularized predictions \mathbf{y}_r . This
748 regularized estimate might be computed via local, recurrent weights
749 \mathbf{R} that encode a fixed prior model of Σ_y :

$$\begin{aligned} \mathbf{y}_f &= \mathbf{W}^\top \mathbf{x} \\ \mathbf{y}_r &= \mathbf{R}^\top \mathbf{y}_f \\ \mathbf{R} &= [\Sigma_y + \kappa I]^{-1} \Sigma_y, \end{aligned} \quad [35] \quad 750$$

751 where κ sets the strength of the regularization in the recurrent
752 dynamics.

753 This pools information across the readout population by linearly
754 predicting the readout's activity from itself, with regularization
755 strength α . This can also be viewed as Gaussian process (GP)
756 smoothing, where Σ_y encodes the GP prior kernel using the "true"
757 tuning curves $\mathbf{y}(\theta)$ to support the function space. Eq. (35) can be
758 computed as a steady-state solution of a recurrent network that
759 computes a prediction error $\mathbf{W}^\top \mathbf{x} - \mathbf{y}$ using inhibitory feedback:

$$760 \quad \tau \dot{\mathbf{y}} = -\mathbf{y} + \tau \Sigma_y [\mathbf{W}^\top \mathbf{x} - \mathbf{y}], \quad [36]$$

761 where $\tau = 1/\kappa$. If \mathbf{x} varies slowly relative to the time constant τ ,
762 and if Eq. (36) converges, then it converges to Eq. (35), and tracks
763 $\mathbf{y}_r(t)$. We stop short of specifying a specific biological realization
764 of Eq. (36), but this feedback-based solution is consistent with
765 the prevailing theory that the brain learns and computes using
766 prediction errors (e.g. 58). Recurrent feedback yields a new error
767 signal, $\mathbf{y}_r - \mathbf{y}_f$ that detects when the decoded activity strays outside
768 of the low-dimensional subspace of the initial code, $\mathbf{y}_0(\theta)$. This error
769 can be added to the weight update Eq. (34) to yield a combined
770 update that reflects two constraints: Hebbian homeostasis, and
771 local recurrent dynamics (Results, Eq. 13).

772 In this form, it becomes clear that the recurrent dynamics in $\mathbf{y}(\theta)$
773 truly are predictive dynamics. A Hebbian rule which tracks drift is,
774 essentially, minimizing the errors in the online predictions that \mathbf{y}
775 makes about the activity \mathbf{x} . In this paper, we consider only the case
776 where \mathbf{x} changes so slowly that this prediction should be the identity
777 map. However, in a scenario where \mathbf{x} has nontrivial temporal
778 dynamics, such recurrent computations and learning inherently
779 learn an asymmetric model that captures how θ evolves in time.

780 **Linear simulations.** We simulated a self-healing linear network en-
781 coding a circular latent variable $\theta \in [0, 2\pi)$, discretized into $L=60$
782 spatial bins. We sampled $K=200$ randomly-drifting spatial features
783 $\mathbf{x}(\theta)$ from a Gaussian process on θ , with an exponentiated quadratic
784 (i.e. radial basis; Gaussian) covariance kernel with a spatial stan-
785 dard deviation of $\sigma_l=9$ bins, scaled so that the standard deviation
786 of each feature was $s=0.15$. These features underwent Ornstein Uh-
787 lenbeck drift over time, with a time-constant of $\tau=50$ days (Eq. 16).
788 $M=50$ readout units $\mathbf{y}(\theta)$ were initialized with bump-like tuning
789 curves, modeled as Gaussians with $\sigma_y=9$ bins, evenly distributed
790 over a range of preferred tunings θ_0 . These readouts were given a
791 homeostatic target variance of $\sigma_t=1$.

792 We simulated 500 days of drift—ten times of the correlation time
793 for the drifting encoding features. This allowed multiple complete
794 reconfigurations of the encoding population. We simulated Hebbian
795 homeostasis using a continuous-time learning rule (Eq. 34) applied
796 for 500 time-steps on each day, with a learning rate of 1×10^{-5} per
797 step. These updates were batched: rather than sampling individual
798 stimuli and using $\mathbf{x}\mathbf{x}^\top \mathbf{W}$ to calculate updates in stochastic gradient
799 descent, we directly apply the expectation $\Sigma_x \mathbf{W}$.

800 We evaluated three scenarios: fixed weights, Hebbian homeo-
801 stasis, and Hebbian homeostasis with recurrent prediction errors
802 (Figure 3). We modeled recurrence as an additional linear map
803 $\mathbf{y}_r = \mathbf{R}^\top \mathbf{y}_f$ as in Equation Eq. (35), and the resulting \mathbf{y}_r was used
804 as a training signal in a batched least-mean-squares continuous-time
805 gradient update (Eq. 29). To summarize the relative performance of
806 these three scenarios (Fig. 3d), we sampled 20 random realizations
807 of the aforementioned simulations.

808 The ability of the linear model to error-correct is limited by the
809 amount of drift that projects onto the low-dimensional subspace
810 in $\mathbf{x}(\theta)$ that encodes θ . While the total amount of drift increases
811 for larger populations, averaging predicts that the disruptive effect

of drift (in terms of squared error) should scale inversely with population size. To verify this, we simulated a range of models with different degrees of redundancy. We simplified the input features $\mathbf{s}(\theta)$ to reflect a K -dimensional Gaussian variable θ , encoded in an $N > K$ population. The readout $\mathbf{y}(\theta)$ was initialized to recover θ via linear regression. As above, we simulated 500 days of random drift as an O.U. process on the encoding weights, using Hebbian homeostasis (without recurrence). For each network realization, we sampled ten instances of the initial features and network, and then five independent realizations of random drift for each instance.

Nonlinear simulations. For the nonlinear readout, we simulated a circular variable $\theta \in [0, 2\pi)$ divided into $L=60$ discrete bins. We sampled $K=60$ features $\mathbf{s}(\theta)$ from a Gaussian process on θ , with zero mean and an exponentially quadratic covariance kernel with standard deviation $\sigma_l=15$. We allowed individual encoding units $x_n(\theta)$ to change abruptly, rather than undergo a continuous random walk. We did this by re-sampling features one-at-a-time, and running Hebbian homeostasis each time 8% of the encoding features changed. This approach emphasized that the nonlinear readout can track drift through multiple complete reconfigurations of the encoding population. Encoding features were normalized to range from 0 to 1, then passed through a nonlinearity $\mathbf{x}(\theta) = \exp[z(\theta) - \frac{1}{2}]$ to simulate sparse, non-negative network inputs.

We initialized $N=60$ linear-nonlinear readout neurons ($\mathbf{y}(\theta)$; Eq. 2) with Gaussian tuning curves $\mathbf{y}_0(\theta)$ (standard deviation $\sigma_y=5$ bins), and with preferred tunings θ_0 evenly distributed on $[0, 2\pi)$. Readout weights \mathbf{W} were trained via gradient descent to minimize a loss similar to a log-linear Poisson model.

$$\mathbf{W} = \underset{\mathbf{W}}{\operatorname{argmin}} \left\langle \exp[\mathbf{W}^\top \mathbf{x}] - \mathbf{y}_0 \circ \mathbf{W}^\top \mathbf{x} \right\rangle + \kappa \|\mathbf{W}\|^2, \quad [37]$$

where \circ denotes element-wise multiplication, the expectation $\langle \cdot \rangle$ is taken over θ and the readout population, and the regularization strength is $\kappa=10^{-2}$. The homeostatic set-points for the mean and the variance of the firing rate (μ_t, σ_t^2) were taken from the statistics of these initial tuning curves.

We implemented Hebbian homeostasis by defining slow variables γ and β , which track the deviations of the neuron's firing rate statistics from its homeostatic set points. Weights were trained to restore these set-points via a continuous-time Hebbian learning rule (Eq. 7). 50 iterations of this learning rule were applied each time 8% (5 out of 60) of the encoding population had reconfigured. For nonlinear neurons, homeostasis of the mean-rate and variability interact. Controlling the variability can change the overall excitability of the neuron, and can lead to instability. To address this, we used different learning rates $\eta_\beta=0.9$ and $\eta_\sigma=0.1$ for the mean-rate and variability, respectively.

To simulate response normalization, we divided the response $\mathbf{y}(\theta)$ by the average population rate, scaled to preserve the population rates seen in the initial network configuration, as in Equation Eq. (14). To model recurrent dynamics, we trained another set of fixed recurrent weights \mathbf{R} as in Equation Eq. (15), with a gradient descent objective similar to the one used to initialize the decoding weights (Eq. 37).

$$\mathbf{R} = \underset{\mathbf{R}}{\operatorname{argmin}} \left\langle \exp[\mathbf{R}^\top \mathbf{y}_0] - \mathbf{y}_0 \circ \mathbf{R}^\top \mathbf{y}_0 \right\rangle + \kappa_r \|\mathbf{R}\|^2 \quad [38]$$

with regularization strength of $\kappa_r=10^{-4}$.

These recurrent predictions yield a revised prediction $\mathbf{y}_r(\theta)$ after applying response normalization. For both response normalization and the recurrent model, "error-corrected" estimates $\hat{\mathbf{y}}=\mathbf{y}_d$ or $\hat{\mathbf{y}}=\mathbf{y}_r$ were used to retrain the decoding weights via Hebbian learning, with regularizing weight decay rate of $\rho_d=\frac{1}{3} \times 10^{-3}$:

$$\begin{aligned} \mathbf{y}_f &= \exp(\mathbf{U}^\top \mathbf{x}) \\ \Delta \mathbf{U} &= \eta \left\langle \mathbf{x} \left[\hat{\mathbf{y}} - \mathbf{y}_f \right]^\top \right\rangle - \rho_d \mathbf{U}, \end{aligned} \quad [39]$$

with a learning rate of $\eta=0.5$. The above Eq. (39) corresponds to online gradient descent of an objective similar to those used to train the initial forward and recurrent weights (Eqs. 37, 38).

ACKNOWLEDGMENTS. This work was supported by ERC grant StG 716643 FLEXNEURO and HFSP grant RGY0069.

1. AR Chambers, S Rumpel, A stable brain from unstable components: emerging concepts and implications for neural computation. *Neuroscience* **357**, 172–184 (2017). 877
2. Y Ziv, et al., Long-term dynamics of ca1 hippocampal place codes. *Nat. neuroscience* **16**, 264 (2013). 878
3. SJ Levy, NR Kinsky, W Mau, DW Sullivan, ME Hasselmo, Hippocampal spatial memory representations in mice are heterogeneously stable. *bioRxiv*, 843037 (2019). 880
4. LN Driscoll, NL Pettit, M Minderer, SN Chetlin, CD Harvey, Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell* **170**, 986–999 (2017). 881
5. A Singh, A Peyrache, MD Humphries, Medial prefrontal cortex population activity is plastic irrespective of learning. *J. Neurosci.* **39**, 3470–3483 (2019). 882
6. BR Cowley, et al., Slow drift of neural activity as a signature of impulsivity in macaque visual and prefrontal cortex. *Neuron* **108**, 551–567 (2020). 883
7. D Deitch, A Rubin, Y Ziv, Representational drift in the mouse visual cortex. *bioRxiv* (2020). 884
8. TD Marks, MJ Goad, Stimulus-dependent representational drift in primary visual cortex. *bioRxiv* (2020). 885
9. CE Schoonover, SN Ohashi, R Axel, AJ Fink, Representational drift in primary olfactory cortex. *bioRxiv* (2020). 886
10. E Marder, JM Goaillard, Variability, compensation and homeostasis in neuron and network function. *Nat. Rev. Neurosci.* **7**, 563–574 (2006). 887
11. GG Turrigiano, KR Leslie, NS Desai, LC Rutherford, SB Nelson, Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* **391**, 892–896 (1998). 888
12. J Cannon, P Miller, Stable control of firing rate mean and variance by dual homeostatic mechanisms. *The J. Math. Neurosci.* **7**, 1–38 (2017). 889
13. GG Turrigiano, The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell* **135**, 422–435 (2008). 890
14. F Zenke, W Gerstner, Hebbian plasticity requires compensatory processes on multiple timescales. *Philos. Transactions Royal Soc. B: Biol. Sci.* **372**, 20160259 (2017). 891
15. YK Wu, KB Hengen, GG Turrigiano, J Gjorgjieva, Homeostatic mechanisms regulate distinct aspects of cortical circuit dynamics. *Proc. Natl. Acad. Sci.* **117**, 24514–24525 (2020). 892
16. T O'Leary, AH Williams, JS Caplan, E Marder, Correlations in ion channel expression emerge from homeostatic tuning rules. *Proc. Natl. Acad. Sci.* **110**, E2645–E2654 (2013). 893
17. ME Rule, T O'Leary, CD Harvey, Causes and consequences of representational drift. *Curr. opinion neurobiology* **58**, 141–147 (2019). 894
18. ME Rule, et al., Stable task information from an unstable neural population. *Elife* **9**, e51121 (2020). 895
19. M Gillett, U Pereira, N Brunel, Characteristics of sequential activity in networks with temporally asymmetric hebbian learning. *Proc. Natl. Acad. Sci.* **117**, 29948–29958 (2020). 896
20. DV Raman, T O'Leary, Optimal synaptic dynamics for memory maintenance in the presence of noise. *BioRxiv* (2020). 897
21. D Acker, S Paradis, P Miller, Stable memory and computation in randomly rewiring neural networks. *J. neurophysiology* (2019). 898
22. L Susman, N Brenner, O Barak, Stable memory with unstable synapses. *Nat. communications* **10**, 1–9 (2019). 899
23. MJ Fauth, MC van Rossum, Self-organized reactivation maintains and reinforces memories despite synaptic turnover. *ELife* **8**, e43717 (2019). 900
24. FYK Kossio, S Goedeke, C Klos, RM Memmesheimer, Drifting assemblies for persistent memory. *bioRxiv* (2020). 901
25. C Stringer, M Pachitariu, N Steinmetz, M Carandini, KD Harris, High-dimensional geometry of population responses in visual cortex. *Nature* **571**, 361–365 (2019). 902
26. ZP Kilpatrick, B Ermentrout, B Doiron, Optimizing working memory with heterogeneity of recurrent cortical excitation. *J. neuroscience* **33**, 18999–19011 (2013). 903
27. MF Panichello, B DePasquale, JW Pillow, TJ Buschman, Error-correcting dynamics in visual working memory. *Nat. communications* **10**, 1–11 (2019). 904
28. JA Gallego, MG Perich, LE Miller, SA Solla, Neural manifolds for the control of movement. *Neuron* **94**, 978–984 (2017). 905
29. A Rubin, et al., Revealing neural correlates of behavior without behavioral measurements. *Nat. Commun.* **10**, 4745 (2019). 906
30. JA Gallego, MG Perich, RH Chowdhury, SA Solla, LE Miller, Long-term stability of cortical population dynamics underlying consistent behavior. *Nat. neuroscience* **23**, 260–270 (2020). 907
31. E Sorrell, ME Rule, T O'Leary, Brain-machine interfaces: Closed-loop control in an adaptive system. *Annu. Rev. Control. Robotics, Auton. Syst.* **4** (2021). 908
32. NA Cayco-Gajic, C Clopath, RA Silver, Sparse synaptic connectivity is required for decorrelation and pattern separation in feedforward networks. *Nat. communications* **8**, 1–11 (2017). 909
33. L Chen, et al., The role of intrinsic excitability in the evolution of memory: Significance in memory allocation, consolidation, and updating. *Neurobiol. Learn. Mem.* **173**, 107266 (2020). 910
34. JC Whittington, R Bogacz, An approximation of the error backpropagation algorithm in a predictive coding network with local hebbian synaptic plasticity. *Neural computation* **29**, 1229–1262 (2017). 911
35. P Berkes, G Orbán, M Lengyel, J Fiser, Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* **331**, 83–87 (2011). 912
36. SE Palmer, O Marre, MJ Berry, W Bialek, Predictive information in a sensory population. *Proc. Natl. Acad. Sci.* **112**, 6908–6913 (2015). 913
37. M Wilf, et al., Spontaneously emerging patterns in human visual cortex reflect responses to naturalistic sensory stimuli. *Cereb. cortex* **27**, 750–763 (2017). 914
38. V Koren, S Denève, Computational account of spontaneous activity as a signature of predictive coding. *PLoS computational biology* **13**, e1005355 (2017). 915
39. GB Keller, TD Mrisic-Flogel, Predictive processing: a canonical cortical computation. *Neuron* **100**, 424–435 (2018). 916
40. E Oja, Simplified neuron model as a principal component analyzer. *J. mathematical biology* **15**, 267–273 (1982). 917
41. M Carandini, DJ Heeger, Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13**, 51–62 (2012). 918
42. A Sengupta, C Pehlevan, M Tepper, A Genkin, D Chklovskii, Manifold-tiling localized receptive fields are optimal in similarity-preserving neural networks in *Advances in Neural Information* 919

- 961 *Processing Systems*. pp. 7080–7090 (2018).
- 962 43. D Krotov, JJ Hopfield, Unsupervised learning by competing hidden units. *Proc. Natl. Acad.*
963 *Sci.* **116**, 7723–7731 (2019).
- 964 44. KA Katlowitz, MA Picardo, MA Long, Stable sequential activity underlying the maintenance
965 of a precisely executed skilled behavior. *Neuron* **98**, 1133–1140 (2018).
- 966 45. T Hainmueller, M Bartos, Parallel emergence of stable and dynamic memory engrams in the
967 hippocampus. *Nature* **558**, 292–296 (2018).
- 968 46. RM French, Catastrophic forgetting in connectionist networks. *Trends cognitive sciences* **3**,
969 128–135 (1999).
- 970 47. L Duncker, L Driscoll, KV Shenoy, M Sahani, D Sussillo, Organizing recurrent network dy-
971 namics by task-computation to enable continual learning. *Adv. Neural Inf. Process. Syst.* **33**
972 (2020).
- 973 48. S Káli, P Dayan, Off-line replay maintains declarative memories in a model of hippocampal-
974 neocortical interactions. *Nat. neuroscience* **7**, 286 (2004).
- 975 49. OC González, Y Sokolov, GP Krishnan, JE Delanois, M Bazhenov, Can sleep protect memo-
976 ries from catastrophic forgetting? *Elife* **9**, e51005 (2020).
- 977 50. E Bullmore, O Sporns, The economy of brain network organization. *Nat. Rev. Neurosci.* **13**,
978 336–349 (2012).
- 979 51. EJ Crowley, J Turner, A Storkey, M O’Boyle, Pruning neural networks: is it time to nip it
980 in the bud? in *Workshop on Compact Deep Neural Network Representation with Industrial*
981 *Applications at the Thirty-second Conference on Neural Information Processing Systems*.
982 Vol. 32, (2018).
- 983 52. C Scholl, ME Rule, MH Hennig, The information theory of developmental pruning: Optimizing
984 global network architecture using local synaptic rules. *bioRxiv* (2020).
- 985 53. G Chechik, I Meilijson, E Ruppin, Synaptic pruning in development: a computational account.
986 *Neural computation* **10**, 1759–1777 (1998).
- 987 54. DV Raman, AP Rotondo, T O’Leary, Fundamental bounds on learning performance in neural
988 circuits. *Proc. Natl. Acad. Sci.* **116**, 10537–10546 (2019).
- 989 55. J Steinberg, M Advani, H Sompolinsky, New role for circuit expansion for learning in neural
990 networks. *Phys. Rev. E* **103**, 022404 (2021).
- 991 56. LN Driscoll, NL Pettit, M Minderer, SN Chettih, CD Harvey, Data from: Dynamic re-
992 organization of neuronal activity patterns in parietal cortex dataset. *Dryad (Dataset)*
993 <https://doi.org/10.5061/dryad.gqnk98sjq> (2020).
- 994 57. B Widrow, ME Hoff, Adaptive switching circuits, (Stanford Univ Ca Stanford Electronics Labs),
995 Technical report (1960).
- 996 58. C Savin, S Deneve, Spatio-temporal representations of uncertainty in spiking neural networks.
997 in *NIPS*. Vol. 27, pp. 2024–2032 (2014).

DRAFT