Polyphonic Intelligence: Constraint-Based Emergence, Pluralistic Inference, and Non-Dominating Integration: Supplementary

Alexander D. Shaw
Computational Psychiatry & Neuropharmacological Systems (CPNS) Lab
Department of Psychology, Faculty of Health & Life Sciences
University of Exeter, UK
https://cpnslab.com

Conventional Active Inference (single generative model)

Generative model. Let $o_{1:T}$ denote observations, $s_{1:T}$ hidden states, and $a_{1:T}$ actions. Under a single generative model m, we assume a joint density

$$p(o_{1:T}, s_{1:T} \mid m) = p(s_1 \mid m) \prod_{t=1}^{T} p(o_t \mid s_t, m) \prod_{t=1}^{T-1} p(s_{t+1} \mid s_t, a_t, m). \tag{1}$$

This specifies (i) a prior over initial states, (ii) an observation model (likelihood), and (iii) controlled dynamics.

Variational free energy (VFE). Active Inference typically performs approximate Bayesian inference by maintaining a variational posterior $q(s_{1:T})$ and minimising variational free energy

$$F[q] = \mathbb{E}_{q(s_{1:T})}\big[\ln q(s_{1:T}) - \ln p(o_{1:T}, s_{1:T} \mid m)\big]. \tag{2}$$

Minimising $F[q]$ tightens a bound on surprise $-\ln p(o_{1:T} \mid m)$ and makes $q(s_{1:T})$ approximate $p(s_{1:T} \mid o_{1:T}, m)$.

Equivalent decomposition (energy–entropy form).

$$F[q] = \underbrace{\mathbb{E}_q\big[-\ln p(o_{1:T}, s_{1:T} \mid m)\big]}_{\text{expected energy}} - \underbrace{\mathbb{E}_q\big[-\ln q(s_{1:T})\big]}_{\text{entropy of } q}. \tag{3}$$

Posterior update (perception).

$$q^*(s_{1:T}) = \arg\min_q F[q]. \tag{4}$$

In practice this is implemented via gradient flows, Laplace/VL updates, variational message passing, or other approximate inference schemes.

Policies and expected free energy (planning). Let π denote a policy (a sequence of future actions). Active Inference selects policies by minimising expected free energy

$$G(\pi) = \mathbb{E}_{q(o_\tau, s_\tau \mid \pi)}\big[\ln q(s_\tau \mid \pi) - \ln p(o_\tau, s_\tau \mid m)\big], \tag{5}$$

where τ indexes future time points and $q(o_\tau, s_\tau \mid \pi)$ is the policy-conditioned predictive density.

Common decomposition of expected free energy. A widely used decomposition is

$$G(\pi) = \underbrace{\mathbb{E}_{q(o_\tau \mid \pi)}\big[-\ln p(o_\tau)\big]}_{\text{risk (preference violation)}} + \underbrace{\mathbb{E}_{q(o_\tau \mid \pi)}\big[H(p(o_\tau \mid s_\tau))\big]}_{\text{ambiguity}} - \underbrace{\mathbb{E}_{q(o_\tau \mid \pi)}\big[\mathrm{IG}(s_\tau; o_\tau \mid \pi)\big]}_{\text{epistemic value}}, \tag{6}$$

where $p(o_\tau)$ encodes prior preferences over outcomes, $H(\cdot)$ is entropy, and IG denotes information gain (e.g., $\mathrm{IG}(s; o) = \mathrm{KL}(q(s \mid o) \| q(s))$). Exact forms vary with factorisation assumptions.

Policy posterior with precision. Policies are typically selected using a softmax (Boltzmann) distribution

$$p(\pi) = \sigma(-\beta\, G(\pi)) \quad \propto \quad \exp(-\beta\, G(\pi))\,, \tag{7}$$

where β is an inverse temperature (policy precision) controlling stochasticity of policy selection.

Polyphonic Active Inference (multiple generative models / "voices")

Ensemble of generative models (voices). Assume a set of K generative models (voices)

$$\mathcal{M} = \{m_1, m_2, \ldots, m_K\}. \tag{8}$$

Each voice k supports its own latent-state posterior $q_k(s_{1:T})$ under its own assumptions (e.g., precisions, priors, dynamics, observation mappings).

Voice-specific variational free energy.

$$F_k[q_k] = \mathbb{E}_{q_k(s_{1:T})}\big[\ln q_k(s_{1:T}) - \ln p(o_{1:T}, s_{1:T} \mid m_k)\big]. \tag{9}$$

Each voice can be updated using the same inference machinery as conventional Active Inference (e.g., Laplace/VL, message passing), but applied independently within each m_k.

Non-dominating integration: polyphonic free energy. Polyphonic intelligence combines local objectives while penalising destructive inconsistency:

$$\mathcal{F}_{\text{poly}} = \sum_{k=1}^{K} \pi_k\, F_k[q_k] \; + \; \sum_{i<j} \lambda_{ij}\, C(q_i, q_j)\,. \tag{10}$$

Here, $\pi_k \geq 0$ are *credence weights* (with $\sum_k \pi_k = 1$), $C(q_i, q_j)$ is a consistency cost (soft alignment), and $\lambda_{ij} \geq 0$ are coupling strengths. Crucially, there is no hard model selection (no winner-takes-all pruning).

Examples of consistency costs. A simple choice is alignment in predicted outcomes:

$$C(q_i, q_j) = \mathbb{E}_{q_i(o_\tau)}[\phi(o_\tau)] - \mathbb{E}_{q_j(o_\tau)}[\phi(o_\tau)] \quad \Rightarrow \quad C_{ij} = \|\mu_i - \mu_j\|^2 , \tag{11}$$

where $\phi(\cdot)$ is a feature map (e.g., goal-relevant summaries) and μ_k denotes the corresponding predicted feature mean. Alternative choices include KL divergences between predictive distributions, or penalties on incompatible latent factors.

Polyphonic inference (local updates under coupling). A generic coupled update can be written as

$$q_k^* = \arg\min_{q_k} \left[\pi_k \, F_k[q_k] + \sum_{j \neq k} \lambda_{kj} \, C(q_k, q_j) \right] , \tag{12}$$

which reduces to standard Active Inference when $K = 1$ and all coupling terms vanish.

Voice-specific expected free energy (planning). Each voice evaluates policies using its own predictive density:

$$G_k(\pi) = \mathbb{E}_{q_k(o_\tau, s_\tau | \pi)} \big[\ln q_k(s_\tau \mid \pi) - \ln p(o_\tau, s_\tau \mid m_k) \big] . \tag{13}$$

As in the single-model case, $G_k(\pi)$ can be decomposed into risk, ambiguity, and epistemic value with respect to the voice's generative assumptions.

Polyphonic policy value (non-dominating integration for control). Define a population-level policy objective

$$G_{\text{poly}}(\pi) = \sum_{k=1}^{K} \pi_k^{\text{ctrl}} \Big(G_k(\pi) + \lambda \, C_k(\pi) \Big), \tag{14}$$

where π_k^{ctrl} are *control influence weights* (not necessarily equal to π_k), and $C_k(\pi)$ is a policy-level alignment penalty. For example, aligning on a goal-progress statistic $r_k(\pi)$:

$$C_k(\pi) = (r_k(\pi) - \bar{r}(\pi))^2 , \qquad \bar{r}(\pi) = \sum_{j=1}^{K} \pi_j \, r_j(\pi). \tag{15}$$

This encourages agreement on coarse progress signals while allowing persistent disagreement elsewhere.

Decoupling credence from control. A simple convex mixing that prevents executive domination is

$$\pi_k^{\text{ctrl}} = \alpha \, \pi_k + (1 - \alpha) \, \frac{1}{K}, \qquad 0 \leq \alpha \leq 1, \tag{16}$$

so that even minority voices retain some influence over action selection.

Policy posterior under polyphonic EFE.

$$p(\pi) \propto \exp(-\beta_{\text{eff}} \, G_{\text{poly}}(\pi)) . \tag{17}$$

This retains the standard Active Inference softmax form, but replaces G with G_{poly}.

Precision as diplomacy (adaptive action precision). In polyphonic settings, commitment can be controlled by precision as a function of cross-voice agreement:

$$\beta_{\text{eff}}(t) = \beta_0\, f(\mathcal{A}(t))\,, \tag{18}$$

where $\mathcal{A}(t)$ is an agreement (or coalition) index and $f(\cdot)$ is monotone increasing. One simple choice uses the variance of predicted progress across voices:

$$\mathcal{A}(t) = -\frac{1}{|\Pi|}\sum_{\pi\in\Pi} \text{Var}_k\big(r_k(\pi,t)\big), \qquad \beta_{\text{eff}}(t) = \text{clip}\big(\beta_0 \exp\big(\kappa\,\mathcal{A}(t)\big), \beta_{\min}, \beta_{\max}\big)\,, \tag{19}$$

where Π is the candidate policy set and clip bounds precision.

Updating credence weights from model evidence (without collapse). Instead of hard model selection, credence can be updated via a soft evidence accumulator:

$$\ell_k(t) = \rho\,\ell_k(t-1) - F_k(t), \qquad \pi_k(t) = \epsilon\frac{1}{K} + (1-\epsilon)\,\frac{\exp(\gamma\,\ell_k(t))}{\sum_{j=1}^{K}\exp(\gamma\,\ell_j(t))}, \tag{20}$$

where $\ell_k(t)$ is a leaky log-evidence proxy, $\rho \in (0,1)$ sets the timescale, γ controls sharpness, and ϵ enforces a floor (pluralism guarantee).

Viability. Let $\mathcal{V}$ denote a viability set over internal and external states (e.g., physical constraints, bounded energy, bounded uncertainty, bounded coupling). Polyphonic control can be cast as maintaining viability by modulating coupling and precision:

$$\mathcal{V} = \{x:\ g_r(x) \le 0\ \forall r\}, \qquad \dot{\lambda}_{ij} = h_{ij}(\text{slack}(x))\,, \qquad \dot{\beta}_{\text{eff}} = u(\text{slack}(x), \mathcal{A}(t))\,, \tag{21}$$

so that the system becomes more decisive (higher coupling/precision) near constraint boundaries and more plural/exploratory (lower coupling/precision) when safely within $\mathcal{V}$.

Reduction to conventional Active Inference. The polyphonic formulation reduces to standard Active Inference under any of the following conditions:

- $K = 1$ (a single generative model),
- $\lambda_{ij} = 0$ for all i, j (no coupling),
- $\pi_k = \pi_k^{\text{ctrl}} = \delta_{k,k^*}$ (hard model selection),
- or $\beta_{\text{eff}} = \beta_0$ is fixed and independent of inter-voice agreement.

In these limits, $\mathcal{F}_{\text{poly}} \to F$ and $G_{\text{poly}}(\pi) \to G(\pi)$, recovering the standard Active Inference scheme.

Interpretational summary. Polyphonic Active Inference preserves the normative objective of free energy minimisation, while relaxing the organisational assumption that inference and control must be governed by a single dominant generative model. Multiple models remain concurrently viable, coordination is achieved through soft alignment rather than elimination, and commitment is regulated by adaptive precision. In this sense, polyphony specifies a mode of inference–control organisation compatible with Active Inference, rather than an alternative objective function.