# Bayesian Literature Review

ar007 · 19 May 2026

| Paper | Year |
|---|---|
| Ma, Beck, Latham & Pouget — *Bayesian Inference with Probabilistic Population Codes* | 2006 |
| Pouget, Dayan & Zemel — *Inference and Computation with Population Codes* | 2003 |
| Fiser, Berkes, Orbán & Lengyel — *Statistically Optimal Perception and Learning: From Behavior to Neural Representations* | 2010 |
| Orbán, Berkes, Fiser & Lengyel — *Neural Variability and Sampling-Based Probabilistic Representations in the Visual Cortex* | 2016 |
| Aitchison & Lengyel — *The Hamiltonian Brain: Efficient Probabilistic Inference with Excitation-Inhibition Networks* | 2016 |
| Echeveste, Aitchison, Hennequin & Lengyel — *Cortical-like Dynamics in Recurrent Circuits Optimized for Sampling-Based Probabilistic Inference* | 2020 |
| Padamsey, Katsanevaki, Dupuy & Rochefort — *Neocortex Saves Energy by Reducing Coding Precision during Food Scarcity* | 2022 |

## Terms to know first

The papers above share a small vocabulary that runs across all of them. If a term feels hazy, look it up before reading — they are not redefined in each paper.

**Bayesian fundamentals.**

- *Prior, likelihood, posterior.* $p(\theta)$, $p(x \mid \theta)$, $p(\theta \mid x) \propto p(x \mid \theta)p(\theta)$. The posterior is what every paper here claims the brain represents in some form.
- *Generative model.* The forward model the brain is assumed to have inverted: latent causes → sensory observations.
- *Recognition / inference.* The inverse problem — recovering $p(\theta \mid x)$ from $x$ — performed approximately by the circuit.
- *Marginalisation.* Integrating out nuisance variables. Cortical circuits are claimed to do this in different ways depending on whether the representation is parametric or sample-based.

**Two flavours of neural representation of probability.** This is the central split in the reading list and the two camps disagree on it.

- *Parametric (probabilistic population code, PPC).* The instantaneous population firing rate $\boldsymbol{r}$ *encodes the parameters* of a distribution over $\theta$ — for an exponential-family code, $p(\theta \mid x) \propto \exp\big(\boldsymbol{h}(\theta)^\top \boldsymbol{r}\big)$. Posterior uncertainty is read off the population gain. Ma et al. 2006 and Pouget et al. 2003 are the canonical PPC papers.

- *Sampling-based.* The instantaneous activity $\boldsymbol{r}(t)$ *is* a sample from the posterior, $\boldsymbol{r}(t) \sim p(\theta \mid x)$. Uncertainty is read off the trial-to-trial variability over time. Fiser et al. 2010, Orbán et al. 2016, and the Lengyel-group papers (Aitchison & Lengyel 2016, Echeveste et al. 2020) are the canonical sampling-camp papers.
- *Variational inference (background).* A third flavour — fit a parametric $q(\theta)$ to the posterior by minimising KL divergence. Worth knowing as the ML reference point, though none of these papers commit to it as the brain's algorithm.

**Population-code mechanics.**

- *Tuning curve* $f_i(\theta)$ — neuron $i$'s mean rate as a function of stimulus $\theta$.
- *Fisher information* $I_F(\theta)$ — the precision a code carries about $\theta$; for an unbiased estimator, $\mathrm{Var}\left(\hat{\theta}\right) \geq 1/I_F(\theta)$ (Cramér–Rao). Often discussed per spike, to compare codes at matched energy.
- *Linear Fisher information.* The piece of $I_F$ accessible by a linear readout — what a downstream neuron with synaptic weights can actually extract.
- *Noise correlations* $c_{ij}$ — trial-to-trial covariance between neurons at fixed stimulus. Information-limiting correlations are the slice of $c_{ij}$ aligned with the tuning gradient $\partial f/\partial\theta$; they cap $I_F$ no matter how many neurons you average.
- *Poisson variability.* Spike-count variance ≈ mean. The baseline noise model PPCs are built on.

**Sampling-camp specifics.**

- *Langevin dynamics.* Stochastic gradient on $-\log p(\theta \mid x)$ plus white noise — the simplest dynamical system whose stationary distribution is the posterior. The recurrent-circuit papers are versions of this idea with biological constraints.
- *Hamiltonian Monte Carlo (HMC).* Sampling with momentum — propose moves along trajectories of a fictitious Hamiltonian, accept by Metropolis. Aitchison & Lengyel argue E/I circuits implement an HMC-like sampler with the inhibitory population playing the momentum role.
- *Inhibition-stabilised network (ISN).* A recurrent E/I circuit whose excitatory subnetwork is *unstable* on its own and stabilised only by inhibition. Empirically common in cortex; the sampler papers exploit ISN dynamics for fast, well-mixed sampling.
- *Mixing time, autocorrelation time.* How quickly a sampler decorrelates from its current state — sets how fast a circuit can revise its uncertainty estimate. Echeveste et al. argue cortical-like transients arise from optimising for fast mixing.
- *Stochastic stability vs. transients.* A sampler must be globally stable but locally fluctuating. The non-trivial transient responses cortical circuits show are read as evidence of a sampler near, not at, equilibrium.

**Coding-cost and metabolism.**

- *Bits per spike.* Mutual information per spike between stimulus and response — the efficiency the energetics papers ask about.
- *Coding precision.* The width of the represented posterior; can be widened to save metabolic cost without changing the represented mean. Padamsey et al. 2022 show this

happens *in vivo* under food scarcity — directly tying the Bayesian and metabolic stories together.

- *Metabolic cost of a spike.* Sodium-pump ATP per action potential; the currency the precision–cost trade-off is denominated in.

**Useful background that none of these papers stop to define.**

- *Exponential family, sufficient statistics, natural parameters.* PPCs lean on this — the firing rates are the natural parameters of an exponential-family posterior.
- *KL divergence.* The asymmetric distance between distributions, used to score how good an approximate posterior is.
- *Stationary distribution of a Markov chain.* The distribution the chain converges to and samples from at long times. The sampling-camp claim is that the cortical chain's stationary distribution is the posterior.