

Neural Policy Composition from Free Energy Minimization

Francesca Rossi ^{1,*}

Veronica Centorrino ^{2,*}

Francesco Bullo ^{3,†}

Giovanni Russo ^{4,†} ✉

December 5, 2025

Abstract. The ability to compose acquired skills to plan and execute behaviors is a hallmark of natural intelligence. Yet, despite remarkable cross-disciplinary efforts, a principled account of how task structure shapes gating and how such computations could be delivered in neural circuits, remains elusive. Here we introduce GateMod, an interpretable theoretically grounded computational model linking the emergence of gating to the underlying decision-making task, and to a neural circuit architecture. We first develop GateFrame, a normative framework casting policy gating into the minimization of the free energy. This framework, relating gating rules to task, applies broadly across neuroscience, cognitive and computational sciences. We then derive GateFlow, a continuous-time energy based dynamics that provably converges to GateFrame optimal solution. Convergence, exponential and global, follows from a contractivity property that also yields robustness and other desirable properties. Finally, we derive a neural circuit from GateFlow, GateNet. This is a soft-competitive recurrent circuit whose components perform local and contextual computations consistent with known dendritic and neural processing motifs. We evaluate GateMod across two different settings: collective behaviors in multi-agent systems and human decision-making in multi-armed bandits. In all settings, GateMod provides interpretable mechanistic explanations of gating and quantitatively matches or outperforms established models. GateMod offers a unifying framework for neural policy gating, linking task objectives, dynamical computation, and circuit-level mechanisms. It provides a framework to understand gating in natural agents beyond current explanations and to equip machines with this ability.

Introduction

Humans and other animals can dynamically combine previously acquired skills to plan and execute complex behaviors [98]. This ability – widely regarded as a hallmark of natural intelligence – is crucial to survival and flexible problem solving [52, 64]. Yet, understanding how the brain implements this capability [33, 83, 98] and, consequently, how it could inspire computational models for autonomous decision-making agents [77, 64, 103] is a central challenge with broad implications across neuroscience, engineering and artificial intelligence (AI).

In neuroscience, a growing body of experimental evidence suggests that the prefrontal cortex (PFC) may play a key role in synthesizing complex decision policies for a given task by combining behavioral schemas. This process may be implemented via a gating mechanism that regulates information flow

¹ Scuola Superiore Meridionale, Italy. ² ETH, Zürich. ³ Center for Control, Dynamical Systems, and Computation, UC Santa Barbara, CA, USA. ⁴ Department of Information and Electrical Engineering and Applied Mathematics, University of Salerno, Italy. ^{*,†} These authors contributed equally. ✉ e-mail: giovarusso@unisa.it

across brain circuits [81, 45]. Theoretical work further suggests that compositional mechanisms can be modeled via architectures that combine Recurrent Neural Networks (RNNs) and mixture-of-experts (MoE) frameworks [104, 98, 48]. Here, a gating network – often implementing a softmax rule – modulates the use of the appropriate schemas/skills (i.e., the experts) based on the environment inputs and the underlying task. While impressive, these approaches yield insights that are often specific to the data, the task, and the network architecture considered.

To both advance a more general understanding of knowledge composition mechanisms and engineer such principles in autonomous agents, researchers have increasingly turned to robotics as a testbed for designing, validating, and benchmarking gating-based computational models, often using motor control tasks as reference problems [77]. Given a task, behavioral schemas are associated to primitives (reusable policies) that are combined into a single policy, with weights assigned by a gating mechanism. This approach has inspired the design of layered architectures such as MOSAIC and Hammer [107, 27]. In these architectures, the output of fast, specialized, controllers is linearly combined by a slower, more flexible, mechanism. Weights are determined by a gating rule – again, a softmax. This gating rule is also central to sensorimotor control schemes based on the MoE framework [43], such as the one in [95]. Beyond robotics, the MoE framework [60, 108, 46] has become a cornerstone of modern AI systems, including Large Language Models (LLMs) and in-context decision-making methods [16, 67]. Compared with the neuroscience literature, these advances focus primarily on computational principles, typically implementing them in artificial networks. Two main architectures to combine expert outputs are [16, 60]: (i) dense, fully activated MoE, where all experts contribute to the final output, typically via softmax gating [43]; (ii) sparse, selective MoE, where only a subset of experts is chosen, typically determined via an argmax-based selection [90, 30] or Gumbel-softmax [44, 68]. In brief, the gating rule, a core determinant of the performance, is selected by the network designer rather than emerging from the properties of the task.

Despite remarkable cross-disciplinary efforts, most explanations remain tied to specific network architectures, tasks, and datasets. As a result, a principled account of how the task shapes gating computations and, in turn, how such computations drive the organization of the neural circuits that implement them, remains elusive. What appears to be missing is a theoretically grounded and interpretable computational model that applies broadly across neuroscience, cognitive science, and machine learning. Such a model should provide a unifying account of gating and integrate naturally with existing conceptual frameworks. At present, it remains unclear what general and broadly applicable objective a given gating rule is optimizing, nor how its functional requirements are instantiated mechanistically in neural circuits.

To address this gap, we develop GateMod, a computational model grounded into the minimization of the free energy. GateMod yields a quantitative characterization of gating as an energy model. It casts gating within a variational, normative, formulation that explicitly relates task to gating mechanism to the energy landscape. It further specifies a neural circuit capable of implementing these computations, making the role of each circuit element interpretable in view of the task. In doing so, GateMod yields several implications, including highlighting the central role of in-context computation and suggests that dendritic processing may be a key biological substrate supporting gating mechanisms in natural circuits.

GateMod consists of three key components. The first is GateFrame, a normative framework for primitives gating formalized via an optimization problem. In GateFrame, the policy is computed by combining a set of primitives, e.g., schemas, skills for natural agents, or reusable sensorimotor controllers for artificial agents. The decision variables are the gating weights of the primitives. These weights are determined by minimizing a cost functional that balances a statistical complexity term and an entropy regularizer.

Remarkably, GateFrame can be cast into the minimization of the (expected) free energy, a unifying account across neuroscience [32, 61], artificial intelligence [91], and control [89, 86]. As such, GateFrame objective subsumes as special cases a broad range of decision-making frameworks, such as maximum entropy [11], broadly applicable across neuroscience, cognitive science and machine learning.

The second component of GateMod is GateFlow, a continuous-time dynamical system to provably find the optimal solution of GateFrame. GateFlow is an energy model featuring two distinct energy functions. We rigorously demonstrate that the unique equilibrium of GateFlow is also the optimal solution of GateFrame. Then, we prove that GateFrame exhibits highly ordered transient and asymptotic behaviors, ensuring convergence to the equilibrium regardless of initial conditions. More precisely, GateFlow benefits from a contrativity property, implying that the distance between any trajectory of its trajectories exponentially shrinks in time. This property also yields explicit convergence rates together with robustness and other desirable properties [57, 14]. In deriving these results, we show that GateFlow belongs to a class of dynamical systems – which we term as softmax flows – that can tackle a broad class of entropy maximization problems. Our characterize in full softmax flows.

The final component of GateMod is a neural circuit implementing GateFlow, which we term GateNet. Its architecture is transcribed directly from GateFrame, making the role of each element of the circuit clear and its computations interpretable. Remarkably, this top-down approach yields a soft-recurrent neural circuit with contextual computations [59] that can be implemented via dendritic computations [65, 62]. Moreover, we also show that our circuit not only can implement both dense and sparse gating rules, but it also features non-negative neuronal variables that can naturally be interpreted as firing rates.

To evaluate our model and its implications, we consider two applications across different domains. First, we examine flocking, a well-studied phenomenon in both nature and technology [93, 7, 102, 73, 8, 100, 6]. Here, agents follow social forces/primitives that are well documented and established in the literature [78, 23]. We show that GateMod, building on known social forces, can dynamically modulate their use, recovering classic collective behaviors such as polarization, plasticity, and leader-guided, soft-controlled, navigation [39, 105, 40, 28]. In this last setting, GateMod successfully steers emergent group dynamics, highlighting its potential for applications such as coordinated migrations and cooperative tasks [34, 35, 51, 18, 71]. Second, we evaluate GateMod ability to interpret decision-making behaviors in humans involved in multi-armed bandit tasks. By evaluating our model on a publicly available dataset, and benchmarking it with excellent methods, we show that GateMod not only better explains the data, but it also allows to gain insights that remain elusive for related excellent methods.

GateMod is the first computational model revealing how gating relates to agent tasks and to the underlying neural circuit to implement this functionality. Unlike other models, where the gating function is often imposed by the designer, in GateMod gating – a softmax – emerges from an optimization problem that involves minimizing the free energy. The flow that we introduce to solve the minimization, not only provably finds the optimal gating weights and exhibits desirable dynamic properties, but it also admits a neural implementation. GateMod establishes a framework to both empower the design of policy composition schemes in artificial agents and understand policy composition in natural behaviors beyond current explanations. Despite the success of state-of-the-art frameworks, there is no general theory mechanistically relating gating, to agent task, to a neural circuit. GateMod provides these explanations.




Figure 1: GateMod Set-up. **A** At time step $k - 1$, an agent (e.g., a boid in a flock, or a person in a multi-armed bandit task, or an autonomous agent) receives the state $\mathbf{x}_{k-1}$ from the environment and determines action $\mathbf{u}_k$. Both $\mathbf{x}_{k-1}$ and $\mathbf{u}_k$ are realizations of random variables, $\mathbf{X}_{k-1}$ and $\mathbf{U}_k$. We denote random variables with upper-case letters and their realizations with lower-case letters. Bold means that the variable is, in general, a vector. **B** At each time step, the agent computes the optimal policy $p_{\mathbf{u}}^*(\mathbf{u}_k | \mathbf{x}_{k-1})$ by combining a set of available primitives $\pi^1(\mathbf{u}_k | \mathbf{x}_{k-1}), \dots, \pi^{n_\pi}(\mathbf{u}_k | \mathbf{x}_{k-1})$ via a gating mechanism. GateMod provides a normative framework (GateFrame) to optimally combine the weights, a continuous-time dynamics (GateFlow) to provably find the weights, and a neural circuit (GateNet) implementing this continuous-time solver. Intuitively, given a task – formalized via a generative model – and a model of the environment, GateMod computes the weights, $\mathbf{w}_k^*$, to linearly combine the primitives.

GateMod Set-up

Fig. 1A illustrates an agent interacting with a stochastic environment. At each time-step, the agent determines an action $\mathbf{u}_k$ based on the current state $\mathbf{x}_{k-1}$. The environment transition kernel $p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ transitions from state $\mathbf{x}_{k-1}$ to $\mathbf{x}_k$ in response to $\mathbf{u}_k$. The action is sampled from the agent stochastic policy $p(\mathbf{u}_k | \mathbf{x}_{k-1})$ so that the (closed-loop) agent-environment dynamics is described by the joint probability $p(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1})$. For example, as in our first application example, the agent could represent a boid in a flock, with state defined by its position and velocity and action by its steering force. Other examples include the agent being a person in a multi-armed band task (as in our second application) with state associated to the history of rewards received by each arm.

Given a task, GateMod computes the optimal policy $p_{\mathbf{u}}^*(\mathbf{u}_k | \mathbf{x}_{k-1})$ arising from a gating mechanism that linearly combines a set of n_π primitives (Fig. 1B). Each primitive, $\pi^i(\mathbf{u}_k | \mathbf{x}_{k-1})$, is a reusable randomized policy with support spanning the full action space. For boids in a flock, primitives may be mapped onto social forces, while for a person in a multi-armed bandit task these could be behavioral schemas or, for an autonomous agent, previously acquired policies.

The optimal n_π -dimensional weight vector combining primitives is $\mathbf{w}_k^*$. The weights (see Results) are obtained by solving an optimization problem that minimizes the statistical complexity (the discrepancy) between the agent-environment dynamics and a reference probability $q(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1})$. We use the wording *generative model* to denote this probability, using this term broadly: in GateMod, it may represent a time-series model, specify a target behavior, encode a cost function, or combine these elements. For a boid in a flock, the generative model could be a time-series model describing the aggregate trajectory [42] of the neighbors. For a multi-armed bandit task, the model may capture predictions about uncertainty and value inferred from experimental data [36]. For an autonomous agent, the model may incorporate the task cost through an exponential kernel [38, 10, 89].

We next present our main results. First, we introduce a broadly applicable normative approach, cast as an optimization problem, for composing primitives. Second, we show that the optimal weights can be computed via a continuous-time dynamics with guaranteed convergence. Third, we derive a neural circuit that inherits these properties. Finally, we evaluate our model on two different domains.




Figure 2: GateMod. **A** GateFrame normative framework. At each time step, the agent computes optimal policy weights $\mathbf{w}_k^*$ by solving an entropy-regularized optimization problem that minimizes a trade-off between statistical complexity and entropy. The constraints formalize the fact that the resulting policy is a linear, and hence convex, combination of primitives. The optimal weights correspond to the equilibrium of GateFlow: a continuous-time dynamical system defined by a softmax gradient flow. This is an energy model that provably converges to GateFrame optimal solution. **B** GateFlow is an energy model featuring highly ordered behaviors with guaranteed and explicit exponential convergence rate to the optimal solution. GateFrame objective is an energy function for GateFlow so that the energy decreases along its trajectories towards the optimal solution (top). Convergence is global (bottom): regardless of the initial conditions (initialization value for the weights) GateFlow trajectories converge to $\mathbf{w}_k^*$. Convergence follows from a stronger contractivity property that also confers robustness and other desirable properties. **C** GateFlow admits a neural implementation. The architecture consists of two coupled modules operating at different timescales: a fast subsystem that computes the gradient of the objective using local operations (linear summation, logarithmic activation) and a slower subsystem that implements the softmax activation function featuring exponential and logarithmic activation functions. The fast unit features contextual computations that are based on the current state. These computations can be implemented via the Sigma-Pi model. The input to the fast unit is a cost combining a mismatch from the generative model and a log-likelihood. The result, aligned with literature on the distributional costs in biological neural circuits, is a vector associated to the action space rather than a single mean value.

Results

GateFrame Optimization

Within GateMod, GateFrame provides the top-down normative framework for primitives gating, formalized as an optimization problem. The decision variables (Fig. 2A, top) are the weights associated, at each time-step, to each of the primitives. The constraints enforce that the policy is a linear mixture of primitives. The cost functional consists of two terms. The first is the statistical complexity between the agent-environment dynamics and the generative model; minimizing this term aligns $p(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1})$ with $q(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1})$, which is therefore a bias. The second term is an entropic regularizer, scaled by a temperature-like parameter ε . As a result, GateFrame maximizes a reward (negative of the complexity) regularized with a widely-adopted entropy maximization term, see, e.g., [38, 11, 106, 29].

Statistical complexity is defined as the Kullback-Leibler (KL) divergence between $p(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}) = p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) p(\mathbf{u}_k | \mathbf{x}_{k-1})$ and $q(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}) = q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) q(\mathbf{u}_k | \mathbf{x}_{k-1})$. The resulting problem is

$$\begin{aligned} \min_{\mathbf{w}_k \in \Delta_{n_\pi}} & \overbrace{D_{\text{KL}}(p(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}) || q(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}))}^{\text{Statistical Complexity}} - \varepsilon \overbrace{\mathcal{H}(\mathbf{w}_k)}^{\text{Entropy}} \\ \text{s.t. } & p(\mathbf{u}_k | \mathbf{x}_{k-1}) = \underbrace{\sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})}_{\text{Mixture-of-primitives}}. \end{aligned} \quad (1)$$

The weights vector $\mathbf{w}_k$ belongs to the probability simplex Δ_{n_π} and $\mathcal{H}(\cdot)$ is its entropy. Embedding the constraint in the cost (see Methods for details) renders the dependency on the decision variables explicit and reveals that the optimization problem is convex. To stress the dependency of the complexity on the decision variables we use the notation $\mathbf{F}(\mathbf{w}_k)$. The expression of the cost obtained after substituting the mixture-of-primitives constraint into the objective of Eq. [1] can then be written compactly as $\mathbf{F}(\mathbf{w}_k) - \varepsilon \mathcal{H}(\mathbf{w}_k)$. This cost is convex in $\mathbf{w}_k$ even when the environment, generative model, and primitives are nonlinear and nonstationary.

We now show that GateFrame formulates a free energy minimization problem. The minimization of the KL divergence naturally arises in the context of variational inference, where the second probability is a posterior to be approximated, see, e.g., [69, Chapter 10] and references therein. Beyond inference, the link with free energy minimization becomes apparent when we consider the choice of $q(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1})$ as $\frac{1}{Z} \tilde{q}(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}) \exp(-c(\mathbf{x}_k, \mathbf{u}_k))$. Here, Z is a normalization constant, $\tilde{q}(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1})$ is a probability, and $c(\mathbf{x}_k, \mathbf{u}_k)$ is a state/action cost. In this case, the complexity term in Eq. [1] becomes

$$D_{\text{KL}}\left(p(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}) || \frac{\tilde{q}(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}) \exp(-c(\mathbf{x}_k, \mathbf{u}_k))}{Z}\right). \quad (2)$$

Then, applying the logarithm product rule, Eq. [2] yields

$$\begin{aligned} & D_{\text{KL}}(p(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}) || \tilde{q}(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1})) \\ & + \mathbb{E}_{p(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1})}[c(\mathbf{x}_k, \mathbf{u}_k)] + \ln Z, \end{aligned} \quad (3)$$

where the second term is the expected cost and the last term, being constant, does not affect the minimization and can be dropped. This is however the expected free energy objective of [89], thereby connecting

GateFrame to active inference [75] and, in turn, to KL control, control as inference [96, 13, 47] and maximum diffusion (MaxDiff) reinforcement learning [11]. This is notable because MaxDiff generalizes the MaxEnt approach [11, 109] and inherits its desirable properties.

Despite its simple formulation, GateFrame provides an overarching framework that embeds a broad class of decision-making problems within a mixture-of-primitives setting. Notably, the optimal policy in GateFrame is obtained by combining primitives. The entropic regularizer is the key theoretical ingredient underlying the emergence of softmax gating mechanisms. Specifically, the entropy term enforces that the optimal weights are determined through a softmax gating mechanism, whereas models without this term (i.e., with $\varepsilon = 0$) lack this desirable structure [82]. Next, we derive the continuous-time dynamical system solver of GateFrame (GateFlow) and its neural implementation (GateNet).

GateFlow Transcribes GateFrame Optimization

We first derive the equations constituting GateFlow (Fig. 2A, bottom). Then, we investigate their structure, revealing two key highlights. First, GateFrame yields a softmax weights selection rule and it encodes the optimal solution of GateFrame via its equilibrium. Second, the agent task – encoded through the statistical complexity term in GateFrame – determines *how much* each primitive contributes to the optimal policy. Derivations are unpacked in Methods.

To establish the results, we reformulate the constrained optimization in Eq. [1] into the unconstrained problem

$$\min_{\mathbf{w}_k \in \mathbb{R}^{n_\pi}} F(\mathbf{w}_k) + \varepsilon H_{\text{barrier}}(\mathbf{w}_k). \quad (4)$$

Here, $F(\mathbf{w}_k)$ is defined as in the previous section and $H_{\text{barrier}}(\mathbf{w}_k)$ is $-\mathcal{H}(\mathbf{w}_k) + \iota_{\Delta_{n_\pi}}$, with $\iota_{\Delta_{n_\pi}}$ being the indicator function (that is, $\iota_{\Delta_{n_\pi}}$ is equal to 0 if $\mathbf{w}_k$ belong to the simplex Δ_{n_π} , and to $+\infty$ otherwise). We term the map $H_{\text{barrier}}(\mathbf{w}_k)$ as *entropic barrier*, since it equals $-\mathcal{H}(\mathbf{w}_k)$ if the decision variables belong to the simplex and $+\infty$ otherwise.

The reformulation in Eq. [4] yields an unconstrained composite optimization problem whose regularity properties enable the design of a continuous-time proximal-gradient dynamics that solves this problem [2, 41, 19, 25]. A key property of this dynamics is that a vector $\mathbf{w}_k^*$ is an optimal solution of the problem in Eq. [4] if and only if it is also an equilibrium of the dynamics. Deriving the continuous-time proximal-gradient dynamics associated to Eq. [4] yields GateFlow:

$$\tau \dot{\mathbf{w}}_k = -\mathbf{w}_k + \text{softmax}(-\varepsilon^{-1} \nabla F(\mathbf{w}_k)). \quad (5)$$

In Eq. [5], $\tau > 0$ is a time-scale parameter and $\nabla F(\mathbf{w}_k)$ is the gradient of $F(\mathbf{w}_k)$ with respect to the decision variables. In Eq. [5], each entry of the state variables – the weights vector – is updated according to the corresponding component of the softmax applied to $-\varepsilon^{-1} \nabla F(\mathbf{w}_k)$. Following this iterative rule, given an initial condition (an initial guess for the weights) $\mathbf{w}_k(0)$, GateFlow iteratively updates the weights. The update at a generic time t is $\mathbf{w}_k(t)$.

GateFlow dynamics is a proximal flow dynamics. For an optimization problem of the form $\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) + g(\mathbf{x})$, where f is convex and g is a (possibly, poorly behaved) regularizer, the proximal flow is $\dot{\mathbf{x}} = -\mathbf{x} + \text{prox}_g(\mathbf{x} - \nabla_x f(\mathbf{x}))$. Just like gradient descent $\dot{\mathbf{x}} = -\nabla_x f$ determined by the energy f , proximal gradient descent is determined by the energy $f + g$. Details in Methods. In GateFlow, the softmax emerges from the proximal operator of the entropic term in Eq. [4]. This enforces the optimal weights to be selected according to a softmax gating rule. In fact, GateFlow equilibrium point (and hence GateFrame optimal so-

lution) $\mathbf{w}_k^*$ must satisfy the equation $\mathbf{w}_k^* = \text{softmax}(-\varepsilon^{-1}\nabla F(\mathbf{w}_k^*))$. This means that the optimal weights, forming a probability vector, are selected according to a softmax rule. As ε decreases, the softmax recovers the argmax arising in, e.g., selective-type MoE architectures. In the Supplementary Information, we show that GateFlow can also recover Gumbel-softmax gating typical of dense MoE architectures and this corresponds to embedding a bias in GateFrame. Moreover, GateFlow elucidates the mechanistic role of the task onto weights selection. The expression for the optimal weights shows that each weight is adjusted according to the scaled negative gradient of the statistical complexity, this is the task-encoding term in GateFrame. Each weight decreases as the corresponding gradient coordinate increases. Consequently, as the weights are the coefficients (Fig. 1B) to combine primitives, a primitive contributes less to the optimal policy as its gradient coordinate increases, therefore allowing to interpret the output of our model.

GateFlow is the continuous dynamical system encoding the optimal solution of GateFrame as its equilibrium. It yields a gating mechanism where the softmax gating rule emerges from the structure of the underlying optimization. Next, we show that not only GateFlow exhibits these features, but it also enjoys highly desirable convergence properties. Crucially, for any initial guess of the weights, GateFlow dynamics converges to the GateFrame optimal solution.

GateFlow is an Energy Model and Provably Converges to GateFrame Optimal Solution

To solve GateFrame optimization, GateFlow not only needs to encode the optimal solution via its equilibrium but it also needs to guarantee convergence to the equilibrium starting from an initial guess for the weights. Not only this is the case, but GateFlow also features highly ordered behaviors robustly converging to $\mathbf{w}_k^*$ from any initial feasible guess with guaranteed convergence rate. Our results also show that GateFlow is an energy model (Fig. 2B) and we highlight two distinct energy functions.

The first energy function is GateFrame cost and this is revealed by the structure of the optimization. Since GateFlow is the proximal-gradient dynamics associated to GateFrame and F is at least convex in the decision variables, from the theory of proximal operators it follows that the non-negative GateFrame cost decreases along the trajectories [37]. More precisely, $F(\mathbf{w}_k(t)) - \varepsilon H(\mathbf{w}_k(t))$ decreases over time to $\mathbf{w}_k^*$, the energy minimum (Fig. 2B). The second energy function arises from GateFlow convergence properties.

Convergence of GateFlow trajectories to $\mathbf{w}_k^*$ is global and exponential. Global means that GateFlow trajectories converge to the optimal solution for any initial guess that belongs to GateFrame feasibility domain Δ_{n_π} . Exponential means that the distance between $\mathbf{w}_k(t)$ and $\mathbf{w}_k^*$ shrinks exponentially in time. More precisely, given any initial condition $\mathbf{w}_k(0)$ belonging to Δ_{n_π} , $\mathbf{w}_k(t)$ always belongs to Δ_{n_π} . This means that GateFlow trajectories always belong to the GateFrame feasibility domain if they are initialized from valid initial conditions. This invariance of the simplex is particularly important, as it guarantees that GateFlow always returns a valid probability distribution. The invariance property is also crucial to prove that for any trajectory starting in the simplex it holds that

$$\|\mathbf{w}_k(t) - \mathbf{w}_k^*\| \leq e^{-t/\tau} \|\mathbf{w}_k(0) - \mathbf{w}_k^*\|, \quad (6)$$

where $\|\cdot\|$ is the Euclidean norm. Eq. [6] brings two key implications. First, the Euclidean distance between any solution of GateFlow and $\mathbf{w}_k^*$ shrinks exponentially with rate $1/\tau$. Second, $\frac{1}{2} \|\mathbf{w}_k(t) - \mathbf{w}_k^*\|^2$ is an additional monotonically decreasing energy function for GateFlow. The convergence property in Eq. [6] follows from a stronger contractivity property [57, 14] that also implies uniqueness of $\mathbf{w}_k^*$ and other desirable dynamic properties. See Methods for details.

Collectively, these results establish GateFlow as a continuous-time energy model solving GateFrame op-

timization. It benefits from explicit convergence rates and highly ordered dynamics. Moreover, it admits a neural circuit implementation. This is derived next.

GateNet Neural Circuit

GateFlow admits a neural implementation, GateNet (Fig. 2C). This is a soft-competitive continuous-time recurrent neural circuit with contextual computations [59]. Here, we show that each element of the architecture is derived from GateFlow. Consequently, each element of the circuit is directly related to GateFrame, making its role clear and computations interpretable.

Given the current state of the agent, $\mathbf{x}_{k-1}$, GateNet returns $\mathbf{w}_k^*$, the equilibrium of GateFlow (hence the optimal solution of GateFrame). GateNet consists of two units (Fig. 2C): a fast unit, computing the exponent of the softmax in Eq. [5] and a slow unit that implements the softmax itself.

To derive the fast unit, in a setting with d_u discrete actions, computing the softmax exponent in Eq. [5] reduces to evaluate the quantity $\mathbf{y}_k$ defined as

$$-\varepsilon^{-1} \Pi(\mathbf{x}_{k-1})^\top (\ln(\Pi(\mathbf{x}_{k-1}) \mathbf{w}_k) + \mathbf{c}(\mathbf{x}_{k-1}, \mathbf{u}_k)). \quad (7)$$

Here, given the agent state $\mathbf{x}_{k-1}$, $\Pi(\mathbf{x}_{k-1})$ is the $d_u \times n_\pi$ dimensional matrix having on its i -th column the i -th primitive $\pi^i(\mathbf{u}_k | \mathbf{x}_{k-1})$, $\mathbf{c}(\mathbf{x}_{k-1}, \mathbf{u}_k)$ is the d_u -dimensional vector containing the cost associated to each of the actions; the logarithm in the expression is component-wise. The term $\mathbf{c}(\mathbf{x}_{k-1}, \mathbf{u}_k)$ is the input to GateNet. This cost combines the statistical complexity between $p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ and $q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ with a log-likelihood. More precisely, $\mathbf{c}(\mathbf{x}_{k-1}, \mathbf{u}_k)$ is $D_{\text{KL}}(p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)) - \ln q(\mathbf{u}_k | \mathbf{x}_{k-1})$. GateNet suggests that, given the current state, in order to perform the requested computations, the neural circuit must receive as input a distributional cost – a cost associated to the action space – and not just a mean value. Remarkably, a growing body of experimental and computational evidence is suggesting that distributional costs may be used in brain circuits, with the thalamus and striatum playing a key role in this computation [58, 21]

The goal of the fast unit is therefore to output $\mathbf{y}_k$. Denoting $\mathbf{c}(\mathbf{x}_{k-1}, \mathbf{u}_k)$ by $\mathbf{c}_k$, this is obtained with the dynamics (in vector form)

$$\begin{aligned} \tau_g \dot{\mathbf{a}} &= -\mathbf{a} + \Pi(\mathbf{x}_{k-1}) \mathbf{w}_k \\ \tau_g \dot{\mathbf{b}} &= -\mathbf{b} + \ln(\mathbf{a}) \\ \tilde{\tau}_g \dot{\mathbf{y}} &= -\varepsilon \mathbf{y} - \Pi(\mathbf{x}_{k-1})^\top (\mathbf{b} + \mathbf{c}_k), \end{aligned} \quad (8)$$

where $\tau_g \ll \tilde{\tau}_g$ are time-scale constants. The first dynamics corresponds to the input layer in Fig. 2C. In the figure, following Eq. [8], the i -th input neuron receives the projected scalar quantity $\Pi_i(\mathbf{x}_k) \mathbf{w}_k$, where $\Pi_i(\mathbf{x}_k)$ is the i -th row of $\Pi(\mathbf{x}_k)$. The globally exponentially stable input dynamics converges to the equilibrium $\bar{\mathbf{a}} = \Pi(\mathbf{x}_{k-1}) \mathbf{w}_k$ and neural variables from this layer feed the second dynamics. This dynamics corresponds to the intermediate/hidden layer in Fig. 2C, featuring a logarithmic activation function. This dynamics, also globally exponentially stable, converges to the equilibrium $\bar{\mathbf{b}} = \ln(\mathbf{a})$. The hidden neural variables feed the last dynamics. This corresponds to the output of the fast unit. Since the input and hidden layers are faster than the output layer due to time-scale separation, the output variables globally exponentially converge to $\bar{\mathbf{y}} = -\varepsilon^{-1} \Pi(\mathbf{x}_{k-1})^\top (\ln(\Pi(\mathbf{x}_{k-1}) \mathbf{w}_k) + \mathbf{c}_k)$. This in fact coincides with $\mathbf{y}_k$ from Eq. [7], the desired network output.

The output neural variable, $\bar{\mathbf{y}}$, depends on the current state of the agent, $\mathbf{x}_{k-1}$, and to stress this aspect we left the dependency of Π on $\mathbf{x}_{k-1}$. This highlights that the computations carried out by GateNet fast

unit are contextual and depend on the agent state. In natural neural circuits, these contextual computations might be implemented at the dendritic level, for instance through Sigma-Pi-type computations posited in cortical circuit models [31, 84, 62, 63, 76, 56, 50]. In artificial neural networks, these computations can be implemented via neuromorphic circuits [49].

The output from the fast unit feeds the slow unit, which implements the softmax rule in GateFlow Eq. [5] and returns the optimal weights. The circuit (Fig. 2C) leverages the neural implementation of the softmax available in the literature [92]. In the same work it is also noted that this implementation has elements of biological plausibility in that – as for the fast unit – each neuron computes its output based only on local information. The dynamics for the slow unit circuit is

$$\begin{aligned}\tau_s \dot{m} &= -m + \sum_{\alpha=1}^{n_\pi} e^{y^\alpha} \\ \tau_s \dot{\mathbf{r}} &= -\mathbf{r} + \mathbf{y} - \mathbb{1}_{n_\pi} \ln(m) \\ \tau \dot{\mathbf{w}}_k &= -\mathbf{w}_k + \mathbf{e}^{\mathbf{r}},\end{aligned}\tag{9}$$

with $\tilde{\tau}_g \ll \tau_s \ll \tau$ ensuring that the first two equations are slower than the fast unit output dynamics. Given the vector $\mathbf{y}$ from the fast unit, having components y^α , neuron m stores the quantity $\sum_{\alpha=1}^{n_\pi} e^{y^\alpha}$, the normalizer in the softmax. This is also the equilibrium of the first globally exponentially stable dynamics. The output of neuron m is received, together with y^α , by neuron r^α , $\alpha = 1, \dots, n_\pi$. This second dynamics is again globally exponentially stable with equilibrium $\mathbf{y} = \mathbb{1}_{n_\pi} \ln(m)$. The last dynamics coincides with the output of GateNet and returns the component-wise exponential of this last quantity. At steady state, this vector corresponds to $\text{softmax}(\bar{y})$, which in turn gives the optimal weights for primitive composition, that is, the optimal solution of GateFrame.

Together, Eq. [8] and Eq. [9] define a recurrent neural dynamics that implements the complete GateFlow dynamics in Eq. [5]. As such, GateNet inherits all the desirable dynamic and convergence properties of GateFlow. Being GateFlow a positive system, GateNet neural variables also remain positive if initialized from positive values. This means that neural variables in GateNet can be interpreted as firing rates. Each element of the network is derived from GateFlow, which in turn is directly linked to GateFrame and hence to the agent task. This makes GateNet computations and the role of each of its neurons interpretable. Next, we evaluate our full computational model on two different domains.

GateMod Evaluation

To evaluate GateMod, we specifically consider a set of experiments from two different application domains, strategically selected to ensure that the effects of our model could be identified, benchmarked against the literature, and measured quantitatively. First, we consider a collective behavior set-up where primitives are associated to the popular notion of social forces. Experiments show that GateMod recovers well-known emerging flocking behaviors, making them interpretable and gaining insights on the dynamics use of these forces. Then, analyzing data collected from people involved in two multi-armed bandits tasks, GateMod shows that behaviors can be interpreted as a combination of simple behavioral primitives. Our model also reveals insights on how these behaviors are used in different tasks.

GateMod Yields Collective Behavior

Coordinated motions in animal groups, ranging from bird flocks to fish schools and insect swarms [87], have long inspired efforts to deepen our understanding of both natural and engineered systems. A central hypothesis in popular frameworks is that collective behaviors in nature emerge from agent-level interaction rules, frequently expressed in terms of social forces [79, 23, 101, 17]. Yet, an open challenge is to establish a framework that links the agent goal, to how it should dynamically combine social forces, to a neural circuit capable of delivering the required computations. GateMod can provide such a framework. Our model not only recovers widely recognized collective behaviors, but it also offers mechanistic insights into how agents in a group dynamically modulate their social forces.

In classical models, a boid (Fig. 3) in a flock determines the acceleration based on its own state (position and velocity) and the state of the neighbors that lie within its field of view. Following popular models [24, 53, 23] the field of view is partitioned in three non-overlapping zones with each zone promoting a behavior through a social force. The innermost region is the separation zone, where a boid attempts to avoid collisions with neighbors in this area by applying a separation force. The alignment zone surrounds the separation zone. Through an alignment force, the agent attempts to achieve directional consensus with boids that are within this area. Finally, the outer zone is termed cohesion zone. Through a cohesion force, the boid is drawn towards the average position of the boids that fall in this area. Separation, alignment, and cohesion [79, 23, 101, 17] have become foundational in various modeling frameworks [3, 66, 5, 102, 17] with studies suggesting that these forces may emerge from surprise minimization [42].

In GateMod the boid is an agent; social forces are naturally mapped onto primitives, $\pi^\alpha(\mathbf{u}_k^i | \mathbf{x}_{k-1}^i)$, $\alpha = 1, 2, 3$, corresponding to separation, alignment and cohesion. The agent has available these social primitives to determine its own acceleration. Specifically, the agent samples from a policy that combines the primitives according to GateFrame optimization. Consequently, the policy from which the i -th boid samples its acceleration is

$$p_{\mathbf{u}}^{i,*}(\mathbf{u}_k^i | \mathbf{x}_{k-1}^i) = \sum_{\alpha=1}^3 \mathbf{w}_k^{i,*} \pi^\alpha(\mathbf{u}_k^i | \mathbf{x}_{k-1}^i). \quad (10)$$

The optimal weights are computed via GateFlow (Fig. 3B) and experiments results are evaluated through polarization – an order parameter capturing the mean heading of the group, similar in spirit to spin systems magnetization [54] – and distance from a goal position. Details of the settings are provided in Methods.

In the first set of experiments, the group of boids has no leaders and we first verify if GateMod can recover polarization, a well-documented phenomenon in the literature. We equip boids with a generative model inspired by [42]. For the i -th boid, $q(\mathbf{x}_k^i | \mathbf{x}_{k-1}^i, \mathbf{u}_k^i)$ is a Gaussian with mean depending on the average position and velocity of the boids within its cohesion and alignment zones, while $q(\mathbf{u}_k^i | \mathbf{x}_{k-1}^i)$ is uniform. Fig. 3C (left and middle) shows that not only GateMod recovers polarization, consistently with the literature [23, 42], but also reveals how weights evolve over time and hence how primitives are modulated (right). We observe that, under our model, weights gradually evolve over time so that, when polarization is achieved, these become approximately uniform, indicating a balanced use of the forces. More precisely, while the initial weights depend on the position of the boid in the flock, GateMod hints that there is a hidden organization principle, linking the global behavior of the flock to an equilibrium between the social forces. Fig. 3C is representative of this phenomenon (Fig. S1 in Supplementary Information shows the same diagram for all boids). Moreover, additional experiments in the Supplementary Information also show that, when the generative model is equipped with a collision-avoidance term, GateMod also recovers

milling behaviors, again consistently with the literature [42]. See Fig. S2 in Supplementary Information and Sec. 5 therein.

Experiments show that GateMod can recover well-established behaviors consistently with the literature. This conclusion is also supported when leaders are included in the flock. Leaders are boids informed of a goal destination and their generative model encodes a goal-directed behavior. Fig. 3D shows the emerging group behavior when a small fraction of boids is goal-informed and seeks to reach the goal position. Consistently with the literature [22, 12], the group achieves goal-directed flocking without loss of cohesion (left and middle). Moreover, GateMod reveals that the use of primitives is conditioned to the agent being goal-directed (right). More precisely, experiments suggest that the behavior of informed boids is more adaptive, as made apparent by the time evolution of the weights. This may be an indicator of a more flexible behavior of informed agents over followers – a behavior increasingly emphasized in the literature [55, 9, 71]. Fig. 3D is obtained with followers (non-informed) boids computing their actions using GateMod. Results are confirmed in an additional set of experiments (Fig. S3 and Sec. 5 in Supplementary Information) for different numbers of informed boids and the temperatures. Finally, to further evaluate the ability of our model to steer the behavior of the flock, in the Supplementary Information we conduct an ablation study where followers determine actions based on three different remarkable models from the literature. Even in these settings, the informed agents, equipped with GateMod, are able to soft-control [39, 105, 40, 28, 4] the flock towards the goal. See Sec. 5 in Supplementary Information.

In brief, experiments confirm that GateMod consistently recovers well-known behaviors in the literature and enables soft-control of flocks. In our model, boids are probabilistic decision makers, integrating surprise minimization with first principles. According to GateFrame, actions are sampled from a policy dynamically built from primitives. Not only this capability can be delivered by a neural circuit, but our model also makes the modulation of the primitives interpretable, enabling insights that – to the best of our knowledge – remain elusive for classic models.

GateMod as a Framework to Interpret Exploration/Exploitation Balance

Gaining a deeper understanding of how humans balance exploration and exploitation in uncertain tasks is a central theme across cognitive sciences and RL [94, 26, 97, 20], with multi-armed bandits tasks often used as benchmarks to develop, test and validate working hypotheses. In this context, a growing body of experimental evidence, with accompanying computational models, suggest that human decision-making is best understood not in terms of a single strategy but rather via a mixture of multiple mechanisms [36, 20]. This hypothesis implies that the policy adopted by humans in a multi-armed bandit task arises from combining simpler policies. Yet, explanations available in the literature rely on complex primitives in an attempt to capture the exploration/exploitation dilemma [36, 20]. A key challenge is to infer from data interpretable quantitative insights that explain choices in simpler terms, using policies that can be mapped onto behavioral (or mental) types. GateMod can provide the framework to gain these insights. Revisiting a classic dataset [36] we show that GateMod not only explains choices in terms of simple categorical behavioral types, but it also quantifies how much each primitive contributes to decision, and how their use changes across tasks.

In [36], decisions are best explained by inferring a linear combination of two families of algorithms based on uncertainty bonuses (Upper Confidence Bound, UCB) and sampling (Thompson). Data are collected from two web-based experiments with participants selecting one of two mutually excluding options (arms). Participants are instructed to maximize the total reward earned. Each experiment is organized into 20

blocks of 10 trials. In Experiment 1 (44 participants), one arm produces stochastic rewards with zero mean while the other always yields a constant reward of zero. In Experiment 2 (45 participants), both arms produce stochastic rewards, each with its own block-specific mean.

For these data, we evaluate if GateMod can provide explanation value beyond the hybrid model from [36] and the standard UCB, Thompson and Value algorithms. For the application of GateMod, at each trial k , the state $\mathbf{X}_k$ is the belief of experiment participants about the mean and variance of rewards associated with each arm. The generative model is the model from [36] itself and primitives encode simple behaviors. Namely: (i) exploitation, favoring the arm with the highest expected reward; (ii) uncertainty-seeking exploration, favoring the arm associated with the highest uncertainty; and (iii) risk aversion, favoring the least uncertain arm. In this way, at each trial, GateMod returns weights that approximate, with simple primitives, the generative model. The sequence of optimal weights captures how much each primitive contributes to the observed choices, thus making the model interpretable. To quantify and benchmark the explanatory value of GateMod we use the popular Protected Exceedance Probabilities (PXP) also used in [36].

First, we compare GateMod with the hybrid model from [36]. Fig. 4A shows that, in this case, GateMod achieves higher PXPs than the hybrid model on the same dataset for both experiments. Results are confirmed in Fig. 4B, where our model also robustly outperforms standard bandit algorithms (UCB, Thompson, and Value). Moreover, GateMod also reveals patterns (encoded in how primitives are orchestrated) that remain elusive for the other models (Fig. 4C). In Experiment 1, the use of primitives identified by GateMod shows that the weights associated to the risk-adverse primitive remain low, with the exploitation primitive being dominant. This suggests that individuals probe the stochastic arm that could potentially lead to higher rewards over the fixed zero reward. Instead, for Experiment 2, where both arms are stochastic, Fig. 4C shows an almost perfect periodicity, with intermittent spikes of uncertainty-seeking behavior suggesting that participants alternate between sampling and consolidating information in a structured way. In Supplementary Information we provide additional validations for these findings together with a 3-fold cross validation to assess predictive capabilities.

Discussion

We introduced GateMod, a theoretically grounded interpretable computational model linking the emergence of gating mechanisms to the underlying decision-making task, and to a neural circuit delivering this functionality. Unlike other models, where gating is imposed by a designer, in GateMod this key mechanism of intelligence is cast under a normative framework, GateFrame, and the resulting optimization is provably solvable by the associated continuous-time dynamics, GateFlow, and the neural circuit derived from it, GateNet.

In GateFrame, gating emerges from a rigorous variational, or Bayesian, normative framework that is flexible enough to capture a wide range of decision-making problems across behavioral/cognitive sciences, neuroscience and machine learning. GateMod not only provides this broad framework, but also the energy model, GateFlow, that provably finds the optimal solution. GateFlow belongs to a class of proximal gradient flows, the softmax flows, that we characterize in full and that can tackle general entropy-regularized problems. GateFlow has two key properties. First, its equilibrium is also GateFrame optimal solution. Second, GateFlow has strong contractivity properties so that convergence to the optimal solution is global, exponential, and with a guaranteed rate. Finally, GateNet is the soft-competitive recurrent neural circuit implementing GateFrame. Each element of the circuit is clear and interpretable in view of the agent

task. Remarkably, GateNet show that contextual computations are essential to deliver gating, and these computations may be implemented via dendritic connections.

Our approach harmonizes decision-making, complex dynamical systems and neural principles. To evaluate our results in full we considered two different domains. Our collective behaviors experiments showed that, in a setting where primitives are naturally mapped into well-documented social forces, GateMod can recover well-known phenomena. In the multi-armed experiments, our model served as a framework to gain insights into the exploration/exploitation balance, using primitive encoding simple behaviors.

Beyond gathering experimental evidence that natural agents can compose knowledge according to GateMod, our results also open a number of interdisciplinary research questions that we reserve for future work and discuss next.

How do we endow GateMod with learning? We considered settings where models and primitives are given. Exploiting the links between GateFrame objective and the variational free energy, integrating GateMod with active inference is a promising path to learn the transition kernels. However, the integration of GateMod within an inference and learning framework will likely prompt additional studies are needed to understand *what* makes for a *good* set of primitives and *how* these should be exploited for inference and learning. This question is reminiscent of *controllability* concepts in control theory. Natural intelligence has evolved methods to learn, evolve, and grow the right set of skills/primitives. We conjecture a suitable mechanism to endow GateMod with this capability is to introduce competition and evolutionary mechanisms for the primitives. From given skills, these would become entities competing to maximize their usage.

What happens when primitives are learned? Introducing competition will cause primitives to evolve and this calls for a normative framework that controls this process. An unregulated competition between primitives may lead to *rich-get-richer* effects where only a few primitives end up being used. In a game-theoretical context, the other primitives may either end-up being forgotten or would have to evolve, and become similar to the ones used the most. However, in this setting the agent may become fragile to (possibly, adversarial) changes in the environment. Diversity of primitives may confer robustness to an organism. Learning and evolution of the primitives may be studied through the hierarchical RL and game theoretical approaches. GateFlow shows a connection between our results and game theory. In Supplementary Information, we further show that dynamical systems close in spirit to GateFlow arise from best response maps in game-theoretical frameworks – this may be a starting point to address this question.

What about other composition rules? GateMod rigorously relates the onset of softmax gating mechanisms to entropic regularization and a statistical complexity objective. We showed this is a broad setting and GateMod can recover both sparse and dense gating architectures. However, the question of how other gating rules can emerge from similar normative frameworks remains. From a mathematical viewpoint, we conjecture that different gating rules will emerge from proximal gradients arising from different regularizers. Then, additional studies would be needed to find the *right* proximal gradients, characterize convergence of their flow and to seek neural implementations.

Our study complements machine learning and neuroscience approaches to understand gating mechanisms in natural and artificial agents. Unlike other approaches, we propose a general framework explaining how gating rules may emerge from decision-making processes under a normative framework, how the underlying optimization may be solved, and how this capability could be delivered via interpretable neural

circuits. Elaborating and adapting our work to different applications and multi-agent settings, addressing the above open questions, will lead to deeper exciting understanding of the role of gating in natural and machine intelligence.

Methods

The agent has access to: (i) the generative model, $q(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}) = q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) q(\mathbf{u}_k | \mathbf{x}_{k-1})$; (ii) the environment model $p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$; (iii) n_π primitives, $\pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})$, $\alpha = 1, \dots, n_\pi$. We also denote the state space by $\mathcal{X}$ and the action space by $\mathcal{U}$ (see Sec. 4 in Supplementary Information for notation and details).

GateFrame Properties

In Eq. [1] both the term $-\mathbf{H}(\mathbf{w}_k)$ in the cost and the constraints are convex in the decision variables $\mathbf{w}_k$. We now show (detailed derivations are provided in the Supplementary Information) that the statistical complexity term in the cost, $\mathbf{F}(\mathbf{w}_k)$ is convex in $\mathbf{w}_k$, thus making the overall problem strongly convex. To show this, we embed the constraints into the expression of $\mathbf{F}(\mathbf{w}_k)$. Then, using the chain rule for the KL divergence (Sec. 2 in Supplementary Information) reveals that the complexity term in the cost can be written as $F(\mathbf{w}_k) = \sum_{\alpha=1}^{n_\pi} w_k^\alpha (\mathbb{E}_{\pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})} [\ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) - \ln q(\mathbf{u}_k | \mathbf{x}_{k-1}) + D_{\text{KL}}(p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k))])$. Here, w_k^α is the element of the vector $\mathbf{w}_k$ corresponding to primitive α . The function $F(\mathbf{w}_k)$ is twice differentiable in $\mathbf{w}_k$. Computing its Hessian (Sections 3 and 4 in Supplementary Information) shows that this is a positive definite matrix, proving that the map $F(\mathbf{w}_k)$ is convex. Derivations are in Supplementary Information, where we also discuss a generalization of GateFrame optimization with a bias vector for the weights.

In Results we related GateFrame objective to other computational models. These connections were drawn by realizing that, if the generative model $q(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1})$ is proportional to $\tilde{q}(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}) e^{-c(\mathbf{x}_k, \mathbf{u}_k)}$, then the statistical complexity term in GateFrame can be written as [3]. In the derivations (see Supplementary Information for details) the normalizing constant Z is $\int_{\mathcal{U}} \int_{\mathcal{X}} \tilde{q}(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}) e^{-c(\mathbf{x}_k, \mathbf{u}_k)} d\mathbf{x}_k d\mathbf{u}_k$. Therefore, when minimizing this function with respect to the decision variables of GateFrame optimization $\mathbf{w}_k$, the term $\ln Z$ in [3] can be dropped as it does not affect the optimal weights. This relates the statistical complexity term in GateFrame objective to the MaxDiff policy computation framework, which in turn can generalize MaxEnt. We can further elaborate on the expression in [3] by leveraging the chain rule for the KL divergence. The first term in [3] becomes

$$D_{\text{KL}}(p(\mathbf{u}_k | \mathbf{x}_{k-1}) || q(\mathbf{u}_k | \mathbf{x}_{k-1})) + \mathbb{E}_{p(\mathbf{u}_k | \mathbf{x}_{k-1})} [D_{\text{KL}}(p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k))]. \quad (11)$$

This is the objective minimized under the expected free energy formulation from [89] and we refer to this work for connections with other decision-making schemes, such as KL control and control as inference, see, e.g., [96, 99]. We unpack the derivations in Supplementary Information.

Why GateFrame Yields Softmax Gating

Our starting point is Eq. [4], which is a reformulation of Eq. [1]. This reformulation follows from the definition of $\mathbf{H}_{\text{barrier}}$ to embed the the simplex constraint into the objective. To find the optimal solution of

the problem in Eq. [4] and show that GateFrame yields a softmax gating rule, we leverage the continuous-time proximal gradient method [1, 41, 25]. Consider a generic composite optimization problem of the form $\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) + g(\mathbf{x})$, where f is a smooth function and g convex, closed, proper, and possibly non-smooth. The continuous-time proximal gradient method continuously updates an estimate of the optimal solution via the dynamics

$$\dot{\mathbf{x}} = -\mathbf{x} + \text{prox}_{\gamma g}(\mathbf{x} - \gamma \nabla f(\mathbf{x})), \quad (12)$$

where $\text{prox}_{\gamma g}(\mathbf{x}) := \arg \min_{\mathbf{z} \in \mathbb{R}^n} g(\mathbf{z}) + \frac{1}{2\gamma} \|\mathbf{x} - \mathbf{z}\|_2^2$, for all $\mathbf{x} \in \mathbb{R}^n$ is the proximal operator of g and γ is a parameter, see, e.g., [74]. GateFlow is the continuous-time proximal-gradient dynamics associated to our GateFrame optimization. The expression reported in Results is obtained by first adding and subtracting the term $\varepsilon \frac{\|\mathbf{w}_k\|_2^2}{2}$ to GateFrame cost, and subsequently obtaining a suitable closed-form expression for the proximal gradient. More precisely, as we unpack in the Supplementary Information, the problem in Eq. [4] is first reformulated as

$$\min_{\mathbf{w}_k \in \mathbb{R}^{n_\pi}} f(\mathbf{w}_k) + g(\mathbf{w}_k),$$

where $f(\mathbf{w}_k) = F(\mathbf{w}_k) + \varepsilon \frac{\|\mathbf{w}_k\|_2^2}{2}$ is continuously differentiable and $g(\mathbf{w}_k) = \varepsilon \left(H_{\text{barrier}}(\mathbf{w}_k) - \frac{\|\mathbf{w}_k\|_2^2}{2} \right)$ is closed, proper and convex. Then, we rigorously show that $\text{prox}_{\gamma g}(\mathbf{w}_k)$ is the softmax function, and this yields GateFlow. Moreover, the equilibrium point of GateFlow is the optimal solution of Eq. [4]. See Supplementary Information for the formal treatment and for derivations of a generalization of GateFlow. This generalization comprises a broad class of continuous-time dynamics, that we term *softmax gradient flow*, for general entropy-regularized composite optimization problems including bias in the weights. Remarkably, as we show in Sec. S4 of Supplementary Information, introducing a bias weights vector in the formulation yields a Gumbel-softmax gating rule that naturally arises in dense MoE architectures. In the Supplementary Information we also derive the accompanying GateFlow and GateNet enabling the implementation of this mechanism.

Deriving GateFlow Convergence Properties

We start with deriving the forward invariance of the probability simplex Δ_{n_π} (that is, GateFrame feasibility domain) under GateFlow dynamics. This property, beyond being important *per se* as shown in Results, is also instrumental to prove GateFlow convergence.

Intuitively, forward invariance of the simplex means that when GateFrame is initialized at some initial condition $\mathbf{w}_k(0) \in \Delta_{n_\pi}$, then its trajectories remains in the simplex for all t , that is, $\mathbf{w}_k(t) \in \Delta_{n_\pi}$, for all $t > 0$. By noticing that $\Delta_{n_\pi} = \mathbb{R}_{\geq 0}^{n_\pi} \cap \left\{ \mathbf{w}_k \in \mathbb{R}^{n_\pi} \mid \sum_{\alpha=1}^{n_\pi} w_k^\alpha = 1 \right\}$, the invariance property can be shown by proving two properties: (i) the positive orthant is forward invariant; (ii) total *mass* is conserved, that is, $\sum_{\alpha=1}^{n_\pi} \frac{dw_k^\alpha}{dt} = 0$ for all $\mathbf{w}_k \in \mathbb{R}^{n_\pi}$ such that $\sum_{\alpha=1}^{n_\pi} w_k^\alpha = 1$. The first property can be shown via Nagumo's theorem [70] (see also [14, Exercise 3.12]), proving that each component of GateFlow vector field – the right hand side of Eq. [5] – is non-negative at the frontier of the positive orthant. This means that trajectories that are on the frontier are subject to a vector field that points them towards the internal of the orthant. This prohibits that trajectories fall outside the positive orthant. The second property follows from direct computation. Summing all the components of GateFlow yields $\sum_{\alpha=1}^{n_\pi} (\text{softmax}(-\varepsilon^{-1} \nabla F(\mathbf{w}_k))_\alpha - w_k^\alpha) = 0$, as desired. In this last expression the subscript α in the softmax is its α -th component.

GateFlow convergence follows from a stronger contractivity property [57, 14]. A dynamical system is strongly contracting if any two of its trajectories converge towards each other exponentially. This means that the distance, defined with respect to some norm, shrinks in time. This property can be shown by

investigating the Jacobian of the vector field. More precisely, contractivity with respect to the Euclidean norm can be established by proving that all the eigenvalues of the symmetric part of the Jacobian matrix are strictly negative. Moreover, the largest eigenvalue provides the guaranteed convergence rate for the trajectories [14]. In Supplementary Information, we rigorously show that this property is satisfied by the Jacobian of GateFlow and prove that the convergence rate is $\frac{1}{\tau}$ as reported in Results. Beyond showing uniqueness of the equilibrium, building on this desirable convergence property for GateFlow, we also show that trajectories are entrained to a periodic solution if parameters are periodic [85], and the dynamics remain contracting when discretized with a suitably chosen step-size [15]. We give the formal treatment of both the statements and the proofs of these properties in the Supplementary Information.

Deriving GateNet

Eq. [8] and Eq. [9] of GateNet neural circuit (Fig. 2C) are obtained from GateFlow. Here we describe how these equations are derived. See Supplementary Information for additional details and formal treatment.

The fast unit (Fig. 2C and Eq. [8]) returns the exponent of the softmax in GateFlow. This is obtained from the gradient of $F(\mathbf{w}_k)$, which we recall being the task-encoding statistical complexity in GateFrame optimization. The negative of the gradient of $F(\mathbf{w}_k)$ is a n_π -dimensional vector and its i -th component is

$$-\mathbb{E}_{\pi^i(\mathbf{u}_k|\mathbf{x}_{k-1})} \left[\ln \mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1}) + \mathbf{c}(\mathbf{x}_{k-1}, \mathbf{u}_k) \right] + 1 \quad (13)$$

From this expression, the first step to obtain the fast unit dynamics Eq. [8] is to note that the constant term can be dropped from the softmax due to its translation invariance property. In the discrete actions setting, this yields the i -th component of the vector $\varepsilon \mathbf{y}$ in Eq. [8]. This is the vector that we want GateNet fast unit to return. The second step to obtain the fast unit dynamics is to conveniently reformulate Eq. [8] into a form that is suitable for a neural dynamics representation. To this aim, Eq. [8] can be conveniently written as $-\varepsilon^{-1} \Pi(\mathbf{x}_{k-1})^\top (\mathbf{b} + \mathbf{c})$, where $\bar{\mathbf{b}}$ is $\ln(\bar{\mathbf{a}})$, with the logarithm being component-wise and with $\bar{\mathbf{a}}$ being $\Pi(\mathbf{x}_{k-1})\mathbf{w}_k$. These quantities are the globally stable equilibria of the input and hidden layers dynamics in GateNet circuit (Fig. 2C).

The slow unit in Fig. 2C returns the optimal weights solving GateFrame – it implements the softmax gating rule. The circuit, available in the literature [92], implements the softmax rule by leverage the following general identity (see Supplementary Information for derivation): $\text{softmax}(y)_i = \exp(y_i - \log(m))$.

Experiments Settings

In the collective behavior experiments, the state of the i -th boid is $\mathbf{x}_k^i = [\mathbf{p}_k^i, \mathbf{v}_k^i]^\top$, where $\mathbf{p}_k^i$ and $\mathbf{v}_k^i$ are the 2-dimensional position and velocity vectors. When a boid is equipped with GateMod, at time-step k it determines its acceleration $\mathbf{u}_k^*$ by sampling from the optimal policy Eq. [10]. In the experiments, both maximum velocity and acceleration are bounded. As in, e.g., [42], the boid dynamics $p(\mathbf{x}_k^i | \mathbf{x}_{k-1}^i, \mathbf{u}_k^i)$ is the Gaussian $\mathcal{N}(\bar{\mathbf{x}}_{p,k}^i, \Sigma_p)$. The center of the Gaussian is updated according to the model from, e.g., [24, 53], and Σ_p is $0.01\mathbb{I}_4$ ($\mathbb{I}_4$ is the 4×4 identity matrix). The expressions for the separation, alignment and cohesion primitives are in accordance with prior literature and detailed in Supplementary Information (Sec. 5).

In the multi-armed bandit experiments, the state of participant i is $\mathbf{x}_k = [\mu_{1,k}, s_{1,k}, \dots, \mu_{N,k}, s_{N,k}]^\top$, where $\mu_{i,k}$ and $s_{i,k}$ denote the estimated posterior mean and variance of arm i , respectively. Belief updating over rewards is modeled with a Kalman filter. We use the original code from [36] to compute posterior estimates and to implement the Hybrid, UCB, Thompson and Value policies. The expressions

for the primitives are in Supplementary Information (Sec. 5) where we provide supplementary details and experiments.

Acknowledgments. We acknowledge the use of ChatGPT for assistance in improving the wording and grammar of this document. VC affiliated with University of Salerno while preparing the theoretical bulk of this work. VC was supported by the European Union-Next Generation EU Mission 4 Component 1 CUP E53D23014640001. FB is supported in part by AFOSR grant FA9550-22-1-0059. GR is supported by the European Union-Next Generation EU Mission 4 Component 1 CUP E53D23014640001.

References

References

- [1] B. Abbas and H. Attouch. Dynamical systems and forward-backward algorithms associated with the sum of a convex subdifferential and a monotone cocoercive operator. *Optimization*, 64(10):2223–2252, 2014.
- [2] B. Abbas and H. Attouch. Dynamical systems and forward–backward algorithms associated with the sum of a convex subdifferential and a monotone cocoercive operator. *Optimization*, 64(10):2223–2252, 2015.
- [3] G. Albi, D. Balagué, J. A. Carrillo, and J. von Brecht. Stability Analysis of Flock and Mill Rings for Second Order Models in Swarming. *SIAM Journal on Applied Mathematics*, 74(3):794–818, 2014.
- [4] G. Albi, M. Bongini, E. Cristiani, and D. Kalise. Invisible control of self-organizing agents leaving unknown environments. *SIAM Journal on Applied Mathematics*, 76(4):1683–1710, 2016.
- [5] G. Albi, L. Pareschi, and M. Zanella. Uncertainty quantification in control problems for flocking models. *Mathematical Problems in Engineering*, 2015(1):850124, 2015.
- [6] Z. A. Ali, E. H. Alkhamash, and R. Hasan. State-of-the-art flocking strategies for the collective motion of multi-robots. *Machines*, 12(10):739, 2024.
- [7] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. Empirical investigation of starling flocks: a benchmark study in collective animal behaviour. *Animal Behaviour*, 76(1):201–215, 2008.
- [8] L. E. Beaver and A. A. Malikopoulos. An overview on optimal flocking. *Annual Reviews in Control*, 51:88–99, 2021.
- [9] S. Bernardi, R. Eftimie, and K. J. Painter. Leadership through influence: what mechanisms allow leaders to steer a swarm? *Bulletin of Mathematical Biology*, 83(6):69, 2021.
- [10] T. A. Berrueta, A. Pinosky, and T. D. Murphey. Maximum diffusion reinforcement learning. *Nature Machine Intelligence*, 6(5):504–514, 2024.
- [11] Thomas A. Berrueta, Allison Pinosky, and Todd D. Murphey. Maximum diffusion reinforcement learning. *Nature Machine Intelligence*, 6(5):504–514, May 2024.

- [12] D. Biro, D. J. T. Sumpter, J. Meade, and T. Guilford. From compromise to leadership in pigeon homing. *Current Biology*, 16(21):2123–2128, 2006.
- [13] Matthew Botvinick and Marc Toussaint. Planning as inference. *Trends in Cognitive Sciences*, 16(10):485–488, October 2012.
- [14] F. Bullo. *Contraction Theory for Dynamical Systems*. Kindle Direct Publishing, 1.2 edition, 2024.
- [15] F. Bullo, P. Cisneros-Velarde, A. Davydov, and S. Jafarpour. From contraction theory to fixed point algorithms on Riemannian and non-Euclidean spaces. In *IEEE Conf. on Decision and Control*, December 2021.
- [16] W. Cai, J. Jiang, F. Wang, J. Tang, S. Kim, and J. Huang. A survey on mixture of experts in large language models. *IEEE Transactions on Knowledge and Data Engineering*, page 1–20, 2025.
- [17] J. A. Carrillo, M. Fornasier, G. Toscani, and F. Vecil. Particle, kinetic, and hydrodynamic models of swarming. In *Mathematical Modeling of Collective Behavior in Socio-Economic and Life Sciences*, pages 297–336. 2010.
- [18] H. Celikkanat and E. Şahin. Steering self-organized robot flocks through externally guided individuals. *Neural Computing and Applications*, 19(6):849–865, 2010.
- [19] V. Centorrino, A. Gokhale, A. Davydov, G. Russo, and F. Bullo. Positive competitive networks for sparse reconstruction. *Neural Computation*, 36(6):1163–1197, 2024.
- [20] Anne G. E. Collins. A habit and working memory model as an alternative account of human reward-based learning. *Nature Human Behaviour*, November 2025.
- [21] Antoine Collomb-Clerc, Maëlle C. M. Gueguen, Lorella Minotti, Philippe Kahane, Vincent Navarro, Fabrice Bartolomei, Romain Carron, Jean Regis, Stephan Chabardès, Stefano Palminteri, and Julien Bastin. Human thalamic low-frequency oscillations correlate with expected value and outcomes during reinforcement learning. *Nature Communications*, 14(1), October 2023.
- [22] I. D. Couzin, J. Krause, N. R. Franks, and S. A. Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, 2005.
- [23] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective Memory and Spatial Sorting in Animal Groups. *Journal of Theoretical Biology*, 218(1):1–11, 2002.
- [24] F. Cucker and S. Smale. Emergent behavior in flocks. *IEEE Transactions on Automatic Control*, 52(5):852–862, 2007.
- [25] A. Davydov, V. Centorrino, A. Gokhale, G. Russo, and F. Bullo. Time-varying convex optimization: A contraction and equilibrium tracking approach. *IEEE Transactions on Automatic Control*, 70(11):7446–7460, 2025.
- [26] N. D. Daw, J. P. O’doherly, P. Dayan, B. Seymour, and R.J. Dolan. Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095):876–879, 2006.
- [27] Yiannis Demiris and Bassam Khadhour. Hierarchical attentive multiple models for execution and recognition of actions. *Robotics and Autonomous Systems*, 54(5):361–369, May 2006.

- [28] H. Duan and C. Sun. Swarm intelligence inspired skills and the evolution of cooperation. *Scientific Reports*, 4(1):5210, 2014.
- [29] Benjamin Eysenbach and Sergey Levine. Maximum entropy RL (provably) solves some robust RL problems. In *International Conference on Learning Representations*, 2022.
- [30] W. Fedus, B. Zoph, and N. Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022.
- [31] J.A. Feldman and D.H. Ballard. Connectionist models and their properties. *Cognitive Science*, 6(3):205–254, July 1982.
- [32] K. Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.
- [33] J. M. Fuster. The Prefrontal Cortex—An Update. *Neuron*, 30(2):319–333, May 2001.
- [34] K. Genter. Ad hoc teamwork for leading a flock. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems*, AAMAS ’13, page 1431–1432, 2013.
- [35] K. Genter and P. Stone. Adding Influencing Agents to a Flock. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, AAMAS ’16, page 615–623, 2016.
- [36] S. J. Gershman. Deconstructing the human algorithms for exploration. *Cognition*, 173:34–42, 2018.
- [37] A. Gokhale, A. Davydov, and F. Bullo. Proximal gradient dynamics: Monotonicity, exponential convergence, and applications. *IEEE Control Systems Letters*, 8:2853–2858, 2024.
- [38] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 1861–1870, 10–15 Jul 2018.
- [39] J. Han, M. Li, and L. Guo. Soft control on collective behavior of a group of autonomous agents by a skill agent. *Journal of Systems Science and Complexity*, 19(1):54–62, 2006.
- [40] J. Han and L. Wang. Nondestructive intervention to multi-agent systems through an intelligent agent. *PloS One*, 8(5), 2013.
- [41] S. Hassan-Moghaddam and M. R. Jovanović. Proximal gradient flow and Douglas-Rachford splitting dynamics: Global exponential stability via integral quadratic constraints. *Automatica*, 123:109311, 2021.
- [42] C. Heins, B. Millidge, L. Da Costa, R. P. Mann, K. J. Friston, and I. D. Couzin. Collective behavior from surprise minimization. *Proceedings of the National Academy of Sciences*, 121(17):e2320239121, 2024.
- [43] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton. Adaptive Mixtures of Local Experts. *Neural Computation*, 3(1):79–87, February 1991.
- [44] E. Jang, S. Gu, and B. Poole. Categorical Reparameterization with Gumbel-Softmax. In *International Conference on Learning Representations*, 2017.

- [45] Kevin Johnston, Helen M. Levin, Michael J. Koval, and Stefan Everling. Top-down control-signal dynamics in anterior cingulate and prefrontal cortex neurons following task switching. *Neuron*, 53(3):453–462, February 2007.
- [46] M. I. Jordan and R. A. Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6(2):181–214, March 1994.
- [47] Hilbert J. Kappen, Vicenç Gómez, and Manfred Opper. Optimal control as a graphical model inference problem. *Machine Learning*, 87(2):159–182, February 2012.
- [48] Mikail Khona and Ila R. Fiete. Attractor and integrator networks in the brain. *Nature Reviews Neuroscience*, 23(12):744–766, November 2022.
- [49] Denis Kleyko, Christopher J. Kymn, Anthony Thomas, Bruno A. Olshausen, Friedrich T. Sommer, and E. Paxon Frady. Principled neuromorphic reservoir computing. *Nature Communications*, 16(1), January 2025.
- [50] Christof Koch and Tomaso Poggio. *Single Neuron Computation*, chapter Multiplying with Synapses and Neurons, pages 315–345. 1992.
- [51] J. Krause, A. F. T. Winfield, and J. Deneubourg. Interactive robots in experimental biology. *Trends in Ecology & Evolution*, 26(7):369–375, 2011.
- [52] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, 2016.
- [53] H. Levine, W.J. Rappel, and I. Cohen. Self-organization in systems of self-propelled particles. *Phys. Rev. E*, 63:017101, Dec 2000.
- [54] G. Lin, R. Escobedo, X. Li, T. Xue, Z. Han, C. Sire, V. Guttal, and G. Theraulaz. Experimental Evidence of Stress-Induced Critical State in Schooling Fish. *PRX Life*, 3:033018, Sep 2025.
- [55] H. Ling, G. E. McIvor, J. Westley, K. van der Vaart, R. T. Vaughan, A. Thornton, and N. T. Ouellette. Behavioural plasticity and the transition to order in jackdaw flocks. *Nature Communications*, 10(1):5174, 2019.
- [56] David Lipshutz, Charles Windolf, Siavash Golkar, and Dmitri Chklovskii. A biologically plausible neural network for slow feature analysis. In *Advances in Neural Information Processing Systems*, volume 33, pages 14986–14996, 2020.
- [57] W. Lohmiller and J.-J. E. Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683–696, 1998.
- [58] Adam S. Lowet, Qiao Zheng, Melissa Meng, Sara Matias, Jan Drugowitsch, and Naoshige Uchida. An opponent striatal circuit for distributional reinforcement learning. *Nature*, 639(8055):717–726, February 2025.
- [59] Abdullah Makkeh, Marcel Graetz, Andreas C. Schneider, David A. Ehrlich, Viola Priesemann, and Michael Wibral. A general framework for interpretable neural learning based on local information-theoretic goal functions. *Proceedings of the National Academy of Sciences*, 122(10), March 2025.

- [60] S. Masoudnia and R. Ebrahimpour. Mixture of experts: a literature survey. *Artificial Intelligence Review*, 42(2):275–293, 2014.
- [61] P. Mazzaglia, T. Verbelen, and B. Dhoedt. Contrastive active inference. In *Advances in Neural Information Processing Systems*, volume 34, pages 13870–13882. Curran Associates, Inc., 2021.
- [62] Bartlett Mel and Christof Koch. Sigma-pi learning: On radial basis functions and cortical associative learning. In D. Touretzky, editor, *Advances in Neural Information Processing Systems*, volume 2. Morgan-Kaufmann, 1989.
- [63] Bartlett W. Mel. Synaptic integration in an excitable dendritic tree. *Journal of Neurophysiology*, 70(3):1086–1101, September 1993.
- [64] J. A. Mendez and E. Eaton. How to reuse and compose knowledge for a lifetime of tasks: A survey on continual learning and functional composition. *Transactions on Machine Learning Research*, 2023.
- [65] B. Moore. Principal component analysis in linear systems: Controllability, observability, and model reduction. *IEEE Transactions on Automatic Control*, 26(1):17–32, 1981.
- [66] S. Motsch and E. Tadmor. A new model for self-organized dynamics and its flocking behavior. *Journal of Statistical Physics*, 144(5):923, 2011.
- [67] Siyuan Mu and Sen Lin. A comprehensive survey of mixture-of-experts: Algorithms, theory, and applications, 2025.
- [68] M. Muqeeth, H. Liu, and C. Raffel. Soft Merging of Experts with Adaptive Routing. *arXiv*, 2024.
- [69] Kevin P. Murphy. *Probabilistic Machine Learning: Advanced Topics*. MIT Press, 2023.
- [70] M. Nagumo. Über die Lage der Integralkurven gewöhnlicher Differentialgleichungen. *Proceedings of the Physico-Mathematical Society of Japan. 3rd Series*, 24:551–559, 1942.
- [71] S. Nakayama, M. Ruiz Marín, M. Camacho, and M. Porfiri. Plasticity in leader–follower roles in human teams. *Scientific Reports*, 7(1):14562, 2017.
- [72] A. A. Neath and J. E. Cavanaugh. The Bayesian information criterion: background, derivation, and applications. *WIREs Computational Statistics*, 4(2):199–203, 2012.
- [73] R. Olfati-Saber. Flocking for multi-agent dynamic systems: Algorithms and theory. *IEEE Transactions on Automatic Control*, 51(3):401–420, 2006.
- [74] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):127–239, 2014.
- [75] Thomas Parr and Karl J. Friston. Generalised free energy and active inference. *Biological Cybernetics*, 113(5–6):495–513, September 2019.
- [76] Panayiota Poirazi and Bartlett W. Mel. Impact of active dendrites and structural plasticity on the memory capacity of neural tissue. *Neuron*, 29(3):779–796, March 2001.
- [77] T. J. Prescott and S. P. Wilson. Understanding brain functional architecture through robotics. *Science Robotics*, 8(78), 2023.

- [78] C. W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, 1987.
- [79] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, page 25–34, 1987.
- [80] L. Rigoux, K.E. Stephan, K.J. Friston, and J. Daunizeau. Bayesian model selection for group studies — Revisited. *NeuroImage*, 84:971–985, 2014.
- [81] R. V. Rikhye, A. D. Gilra, and M. M. Halassa. Thalamic regulation of switching between cortical representations enables cognitive flexibility. *Nature Neuroscience*, 21(12):1753–1763, November 2018.
- [82] F. Rossi, E. Garrabé, and G. Russo. Free-Gate: Planning, Control and Policy Composition via Free Energy Gating. In *Proceedings of the 6th International Workshop on Active Inference (to appear)*. Springer, 2025.
- [83] Nicolas P. Rougier, David C. Noelle, Todd S. Braver, Jonathan D. Cohen, and Randall C. O’Reilly. Prefrontal cortex and flexible cognitive control: Rules without symbols. *Proceedings of the National Academy of Sciences*, 102(20):7338–7343, May 2005.
- [84] David E. Rumelhart and James L. McClelland. *A General Framework for Parallel Distributed Processing*, pages 45–76. 1987.
- [85] G. Russo, M. Di Bernardo, and E. D. Sontag. Global entrainment of transcriptional systems to periodic inputs. *PLoS Computational Biology*, 6(4):e1000739, 2010.
- [86] T. D. Sanger. Risk-aware control. *Neural Computation*, 26(12):2669–2691, 2014.
- [87] S. Sayin, E. Couzin-Fuchs, I. Petelski, Y. Günzel, M. Salahshour, C. Lee, J. M. Graving, L. Li, O. Deussen, G. A. Sword, and I. D. Couzin. The behavioral mechanisms governing collective motion in swarming locusts. *Science*, 387(6737):995–1000, 2025.
- [88] G. Schwarz. Estimating the Dimension of a Model. *The Annals of Statistics*, 6(2):461–464, 1978.
- [89] A. Shafiei, H. Jesawada, K. Friston, and G. Russo. Distributionally robust free energy principle for decision-making. *Arxiv report*, 2025.
- [90] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *International Conference on Learning Representations*, 2017.
- [91] O. Simeone. A brief introduction to machine learning for engineers. *Foundations and Trends in Signal Processing*, 12(3-4):200–431, 2018.
- [92] M. Snow and J. Orchard. Biological softmax: Demonstrated in modern Hopfield networks. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 44, 2022.
- [93] D. J. T. Sumpter. The principles of collective animal behaviour. *Philosophical Transactions of The Royal Society B: Biological Sciences*, 361(1465):5–22, 2006.
- [94] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.

- [95] J. Tani and S. Nolfi. Learning to perceive the world as articulated: an approach for hierarchical learning in sensory-motor systems. *Neural Networks*, 12(7–8):1131–1141, October 1999.
- [96] Emanuel Todorov. Efficient computation of optimal actions. *Proceedings of the National Academy of Sciences*, 106(28):11478–11483, 2009.
- [97] M. S. Tomov, V. Q. Truong, R. A. Hundia, and S. J. Gershman. Dissociable neural correlates of uncertainty underlie different exploration strategies. *Nature Communications*, 11(1):2371, 2020.
- [98] B. Tsuda, K. M. Tye, H. T. Siegelmann, and T. J. Sejnowski. A modeling framework for adaptive lifelong learning with transfer and savings through gating in the prefrontal cortex. *Proceedings of the National Academy of Sciences*, 117(47):29872–29882, November 2020.
- [99] Bart van den Broek, Wim Wiegerinck, and Bert Kappen. Risk sensitive path integral control. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*, page 615–622. AUAI Press, 2010.
- [100] G. Vásárhelyi, C. Virágh, G. Somorjai, T. Nepusz, A. E. Eiben, and T. Vicsek. Optimized flocking of autonomous drones in confined environments. *Science Robotics*, 3(20), 2018.
- [101] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6-7):1226–1229, 1995.
- [102] T. Vicsek and A. Zafeiris. Collective motion. *Physics Reports*, 517(3):71–140, 2012.
- [103] P. Vijayaraghavan, J. F. Queisser, S. V. Flores, and J. Tani. Development of compositionality through interactive learning of language and action of robots. *Science Robotics*, 10(98), January 2025.
- [104] Jane X. Wang, Zeb Kurth-Nelson, Dharshan Kumaran, Dhruva Tirumala, Hubert Soyer, Joel Z. Leibo, Demis Hassabis, and Matthew Botvinick. Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, 21(6):860–868, May 2018.
- [105] L. Wang and L. Guo. Robust consensus and soft control of multi-agent systems with noises. *Journal of Systems Science and Complexity*, 21(3):406–415, 2008.
- [106] Grady Williams, Nolan Wagener, Brian Goldfain, Paul Drews, James M. Rehg, Byron Boots, and Evangelos A. Theodorou. Information theoretic mpc for model-based reinforcement learning. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1714–1721, 2017.
- [107] D.M. Wolpert and M. Kawato. Multiple paired forward and inverse models for motor control. *Neural Networks*, 11(7–8):1317–1329, October 1998.
- [108] S. E. Yuksel, J. N. Wilson, and P. D. Gader. Twenty Years of Mixture of Experts. *IEEE Transactions on Neural Networks and Learning Systems*, 23(8):1177–1193, 2012.
- [109] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey. Maximum entropy inverse reinforcement learning. In *Proceedings of the 23rd National Conference on Artificial Intelligence*, page 1433–1438. AAAI Press, 2008.




Figure 3: **A** A boid in a flock of N boids. Position and velocity components form 4-dimensional state $\mathbf{x}_k^i$; $\mathbf{u}_k^i$ is the acceleration vector. We use the superscript to denote that states/actions are those of the i -th boid in the flock. The acceleration is built upon the social forces and a boid can only use information from boids within its field of view. The field angle, α , is set to 320° in the experiments. The radii correspond to [24, 53] three concentric separation, alignment and cohesion zones. **B** GateFrame optimization is solved via GateFlow. Starting from an initial feasible initial condition, GateFlow trajectories converge to GateFrame optimal solution. We recall that GateFlow is an energy model and, along its trajectories, the energy $F - \varepsilon H$ decreases. **C** GateMod recovers polarization. Trajectories of the $N = 40$ boids from random initial conditions (left). The group exhibits polarization and this is confirmed in the middle panel, showing the evolution of the polarization order parameter – the average normalized speed across boids – over simulation time. A value of 1 of this parameter indicates perfect alignment between boids. Right: time evolution of the optimal primitives’ weights from GateFlow. The evolution is shown for a representative agent and reveals that – after an initial transient when cohesion prevails – the weights tend toward a nearly uniform distribution. The weights evolution for all the boids in the experiments are in Fig. S1 of Supplementary Information. See also Sec. 5 therein. GateMod can also recover milling when a collision-avoidance term is introduced in the generative model. See Fig. S2 in Supplementary Information. **D** Collective behavior of boids when the generative model of a few boids (10%) encodes a goal-directed behavior. The other boids have the same generative model from Fig. 3C. Trajectories of the boids from random initial positions (left) illustrate global convergence toward the goal; temporal evolution of group-level metrics (middle) reveals a transition from disordered movement to polarized, goal-directed behavior; time evolution of the optimal primitives’ weights for a representative goal-informed boid and an uninformed one (right): the evolution of the informed boid’s weights suggests an adaptive goal-directed behavior; the uninformed boid’s weights, after an initial transient in which they get closer and aligned to the others, tend toward a uniform distribution. Findings are confirmed for different numbers of informed boids, different temperature values and different models for the followers. See Sec. 5 in Supplementary Information. Simulation parameters are in Tab. S1 of Supplementary Information.




Figure 4: **A** Comparison between Hybrid model from [36] and GateMod in terms of PXP. Higher PXP for a given model suggests that the model provides better explanations for the data. Formally, PXP quantifies the probability that each considered model is the most frequent process that generated the data. To obtain the PXP, we start from GateMod optimal policy. The policy at each trial is used to compute the Bayesian Information Criterion (BIC) [88, 72] values. Then, these are submitted to hierarchical Bayesian model selection [80]. **B** PXP comparison between UCB, Thompson, Value and GateMod. GateMod has robustly achieves the highest PXP across both experiments. **C** Evolution of the mean (bold line) and std (shaded area) of primitives' weights across subjects per trial. Weights show a pattern, suggesting that primitives might encode a *mental schema* adopted in similar ways by humans in the same context.

Supplementary Information: Neural Policy Composition from Free Energy Minimization

Francesca Rossi ^{1,*}

Veronica Centorrino ^{2,*}

Francesco Bullo ^{3,†}

Giovanni Russo ^{4,†} ✉

December 5, 2025

1 Introduction

We provide supplementary figures and formal details for the statements in the main text. After providing some background in Sec. 2, we start with considering a general class of entropy-regularized minimization problems and obtain a dynamical system that provably converges to their optimal solution (Sec. 3). Then, building on these general findings, we develop (Sec. 4) the formal treatment for GateFrame, GateFlow, and GateNet. After reporting the supplementary details of the experiments (Sec. 5) we provide the proofs of all the statements (Sec. 6). Finally, in Sec. 7 and Sec. 8, we report supplementary tables and figures for GateMod experiments in the main text.

2 Background

After introducing notation, definitions and standard results related to KL divergence and entropy, we briefly survey key elements of convex optimization and proximal operator theory relevant to our analysis. Then, we provide a primer on the key tool we use to assess convergence of our model.

Notation

Sets are in *calligraphic* letters and vectors in **bold**. The *probability simplex* in $\mathbb{R}^n$ is the set $\Delta_n := \left\{ p \in \mathbb{R}^n \mid \sum_{i=1}^n p_i = 1, p_i \geq 0, \forall i \in \{1, \dots, n\} \right\}$. Random variables are denoted via capital letters and their realization by lower-case letters. The probability density function (pdf) of $\mathbf{V}$ is denoted by $p(\mathbf{v})$; for discrete variables, $p(\mathbf{v})$ is the probability mass function (pmf). The expectation of a function $\mathbf{h}(\cdot)$ with respect to a random variable $\mathbf{V}$ is defined as $\mathbb{E}_p[\mathbf{h}(\mathbf{V})] := \int_{\mathcal{S}(p)} \mathbf{h}(\mathbf{v}) p(\mathbf{v}) d\mathbf{v}$, where $\mathcal{S}(p)$ denotes the support of $p(\mathbf{v})$. For discrete variables, the integral is replaced by a sum. When clear from context, we omit the domain of integration or summation. The joint pdf/pmf of two random variables $\mathbf{V}_1$ and $\mathbf{V}_2$ is $p(\mathbf{v}_1, \mathbf{v}_2)$ and the conditional pdf/pmf of $\mathbf{V}_1$ with respect to (w.r.t.) $\mathbf{V}_2$ is $p(\mathbf{v}_1 \mid \mathbf{v}_2)$. Given two pdfs/pmfs $p(\mathbf{v})$

¹ Scuola Superiore Meridionale, Italy. ² ETH, Zürich. ³ Center for Control, Dynamical Systems, and Computation, UC Santa Barbara, CA, USA. ⁴ Department of Information and Electrical Engineering and Applied Mathematics, University of Salerno, Italy. ^{*,†} These authors contributed equally. ✉ e-mail: giovarusso@unisa.it

and $q(\mathbf{v})$, we say that $p(\mathbf{v})$ is *absolutely continuous* with respect to (w.r.t.) $q(\mathbf{v})$ if the support of $p(\mathbf{v})$ is contained in the support of $q(\mathbf{v})$. We adopt the standard convention $0 \ln(0) = 0$.

The multivariate Gaussian distribution with mean $\boldsymbol{\mu} \in \mathbb{R}^d$ and covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$ is

$$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{d/2} |\boldsymbol{\Sigma}|^{1/2}} \exp \left(-\frac{1}{2} (\mathbf{v} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{v} - \boldsymbol{\mu}) \right).$$

and the notation $\mathbf{V} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ means that $\mathbf{V}$ is sampled from $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Also, we denote by $\mathbf{1}_n$ the all-ones vector in $\mathbb{R}^n$.

The Softmax Function

The softmax function is the map $\text{softmax}: \mathbb{R}^n \rightarrow \mathbb{R}^n$ defined by $\text{softmax}(\mathbf{x}) = (\text{softmax}(\mathbf{x})_i)_{i \in \{1, \dots, n\}}$, where

$$\text{softmax}(\mathbf{x})_i := \frac{e^{x_i}}{\sum_{j=1}^n e^{x_j}}.$$

This function – providing a differentiable and continuous approximation of the argmax operator [23] – naturally arises in the context of statistical physics (here it is also known as *Boltzmann* or *Gibbs distribution* [57]), economics and game theory (where it is often referred to as *logit response function*, *logit map*, or *perturbed best response* [37, 9, 10, 32]), and machine learning [44, 19, 52, 23, 30]. In neuroscience, this function emerges, e.g., in Winner-Take-All (WTA) networks to promote competition among neurons [17, 62]. Finally, we recall the following.

Lemma 1 (Softmax translation invariance). *If the function F is of the form $F(\mathbf{x}) = \tilde{F}(\mathbf{x}) + u \mathbf{1}_n^\top \mathbf{x}$, for some $u \in \mathbb{R}$, then*

$$\text{softmax}(\nabla F(\mathbf{x})) = \text{softmax}(\nabla \tilde{F}(\mathbf{x}) + u \mathbf{1}_n) = \text{softmax}(\nabla \tilde{F}(\mathbf{x})).$$

The Kullback-Leibler Divergence and Entropy

We recall the definition of Kullback-Leibler divergence, which provides a measure of discrepancy between two probability distributions.

Definition 1 (Kullback-Leibler Divergence [31]). *Given two probability distributions $p(\mathbf{v})$ and $q(\mathbf{v})$, the Kullback-Leibler (KL) divergence of p from q is*

$$D_{\text{KL}}(p \parallel q) := \int_{\mathcal{S}(p)} p(\mathbf{v}) \ln \left(\frac{p(\mathbf{v})}{q(\mathbf{v})} \right) d\mathbf{v}.$$

The $D_{\text{KL}}(p \parallel q)$ is finite only if $p(\mathbf{v})$ is absolutely continuous w.r.t. $q(\mathbf{v})$ (see, e.g., [12, Chapter 8]). The following result, see, e.g., [12, Theorem 2.5.3], is used in the main text.

Lemma 2 (Chain rule for the KL divergence). *Let $p(\mathbf{v}, \mathbf{u})$ and $q(\mathbf{v}, \mathbf{u})$ be joint probability density functions. Then,*

$$D_{\text{KL}}(p(\mathbf{v}, \mathbf{u}) \parallel q(\mathbf{v}, \mathbf{u})) = D_{\text{KL}}(p(\mathbf{v}) \parallel q(\mathbf{v})) + \mathbb{E}_{p(\mathbf{v})} [D_{\text{KL}}(p(\mathbf{v} \mid \mathbf{u}) \parallel q(\mathbf{v} \mid \mathbf{u}))]$$

Definition 2 (Entropy and Cross-Entropy). *The entropy function is $H: \Delta_n \rightarrow [0, \ln n]$, $H(p) := -\sum_{i=1}^n p_i \ln p_i$.*

The cross-entropy of q relative to p is $H: \Delta_n \times \Delta_n \rightarrow [0, +\infty]$, $H(p, q) := -\sum_{i=1}^n p_i \ln q_i$.

Proposition 1 (Relation between KLdivergence and entropy [59]). *The following identity holds: $D_{\text{KL}}(p \parallel q) = H(p, q) - H(p)$.*

Convex Analysis and Proximal Gradient Dynamics

Given a convex set $\mathcal{C}$, the *zero-infinity indicator function on $\mathcal{C}$* is the map $\iota_{\mathcal{C}}: \mathbb{R}^n \rightarrow [0, +\infty]$ defined by $\iota_{\mathcal{C}}(x) = 0$ if $x \in \mathcal{C}$ and $\iota_{\mathcal{C}}(x) = +\infty$ otherwise. We recall the following standard definition.

Definition 3 (*L-Lipschitz function*). *Let $(\Theta, \|\cdot\|_{\Theta})$ be a normed vector space. A map $f: \Theta \rightarrow \mathbb{R}$ is Lipschitz with constant L (*L-Lipschitz*) if for all $\theta_1, \theta_2 \in \Theta$, it holds that*

$$|f(\theta_1) - f(\theta_2)| \leq L \|\theta_1 - \theta_2\|_{\Theta}.$$

We denote by $\text{Lip}(f)$ the smallest such constant L , called the Lipschitz constant of f .

If f is a multivariable function, we write $\text{Lip}_{(\cdot)}(f)$ to specify the variable with respect to which the Lipschitz constant is computed. A map $g: \mathbb{R}^n \rightarrow \overline{\mathbb{R}} := [-\infty, +\infty]$, is (i) *convex* if $\text{epi}(g) := \{(\mathbf{x}, y) \in \mathbb{R}^{n+1} \mid g(\mathbf{x}) \leq y\}$ is a convex set (with epi denoting epigraph); (ii) *proper* if its value is never $-\infty$ and there exists at least one $\mathbf{x} \in \mathbb{R}^n$ such that $g(\mathbf{x}) < \infty$; (iii) *closed* if it is proper and $\text{epi}(g)$ is a closed set. We denote by ∂g the *subdifferential* of g . Moreover, we say that a map g is

1. *strongly convex with parameter $\rho > 0$* if the map $\mathbf{x} \mapsto g(\mathbf{x}) - \frac{\rho}{2} \|\mathbf{x}\|_2^2$ is convex;
2. *L-smooth* if it is differentiable and ∇g is L -Lipschitz.

The following characterization of Lipschitz continuity (see, e.g., [3]) is useful for our analysis.

Lemma 3. [*Lipschitz Continuity and Boundedness of Gradient*] *Let $f: \Theta \rightarrow \mathbb{R}$ be differentiable. Then f is L -Lipschitz if and only if $\|\nabla f(\theta)\|_{\Theta,*} \leq L$ for all $\theta \in \Theta$, where $\|\cdot\|_{\Theta,*}$ is the dual norm of $\|\cdot\|_{\Theta}$.*

Next, we define the proximal operator of g , which maps a point in $\mathbf{x} \in \mathbb{R}^n$ and into a subset of $\mathbb{R}^n$, which can be either empty, contain a single element, or be a set with multiple vectors.

Definition 4 (Proximal Operator). *Let $g: \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ be a proper, closed, and convex function, and let $\gamma > 0$. The proximal operator of g with parameter γ is the map $\text{prox}_{\gamma g}: \mathbb{R}^n \rightarrow \mathbb{R}^n$ defined by*

$$\text{prox}_{\gamma g}(\mathbf{x}) = \arg \min_{\mathbf{z} \in \mathbb{R}^n} g(\mathbf{z}) + \frac{1}{2\gamma} \|\mathbf{x} - \mathbf{z}\|_2^2, \quad \text{for all } \mathbf{x} \in \mathbb{R}^n.$$

For any convex, closed, and proper (CCP) function g , the proximal operator $\text{prox}_{\gamma g}(\mathbf{x})$ exists and is unique for all $\mathbf{x} \in \mathbb{R}^n$ [3, Th. 6.3]. Based on the use of proximal operators, *proximal gradient method* (see, e.g., [43]) can be devised to iteratively solve a class of composite (possibly non-smooth) convex problems of the form

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) + g(\mathbf{x}),$$

where $f: \mathbb{R}^n \rightarrow \mathbb{R}$, $g: \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ are CCP functions, and f is differentiable. At its core, the proximal gradient method updates the estimate of the solution of the optimization problem by computing the proximal operator of αg , where $\alpha > 0$ is a step size, evaluated at the difference between the current estimate and the gradient of αf computed at the current estimate. This method has been extended and generalized to a continuous-time framework [1, 25], resulting in the *continuous-time proximal gradient dynamics*

$$\dot{\mathbf{x}} = -\mathbf{x} + \text{prox}_{\gamma g}(\mathbf{x} - \gamma \nabla f(\mathbf{x})), \quad \gamma > 0. \tag{1}$$

Recent work [7] has provided an alternative interpretation of the dynamics [1], particularly in the case where f is a quadratic function and g is a separable sum of scalar functions. In this setting, the dynamics can be interpreted as a continuous-time firing-rate neural network. This observation is exploited in the results from the main text.

A Primer on Contraction Theory

Contraction theory [35, 5] is a powerful framework to study stability and convergence behaviors of dynamical systems. Traditionally, stability properties are defined in terms of convergence to an invariant set. These methods often require prior knowledge of the system's attractors, which can make them challenging to apply in scenarios – such as the one arising in GateMod – where this information is not known a-priori. Contraction theory focuses on the distance between trajectories. Intuitively, a system is contracting if trajectories rooted from different initial conditions converge towards each other exponentially. Beyond characterizing convergence, contractivity implies a number of desirable transient and asymptotic properties, also ensuring the existence of Lyapunov functions. Key properties are summarized at the end of this section.

We begin by recalling the following.

Definition 5 (Forward invariant set). *A set $\mathcal{A} \subseteq \mathbb{R}^n$ is forward invariant for a dynamical system $\dot{\mathbf{x}} = f(t, \mathbf{x})$ if every trajectory that starts in $\mathcal{A}$ remains in $\mathcal{A}$ for all future times. That is,*

$$\mathbf{x}_0 \in \mathcal{A} \implies \phi_t(\mathbf{x}_0) \in \mathcal{A}, \quad \text{for all } t \geq 0,$$

where $t \mapsto \phi_t(\mathbf{x}_0)$ denotes the flow map of the dynamical system with initial condition $\mathbf{x}_0 := \mathbf{x}(0)$.

Next, consider a dynamical system

$$\dot{\mathbf{x}}(t) = f(t, \mathbf{x}(t)), \quad (2)$$

where $f: \mathbb{R}_{\geq 0} \times \mathcal{C} \rightarrow \mathbb{R}^n$, is a smooth nonlinear function with $\mathcal{C} \subseteq \mathbb{R}^n$ being a forward invariant set for the dynamics. Given a norm $\|\cdot\|$ and a matrix $A \in \mathbb{R}^{n \times n}$, we recall that *logarithmic norm (log-norm)* of A , $\mu(A)$, is defined by $\mu(A) = \lim_{h \rightarrow 0^+} \frac{\|I_n + hA\| - 1}{h}$. As we shall see, our convergence result involve the log-norm associated with the Euclidean norm. This is $\mu_2(A) = \lambda_{\max}\left(\frac{A^\top + A}{2}\right)$, where $\lambda_{\max}(\cdot)$ denotes the largest eigenvalue of a matrix. We are now ready to give the following:

Definition 6 (Contracting systems). *Given a norm $\|\cdot\|$ with associated log-norm μ , a smooth function $f: \mathbb{R}_{\geq 0} \times \mathcal{C} \rightarrow \mathbb{R}^n$, with $\mathcal{C} \subseteq \mathbb{R}^n$ being forward invariant for the dynamics, open and convex, and a constant $c > 0$. Then, f is c -strongly infinitesimally contracting on $\mathcal{C}$ if*

$$\mu(Df(t, \mathbf{x})) \leq -c, \quad \text{for all } \mathbf{x} \in \mathcal{C} \text{ and } t \in \mathbb{R}_{\geq 0}, \quad (3)$$

where $Df(t, \mathbf{x}) := \partial f(t, \mathbf{x}) / \partial \mathbf{x}$ is the Jacobian of f with respect to $\mathbf{x}$.

In [16, Theorem 16] condition [3] is generalized for locally Lipschitz function. One of the benefits of contraction theory is that it characterizes convergence with a single condition. In fact, if f is c -strongly infinitesimally contracting, then for any two trajectories $\mathbf{x}(\cdot)$ and $\mathbf{y}(\cdot)$ of the dynamics [2] it holds that

$$\|\phi_t(\mathbf{x}_0) - \phi_t(\mathbf{y}_0)\| \leq e^{-ct} \|\mathbf{x}_0 - \mathbf{y}_0\|, \quad \text{for all } t \geq 0,$$

i.e., the distance between the two trajectories converges exponentially with rate c . This constant quantifies the convergence rate between trajectories and it is often termed *contraction rate*.

Strongly infinitesimally contracting dynamical systems exhibit highly ordered transient and asymptotic behavior, making them particularly advantageous for studying their convergence. Namely, (i) initial conditions are exponentially forgotten [35]; (ii) for time-invariant dynamics, there exists a unique globally exponential stable equilibrium [35]. Additionally, two natural Lyapunov functions are automatically available: the distance from the equilibrium and norm of the vector field itself; (iii) contraction ensures entrainment to periodic inputs [50] and implies robustness properties such as input-to-state stability, also for delayed dynamics [15, 61]; (iv) contracting systems enjoy equilibrium tracking properties [14]; (v) contraction theory is a modular framework, enabling the stability of interconnected systems to be derived from the properties of their components [51]. Moreover, (vi) efficient numerical algorithms can be devised for numerical integration and fixed point computation of contracting systems [6].

3 Entropy-regularized Optimization and Softmax Gradient Flows

Here, we consider a general class of entropy-regularized minimization problems of the form

$$\min_{\mathbf{x} \in \Delta_n} F(\mathbf{x}) - \varepsilon H(\mathbf{x}), \quad (4)$$

where $F: \mathbb{R}^n \rightarrow \mathbb{R}$ is convex and continuously differentiable, and $\varepsilon > 0$ is a regularization parameter sometimes referred to as *temperature*. This setting embeds as special case GateFrame optimization. Our goal is to obtain a continuous-time dynamical system that provably converges to the optimal solution of Eq. [4]. For reasons that will become apparent in this section, we term this dynamics *softmax gradient flow*. The formal derivations for GateFrame, GateFlow, and GateNet (Sec. 4) build on the results reported here.

Entropy-regularized optimization problems of the form [4] naturally arise across reinforcement learning, game theory, optimal transport, and statistical physics [38, 39, 45, 53]. This class of problems also appears in the form of smooth best response maps, which indeed leads to the logit (i.e., softmax) function as a solution [10, 38, 19]. Finally, optimization over the probability simplex with entropy regularization is also closely related to mirror descent, where the negative entropy acts as the mirror map [4, 3].

An optimal solution for the problem in Eq. [4] exists and is unique. This is discussed in the next:

Remark 1 (Existence and uniqueness of the minimizer of [4]). *Consider the entropy-regularized problem [4]. Since the set Δ_n is compact and the map $x_i \mapsto -x_i \ln(x_i)$ extends continuously to $x_i = 0$ (under the standard convention $0 \ln(0) = 0$), the objective function is continuous on Δ_n . Hence, by the Weierstrass Theorem, there exists a minimizer. Regarding uniqueness, note that the regularization term $-\varepsilon H(\mathbf{x})$ acts as a logarithmic barrier: its gradient satisfies $\partial(-x_i \ln x_i)/\partial x_i \rightarrow +\infty$ as $x_i \rightarrow 0$. Therefore, no minimizer can lie on the boundary of the simplex, and every minimizer must belong to the relative interior of Δ_n where the map $-\varepsilon H(\mathbf{x})$ is strictly convex. Consequently, the cost function $F(\mathbf{x}) - \varepsilon H(\mathbf{x})$ is strictly convex on the feasible region for minimizers. It follows that the entropy-regularized problem admits a unique minimizer. In particular, uniqueness is guaranteed whenever F is convex and continuously differentiable on a neighborhood of Δ_n , even if F is not strictly convex.*

Equipped with this observation, we can now derive that a dynamical system that provably converges to the optimal solution. Namely, we first propose a reformulation of the optimization problem that enables

computing a proximal operator. Then, obtain the corresponding continuous-time proximal dynamics, the softmax gradient flow, and characterize its properties. Namely, we show that the equilibrium of this dynamics is the optimal solution of the problem in Eq. [4] and trajectories provably converge to the equilibrium.

3.1 Reformulating Problem [4]

We introduce an unconstrained reformulation of the problem in Eq. [4]. This is obtained by defining the *entropic barrier function* $H_{\text{barrier}}: \mathbb{R}^n \rightarrow]-\infty, +\infty[$ as $H_{\text{barrier}}(\mathbf{x}) := -H(\mathbf{x}) + \iota_{\Delta_n}$. This function is equal to the entropy in the simplex – the feasibility domain in Eq. [4] – and infinity outside of this set. Note that H_{barrier} is convex, closed and proper, being sum of CCP functions. With this function, problem [4] can be recast as

$$\min_{\mathbf{x} \in \mathbb{R}^n} F(\mathbf{x}) + \varepsilon H_{\text{barrier}}(\mathbf{x}). \quad (5)$$

To obtain the softmax gradient flow associated to Problem [5], we add and subtract $\varepsilon \frac{\|\mathbf{x}\|^2}{2}$ to its cost. This motivates the following result, in which we compute the proximal operator of $H_{\text{barrier}}(\mathbf{x}) - \frac{\|\mathbf{x}\|^2}{2}$. While this result can be proved via an adaptation of [8, Example 2.23], where tools from monotone operator theory are used, we provide a more direct and accessible justification in Sec. 6 through an alternative proof based directly on the definition of the proximal operator.

Lemma 4. *Given the function $H_{\text{barrier}}: \mathbb{R}^n \rightarrow]-\infty, +\infty[$, we have $\text{prox}_{H_{\text{barrier}} - \frac{\|\cdot\|^2}{2}}(\mathbf{x}) = \text{softmax}(\mathbf{x})$.*

Beyond being instrumental to obtain a dynamical system converging to the optimal solution of the problem in Eq. [4], Lemma 4 also reveals a link between proximal operators and smooth best response maps. As discussed in [10, 19], in the context of game theory, the softmax function can be derived by considering the *argmax function* under entropy regularization resulting in what is called *smooth best response maps*. More specifically, let $\mathbf{x} \in \mathbb{R}^n$ and consider the argmax function or best response map (the agent seeks the strategy with the highest score)

$$\mathbf{x} \mapsto \arg \max_{\mathbf{z} \in \Delta_n} \mathbf{x}^\top \mathbf{z}.$$

This best response map carries several challenges (see, e.g., [10, 19] for more details). For example, if two or more components are equal, it leads to multi-valued function, while in many applications it is highly desirable to have a singled-valued map (i.e., having a unique maximizer). To address this, the most common approach is to introduce a regularizer h that acts as a penalty (or *control cost*) to the maximization objective. This yields the *regularized argmax function* or *smooth best response map*

$$\mathbf{x} \mapsto \arg \max_{\mathbf{z} \in \Delta_n} \left\{ \mathbf{x}^\top \mathbf{z} - h(\mathbf{z}) \right\}.$$

A popular choice for the regularizer is the negative entropy function restricted to the simplex. This choice leads to the softmax, also known as *logit map* in the context of game theory, that is, $\arg \max_{\mathbf{z} \in \Delta_n} \left\{ \mathbf{x}^\top \mathbf{z} + H(\mathbf{z}) \right\} = \text{softmax}(\mathbf{x})$.

3.2 The Softmax Gradient Flow

To solve the optimization problem [5] we propose the following *softmax gradient flow*:

$$\tau \dot{\mathbf{x}} = -\mathbf{x} + \text{softmax}(-\varepsilon^{-1} \nabla F(\mathbf{x})) = F_{\text{sm}}(\mathbf{x}). \quad (6)$$

This dynamics is motivated an explicit connection between the entropy-regularized minimization problem [4] and the softmax gradient flow [6] formalized in the next Lemma 5. Before introducing the lemma, we make the following observations:

1. as the parameter ε increases, the dynamics [6] promote solutions with higher entropy. This follows from the fact that as $\varepsilon \rightarrow +\infty$ the softmax converges to the uniform distribution.
2. as ε goes to zero, the dynamics [6] converge to a trajectory that follows the steepest descent direction along the simplex vertices. Specifically, as $\varepsilon \rightarrow 0^+$, the softmax function converges to the vector $e_{i^*(\mathbf{x})}$, where $i^*(\mathbf{x}) \in \arg \min_{1 \leq j \leq n} \nabla F(\mathbf{x})_j$. Thus, the dynamics become $\dot{\mathbf{x}} \approx -\mathbf{x} + e_{i^*(\mathbf{x})}$, which means that the trajectory moves toward the vertex of the simplex corresponding to the coordinate of steepest descent of the function F on the simplex. This type of dynamics is reminiscent of the Frank–Wolfe algorithm [18], where at each step the update direction is toward an extreme point (vertex) of the feasible set that minimizes the linear approximation of the objective function.

The next result formalizes that the softmax gradient flow: (i) is the proximal dynamics associated to problem [4]; (ii) encodes the optimal solution of problem [4] as an equilibrium; (iii) is an energy model.

Lemma 5 (Relations between [4] and [6]). *Let $F: \mathbb{R}^n \rightarrow \mathbb{R}$ be a convex and continuously differentiable function, and let $\varepsilon > 0$. Consider the entropy-regularized optimization problem [4] and the softmax gradient flow [6]. Then*

1. *the softmax gradient flow [6] is the proximal gradient flow associated with the composite optimization problem [4];*
2. *a vector $\mathbf{x}^* \in \mathbb{R}^n$ is an equilibrium point of [6] if and only if $\mathbf{x}^*$ is an optimal solution of [4];*
3. *let $\mathbf{x} \in \mathbb{R}^n$ and let $\phi_t(\mathbf{x}_0)$ be the flow map of the proximal gradient dynamics [6]. Then the map $t \mapsto F(\phi_t(\mathbf{x}(t))) + \varepsilon H_{\text{barrier}}(\phi_t(\mathbf{x}(t)))$ is non-increasing.*

The proof is given in Sec. 6. Here, we note that item 3 of Lemma 5 not only reveals that the softmax gradient flow is an energy, but it also defines an energy function. This is the cost of problem [5], more precisely

$$V(\mathbf{x}) := F(\mathbf{x}) + \varepsilon H_{\text{barrier}}(\mathbf{x}).$$

This in turn means that for, for every initial condition $\mathbf{x}_0$, $V(\mathbf{x})$ is non-increasing along the trajectories $\phi_t(\mathbf{x}_0)$ of the softmax gradient flow. With the results in the next section we rigorously show that trajectories of the softmax converge to the equilibrium in item 2, this is an energy minimum and the optimal solution of [4].

3.3 Properties And Convergence of The Softmax Gradient Flow

Remarkably, the softmax gradient flow is a contracting system (Theorem 1). This not only implies exponential convergence to its equilibrium (which is also the optimal solution of [4]) but also an explicit

guaranteed convergence rate, together with desirable implications (Corollary 1). We start by showing that any trajectory of the softmax gradient flow with initial condition in the simplex never leave this set. Thus, trajectories starting from a feasible point for Eq. [4] remain feasible for all time.

Lemma 6 (Invariance property). *The probability simplex Δ_n is a forward invariant set for the softmax gradient flow [6].*

The formal proof is in Sec. 6. Here, we introduce our convergence result.

Theorem 1 (Contractivity of [6]). *Let $F: \mathbb{R}^n \rightarrow \mathbb{R}$ be convex and L_F -smooth. The softmax gradient flow [6] is strongly infinitesimally contracting with respect to the norm $\|\cdot\|_2$ with rate $c = \frac{1}{\tau}$.*

Theorem 1 reveals that – in a setting that applies to GateMod as we shall see in Sec. 4 – the softmax gradient flow is contracting. Since contractivity implies that any two solutions converge towards each other, this means that, for initial conditions in the simplex, solutions of the softmax gradient flow not only remain inside the simplex (by Lemma 6) but also that trajectories converge to the equilibrium, $\mathbf{x}^*$. As this is also the optimal solution of Eq. [4], this means that the trajectories contract towards the optimal solution of the problem. Convergence is exponential and the contraction rate (that is, the rate of convergence towards the optimal solution) is τ^{-1} and hence it does not depend on F . Next, we summarize some key implications of the contractivity of the softmax gradient flow. More precisely, we establish that the softmax gradient flow has two Lyapunov (and hence energy) functions and assess the behavior of the system in response to time variations in the vector field. The result also establishes that the softmax gradient flow is computationally friendly, in the sense that it remains contracting if it is discretized with a suitable step-size.

Corollary 1. *Let $F: \mathbb{R}^n \rightarrow \mathbb{R}$ be convex and L_F -smooth and consider the softmax gradient flow [6]. Then,*

1. *if F_{sm} is time invariant, then every solution of [6] globally exponentially converge to the unique equilibrium $\mathbf{x}^* \in \Delta_n$ with rate $c = \tau^{-1}$. Additionally, two natural Lyapunov functions are automatically available:*

$$\mathbf{x} \mapsto \|\mathbf{x} - \mathbf{x}^*\|, \quad \text{and} \quad \mathbf{x} \mapsto \|F_{\text{sm}}(\mathbf{x})\|;$$

2. *if ∇F is T -periodic, then there exists a unique periodic solution with period T ;*
3. *the forward Euler discretization, $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha F_{\text{sm}}(\mathbf{x}_k)$, is strongly contracting for every step size $\alpha \in \left]0, \frac{2\tau}{(1+\varepsilon^{-1}L_F)^2}\right[$. Moreover, the step size that minimize the discrete-time contraction rate is $\alpha^* = \frac{\tau}{(1+\varepsilon^{-1}L_F)^2}$.*

We conclude by remarking that, in addition to the properties listed in Corollary 1, contractive systems exhibit several other remarkable behaviors that make this form of stability desirable for optimization solvers and neural implementations. These include stability properties in response to, e.g., disturbances and/or inputs, as well as robustness under delayed dynamics [55, 15, 61], implying that its trajectories cannot go arbitrarily far from each other. Moreover, contracting systems enjoy equilibrium-tracking properties [14] when the dynamics depend on time-varying parameters.

3.4 Introducing biases

The derivations in this section naturally extend to optimization problems of the form:

$$\min_{\mathbf{x} \in \Delta_n} F(\mathbf{x}) + \varepsilon D_{\text{KL}}(\mathbf{x} \parallel \hat{\mathbf{x}}),$$

where $\hat{\mathbf{x}} \in \Delta_n$ is a given vector. This formulation introduces a bias for the optimal solution. Essentially, the entropy is replaced with the KL divergence and this biases the optimal solution towards the vector $\hat{\mathbf{x}}$. The formal treatment presented in this section extends to this setting. In fact, one can rewrite the KL divergence regularized problem as $\min_{\mathbf{x} \in \Delta_n} (\bar{F}_\varepsilon(\mathbf{x}) - \varepsilon H(\mathbf{x}))$, where $\bar{F}_\varepsilon(\mathbf{x}) := F(\mathbf{x}) + \varepsilon H(\mathbf{x}, \hat{\mathbf{x}})$ is again convex and continuously differentiable in $\mathbf{x}$.

4 GateMod details

We describe the formal details for computational model. We first present the derivations for GateMod components: GateFrame, GateFlow, and GateNet. Then, we show why – as remarked in the main text – GateMod yields Gumbel-softmax gating when a weight bias is introduced in the formulation. We recall from the main text that $\mathbf{X}_k$ and $\mathbf{U}_k$ are the random variables of the state and action at time-step k . The state space is $\mathcal{X} \subseteq \mathbb{R}^{n_x}$ and the action space is $\mathcal{U} \subseteq \mathbb{R}^{n_u}$. The time-indexing is chosen so that the environment $p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ transitions from state $\mathbf{x}_{k-1}$ to $\mathbf{x}_k$ when action $\mathbf{u}_k$ is applied. The agent selects its action by sampling from the policy $p(\mathbf{u}_k | \mathbf{x}_{k-1})$ so that the agent-environment dynamics is captured by $p(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}) = p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) p(\mathbf{u}_k | \mathbf{x}_{k-1})$.

GateFrame

For convenience, we report here GateFrame optimization from the main text:

$$\begin{aligned} \min_{\mathbf{w}_k \in \Delta_{n_\pi}} D_{\text{KL}}(p(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1}) || q(\mathbf{x}_k, \mathbf{u}_k | \mathbf{x}_{k-1})) - \varepsilon H(\mathbf{w}_k) \\ \text{s.t. } p(\mathbf{u}_k | \mathbf{x}_{k-1}) = \sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1}). \end{aligned} \quad (7)$$

All symbols are defined in the main text and here, we use the shorthand notation $f(\mathbf{w}_k)$ to denote the cost in GateFrame optimization and formalize the mild assumptions described in the main text:

Assumption 1. *The set $\mathcal{W}_k := \{\mathbf{w}_k | \mathbf{w}_k \text{ is feasible for GateFrame optimization and } f(\mathbf{w}_k) < +\infty\}$ is non-empty.*

Assumption 2. *The primitives are uniformly bounded and have full support over the action space $\mathcal{U}$. Specifically, there exist two finite constants $0 < \pi_{\min} \leq \pi_{\max}$ such that, for all $\alpha \in \{1, \dots, n_\pi\}$ and for all time steps k , $\pi_{\min} \leq \pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1}) \leq \pi_{\max}$.*

Intuitively, Assumption 1 makes the optimization meaningful, in the sense that the optimal value of GateFrame optimization is finite. This assumption is standard in the context of, e.g., entropy-regularized optimization. For example, the assumption is satisfied when: (i) $p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ is absolutely continuous with respect to $q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$; (ii) the primitives are absolutely continuous with respect to $q(\mathbf{u}_k | \mathbf{x}_{k-1})$. Assumption 2 formalizes the fact that primitives cover the action space and have bounded moments (this is always satisfied for, e.g., discrete variables or Gaussians).

Next, we unpack the derivations leading to the unconstrained reformulation of GateFrame optimization, also reported here for convenience

$$\min_{\mathbf{w}_k \in \mathbb{R}^{n_\pi}} F(\mathbf{w}_k) + \varepsilon H_{\text{barrier}}(\mathbf{w}_k). \quad (8)$$

The next result derives the functional expression for F given in Methods, shows the equivalence of [8] with GateFrame and establishes a number of useful properties.

Lemma 1 (GateFrame reformulation). *Let Assumption 1 hold. Then,*

1. *the minimization problem [8] is a reformulation of GateFrame when*

$$F(\mathbf{w}_k) := \sum_{\alpha=1}^{n_\pi} w_k^\alpha \left(\mathbb{E}_{\pi^\alpha(\mathbf{u}_k|\mathbf{x}_{k-1})} \left[\ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) - \ln q(\mathbf{u}_k | \mathbf{x}_{k-1}) + D_{\text{KL}}(p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)) - \ln q(\mathbf{u}_k | \mathbf{x}_{k-1}) \right] + \mathbb{1}_{n_\pi} \right) \quad (9)$$

2. *the gradient of $F(\mathbf{w}_k)$ is:*

$$\nabla F(\mathbf{w}_k) = \mathbb{E}_{\pi(\mathbf{u}_k|\mathbf{x}_{k-1})} \left[\ln \mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1}) + D_{\text{KL}}(p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)) - \ln q(\mathbf{u}_k | \mathbf{x}_{k-1}) \right] + \mathbb{1}_{n_\pi}$$

where we are using the notation $\mathbb{E}_{\pi(\mathbf{u}_k|\mathbf{x}_{k-1})} = (\mathbb{E}_{\pi^\alpha(\mathbf{u}_k|\mathbf{x}_{k-1})})_{\alpha \in \{1, \dots, n_\pi\}} \in \mathbb{R}^{n_\pi}$;

3. *the function F is convex and L_F -smooth in $\mathbf{w}_k$ with constant $L_F = n_\pi \frac{\pi_{\max}^2}{\pi_{\min}}$.*

The proof is in Sec. 6. Here, we note that item 1 shows the reformulation in Results from the main text and derives the functional expression for F given in Methods. Item 2 computes the gradient of F and this expression is used in Results and Methods when we obtain GateFlow. Finally, item 3 characterizes convexity and smoothness of the problem. Inspecting the proof of item 3 reveals that, under a mild condition F is actually strictly convex. In the proof, convexity is assessed by verifying the Hessian of F is positive semi-definite, evaluating the quadratic form $\mathbf{z}^\top \nabla^2 F(\mathbf{w}_k) \mathbf{z}$ for a generic vector $\mathbf{z} \in \mathbb{R}^{n_\pi}$, $\mathbf{z} \neq \mathbf{0}_{n_\pi}$. In the proof, we rewrite the quadratic form as

$$\int \frac{(\mathbf{z}^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1}))^2}{\mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1})} d\mathbf{u}_k,$$

and show that this term is in fact non-negative. However, in order for the integral in the right-hand side to be 0, there should exist a single vector, $\mathbf{z}$, that makes the linear combination of primitives in the integral identically 0 over the whole action space, an unlikely situation. Lemma 1 is leveraged next to derive GateFlow.

GateFlow

Building on the formal treatment developed so far, GateFlow is the softmax gradient flow associated to GateFrame optimization and its dynamics is:

$$\tau \dot{\mathbf{w}}_k = -\mathbf{w}_k + \text{softmax} \left(-\varepsilon^{-1} \mathbb{E}_{\pi(\mathbf{u}_k|\mathbf{x}_{k-1})} \left[\ln \mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1}) + D_{\text{KL}}(p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)) - \ln q(\mathbf{u}_k | \mathbf{x}_{k-1}) \right] \right) \quad (10)$$

or, component-wise:

$$\tau \dot{w}_k^\alpha = -w_k^\alpha + \text{softmax} \left(-\varepsilon^{-1} \mathbb{E}_{\pi(\mathbf{u}_k|\mathbf{x}_{k-1})} \left[\ln \mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1}) + D_{\text{KL}}(p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)) - \ln q(\mathbf{u}_k | \mathbf{x}_{k-1}) \right] \right)$$

with $\alpha \in \{1, \dots, n_\pi\}$. This is the dynamics from the main text, where we unpack the exponent of the softmax and we leverage the invariance property (Lemma 1) to drop the constant term $-\mathbb{1}_{n_\pi}$ in the gradient.

In the main text we explicitly discuss the setting where the action space is discrete so that there are d_u possible actions. In this case, it is possible to define the matrix $\Pi(\mathbf{x}_{k-1}) \in \mathbb{R}^{d_u \times n_\pi}$

$$\Pi(\mathbf{x}_{k-1}) := \begin{pmatrix} \pi^1(\mathbf{u}_{k,1} | \mathbf{x}_{k-1}) & \pi^2(\mathbf{u}_{k,1} | \mathbf{x}_{k-1}) & \dots & \pi^{n_\pi}(\mathbf{u}_{k,1} | \mathbf{x}_{k-1}) \\ \pi^1(\mathbf{u}_{k,2} | \mathbf{x}_{k-1}) & \pi^2(\mathbf{u}_{k,2} | \mathbf{x}_{k-1}) & \dots & \pi^{n_\pi}(\mathbf{u}_{k,2} | \mathbf{x}_{k-1}) \\ \vdots & \vdots & \ddots & \vdots \\ \pi^1(\mathbf{u}_{k,d_u} | \mathbf{x}_{k-1}) & \pi^2(\mathbf{u}_{k,d_u} | \mathbf{x}_{k-1}) & \dots & \pi^{n_\pi}(\mathbf{u}_{k,d_u} | \mathbf{x}_{k-1}) \end{pmatrix} \quad (11)$$

and consequently write GateFlow dynamics as

$$\tau \dot{\mathbf{w}}_k = -\mathbf{w}_k + \text{softmax} \left(-\varepsilon^{-1} \Pi(\mathbf{x}_{k-1})^\top (\ln(\Pi(\mathbf{x}_{k-1}) \mathbf{w}_k) + \mathbf{c}(\mathbf{x}_{k-1}, \mathbf{u}_k)) \right), \quad (12)$$

where the input $\mathbf{c}$ is defined as in the main text as $\mathbf{c}(\mathbf{x}_{k-1}, \mathbf{u}_k) := D_{\text{KL}}(p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)) - \ln q(\mathbf{u}_k | \mathbf{x}_{k-1})$.

All the desirable GateFlow properties from the main text follow from the fact that this dynamics is the softmax gradient flow associated to GateFrame. In particular:

1. the probability simplex Δ_n is a forward invariant set for GateFlow. This means that, if the dynamics is initialized with initial conditions that are feasible for GateFrame, then GateFlow trajectories remain feasible;
2. an equilibrium exists for GateFlow. Moreover, a point in the state space is an equilibrium for GateFlow if and only if it is also GateFrame optimal solution;
3. finally, GateFlow is strongly infinitesimally contracting and the contraction rate is $c = \tau^{-1}$. This means that for any initial condition that is feasible for GateFrame, GateFlow trajectories converge exponentially to optimal solution and the guaranteed convergence rate, τ^{-1} , is independent on $\mathbf{F}$.

The properties imply that GateFlow is an energy model, featuring the two energy functions from the main text. The first energy function is GateFrame cost. The fact that this function decreases along GateFlow trajectories follows from the fact that GateFlow is a proximal gradient flow and $\mathbf{F}$ is at least convex. The second energy function reported in the text is the Euclidean distance from GateFrame optimal solution. This function is exponentially decreasing, and this follows from contractivity.

GateNet

We provide the detailed derivations for GateNet from the main text. We derive separately the fast and slow units (Fig. 2C in the main text) and then integrate them together.

We recall that the matrix of the primitives is conditioned on the current state, available to the agent. In the main text, this aspect is highlighted by explicitly writing $\Pi(\mathbf{x}_{k-1})$. Here, we denote by $\Pi_i(\mathbf{x}_{k-1})$ the i -th row of the matrix $\Pi(\mathbf{x}_{k-1})$, for all $i \in \{1, \dots, d_u\}$, and by $\Pi^\alpha(\mathbf{x}_{k-1})$ is the α -th column of $\Pi(\mathbf{x}_{k-1})$, for all $\alpha \in \{1, \dots, n_\pi\}$.

Fast unit. As detailed in the main text, the goal of the fast unit is to compute the vector $\mathbf{y} = -\varepsilon^{-1} \nabla \mathbf{F}(\mathbf{w}_k)$.

Component-wise:

$$\begin{aligned}
y^\alpha &:= -\varepsilon^{-1} (\nabla F(\mathbf{w}_k))_\alpha = -\sum_{i=1}^{d_u} (\ln \Pi_i(\mathbf{x}_{k-1}) \mathbf{w}_k + c(\mathbf{x}_{k-1}, \mathbf{u}_{k,i})) \pi^\alpha(\mathbf{u}_{k,i} | \mathbf{x}_{k-1}) \\
&= -\varepsilon^{-1} \sum_{i=1}^{d_u} (\ln a^i + c(\mathbf{x}_{k-1}, \mathbf{u}_{k,i})) \pi^\alpha(\mathbf{u}_{k,i} | \mathbf{x}_{k-1}) \tag{13}
\end{aligned}$$

$$\begin{aligned}
&= -\varepsilon^{-1} \sum_{i=1}^{d_u} (b^i + c(\mathbf{x}_{k-1}, \mathbf{u}_{k,i})) \pi^\alpha(\mathbf{u}_{k,i} | \mathbf{x}_{k-1}) \tag{14} \\
&= -\varepsilon^{-1} \sum_{i=1}^{d_u} b^i \pi^\alpha(\mathbf{u}_{k,i} | \mathbf{x}_{k-1}) - \sum_{i=1}^{d_u} c(\mathbf{x}_{k-1}, \mathbf{u}_{k,i}) \pi^\alpha(\mathbf{u}_{k,i} | \mathbf{x}_{k-1}) \\
&= -\varepsilon^{-1} (\Pi^\alpha(\mathbf{x}_{k-1}))^\top \mathbf{b} - (\Pi^\alpha(\mathbf{x}_{k-1}))^\top \mathbf{c}_k,
\end{aligned}$$

where, as in the main text, $\mathbf{a}$ is the stack of $a^i := \Pi_i(\mathbf{x}_{k-1}) \mathbf{w}_k$ and $\mathbf{b}$ is the stack of $b^i := \ln(a_i)$. Letting $\mathbf{b} := [b_1, \dots, b_{d_u}]^\top$, and $\mathbf{c}_k := [c_1, \dots, c_{d_u}]^\top$, then the vector $\mathbf{y}$ can be computed via a dynamical system with: (i) two set of equations, say with state variables $\mathbf{a}$ and $\mathbf{b}$, that converge to the equilibrium $\bar{\mathbf{a}} = \Pi(\mathbf{x}_{k-1}) \mathbf{w}_k$ and $\bar{\mathbf{b}} = \ln(\mathbf{a})$, respectively; (ii) a slower set of dynamic equations, say with state variables $\mathbf{y}$, that implements the formula in the last line of the above derivations. This yields the dynamics:

$$\begin{aligned}
\tau_g \dot{\mathbf{a}} &= -\mathbf{a} + \Pi \mathbf{w}_k, \\
\tau_g \dot{\mathbf{b}} &= -\mathbf{b} + \ln(\mathbf{a}), \\
\tilde{\tau}_g \dot{\mathbf{y}} &= -\varepsilon \mathbf{y} - \Pi^\top (\mathbf{b} + \mathbf{c}_k),
\end{aligned}$$

with $\tau_g \ll \tilde{\tau}_g$. This is the dynamics from the main text corresponding to the circuit for the fast in unit in Fig. 2C.

Slow unit. This unit in Fig. 2C (main text) is the biologically plausible softmax neural circuit from the literature [54]. Complementing the derivations in Methods, the unit receives as input the gradient the components of the vector $\mathbf{y}$, y^α and returns the vector $\mathbf{r}$ having as components $w_k^\alpha = \text{softmax}(\mathbf{y})_\alpha$. In the circuit from the main text, neuron m computes the softmax normalizer, each of the $\mathbf{r}$ -neurons returns $y_\alpha - \ln(m)$ so that the $\mathbf{w}_k$ -neurons effectively return the softmax, providing GateFrame optimal solution. This yields the dynamics in Results, also reported here

$$\begin{aligned}
\tau_s \dot{m} &= -m + \sum_{\alpha=1}^{n_\pi} e^{y_\alpha}, \\
\tau_s \dot{r}^\alpha &= -r^\alpha + y^\alpha - \ln(m), \\
\tau \dot{w}_k^\alpha &= -x^\alpha + e^{r^\alpha},
\end{aligned} \tag{15}$$

where $\alpha = 1, \dots, n_\pi$ and $\tau_s \ll \tau$ to ensure that the dynamics for neuron m and $\mathbf{r}$ -neurons is faster than the output dynamics returning the weights.

GateNet dynamics. The dynamics in the main text, corresponding to the full circuit in Fig. 2C, is obtained combining the dynamics for the fast and slow units, with $\tilde{\tau}_g \ll \tau_s$ so that the former dynamics is effectively faster than the latter.

Why introducing biases in GateMod yields Gumbel-softmax gating

Following the arguments at the end of Sec. 3, biases can be introduced in GateFrame by replacing the entropic regularizer with the KL divergence. GateFrame optimization becomes:

$$\begin{aligned} \min_{\mathbf{w}_k \in \Delta_{n_\pi}} \quad & D_{\text{KL}}(p(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1}) \parallel q(\mathbf{x}_k, \mathbf{u}_k \mid \mathbf{x}_{k-1})) + \varepsilon D_{\text{KL}}(\mathbf{w}_k \parallel \hat{\mathbf{w}}_k) \\ \text{s.t.} \quad & p(\mathbf{u}_k \mid \mathbf{x}_{k-1}) = \sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha(\mathbf{u}_k \mid \mathbf{x}_{k-1}), \end{aligned} \quad (16)$$

where $\hat{\mathbf{w}}_k$ is a bias vector. We have seen that the identity $D_{\text{KL}}(\mathbf{w}_k \parallel \hat{\mathbf{w}}_k) = \mathbf{H}(\mathbf{w}_k, \hat{\mathbf{w}}_k) - \mathbf{H}(\mathbf{w}_k)$ holds and the problem remains convex and the same derivations presented above can be followed. By doing so, when the bias $\hat{\mathbf{w}}_k$ is introduced, Eq. [12] is modified into

$$\tau \dot{\mathbf{w}}_k = -\mathbf{w}_k + \text{softmax}(\varepsilon^{-1}(\ln \hat{\mathbf{w}}_k + \mathbf{g}(\mathbf{w}_{k-1}))) . \quad (17)$$

In this case, GateFlow equilibrium is still the optimal GateFrame solution and is equal to

$$\mathbf{w}_k^* = \text{softmax}(\varepsilon^{-1}(\ln \hat{\mathbf{w}}_k + \mathbf{g}(\mathbf{w}_{k-1}^*))) , \quad (18)$$

where $\mathbf{g}$ is the vector $-\Pi(\mathbf{x}_{k-1})^\top (\ln(\Pi(\mathbf{x}_{k-1})\mathbf{w}_k) + \mathbf{c}(\mathbf{x}_{k-1}, \mathbf{u}_k))$ from Eq. [12]. The above expression is a Gumbel-softmax distribution [29]. This distribution not only naturally arises as a gating mechanism in dense mixture-of-experts architectures [26], but also in the context of inference [40, Chapter 6] and in policy computation methods based on the minimization of free-energy type functionals. In fact, the right hand side in Eq. [18] can be written component-wise as

$$\frac{\hat{w}_k^\alpha \exp(\varepsilon^{-1} g_\alpha(w_k^\alpha))}{\sum_{\beta=1}^{n_\pi} \hat{w}_k^\beta \exp(\varepsilon^{-1} g_\beta(w_k^\beta))}.$$

This is a probability mass function that twists the probability vector $\hat{\mathbf{w}}_k$ with an exponential [24] kernel appearing in policies that minimize the variational free energy, including KL control. See, e.g., the surveys in [53, 20]. Moreover, Eq. [17] also reveals that GateNet circuit does not necessarily need to be modified to account for biases. These can be included in the input $\mathbf{c}$ to the network.

5 Supplementary Details for the Experiments

In the experiments reported in the main text, the optimal GateMod weights were computed via GateFlow, after verifying on subset of experiments that it closely matches GateNet trajectories.

5.1 Collective Behaviors Experiments

We provide supplementary details for the experiments results and refer to Tab. 1 for the values of the parameters. We recall from the main text that the state of boid (agent) i , $\mathbf{x}_k^i$, is the stack of the boid's position $\mathbf{p}_k^i = [p_{x,k}^i, p_{y,k}^i]^\top$ and velocity $\mathbf{v}_k^i = [v_{x,k}^i, v_{y,k}^i]^\top$. The maximum velocity is bounded; more precisely, $\|\mathbf{v}_k^i\| \leq v_{\max}$. The action of boid i , $\mathbf{u}_k^i = [u_{x,k}^i, u_{y,k}^i]^\top$, is the boid acceleration. This is also bounded, so that $\|\mathbf{u}_k^i\| \leq u_{\max}$. The dynamics of boid i , $p(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i, \mathbf{u}_k^i)$, is the Gaussian defined in Methods with

mean given by the discrete-time second order model [13, 33, 42]:

$$\begin{cases} \mathbf{p}_k^i &= \mathbf{p}_{k-1}^i + \mathbf{v}_{k-1}^i dt \\ \mathbf{v}_k^i &= \mathbf{v}_{k-1}^i + \mathbf{u}_k^i dt \end{cases}, \quad i = 1, \dots, N. \quad (19)$$

In Fig. 3A, a neighboring boid j is visible to boid i if the relative angle $\theta_{ij,k} = \arccos \left(\frac{(\mathbf{v}_k^i)^\top (\mathbf{p}_k^j - \mathbf{p}_k^i)}{\|\mathbf{v}_k^i\| \|\mathbf{p}_k^j - \mathbf{p}_k^i\|} \right)$ lies in the cone $[-\frac{\alpha}{2}, \frac{\alpha}{2}]$. The three concentric zones in Fig. 3A are from the literature, see [11] and related references in Results. The mean of the Gaussian primitives in the main text is obtained from the social forces widely reported in the literature [47, 11]. In particular, following [36], in the main text the centers of the Gaussians are:

- $\mathbf{u}_{\text{sep}}^i$ for the separation primitive is

$$\mathbf{u}_{\text{sep}}^i = \begin{cases} \frac{1}{|\mathcal{S}_k^i|} \sum_{j \in \mathcal{S}_k^i} g(\|\mathbf{p}_k^i - \mathbf{p}_k^j\|) \frac{\mathbf{p}_k^i - \mathbf{p}_k^j}{\|\mathbf{p}_k^i - \mathbf{p}_k^j\|}, & \text{if } |\mathcal{S}_k^i| > 0, \\ 0_2, & \text{otherwise,} \end{cases}$$

with $\mathcal{S}_k^i$ being the set of neighbors visible to boid i in the separation zone;

- $\mathbf{u}_{\text{ali}}^i$ for alignment primitive is

$$\mathbf{u}_{\text{ali}}^i = \begin{cases} \frac{1}{|\mathcal{A}_k^i|} \sum_{j \in \mathcal{A}_k^i} g(\|\mathbf{p}_k^i - \mathbf{p}_k^j\|) \frac{\mathbf{v}_k^j}{\|\mathbf{v}_k^j\|}, & \text{if } |\mathcal{A}_k^i| > 0, \\ u_{\text{max}} \frac{\mathbf{v}_k^i}{\|\mathbf{v}_k^i\|}, & \text{otherwise,} \end{cases}$$

where $\mathcal{A}_k^i$ is the set of neighbors visible to boid i in the alignment zone;

- $\mathbf{u}_{\text{coh}}^i$ for the cohesion primitive is

$$\mathbf{u}_{\text{coh}}^i = \begin{cases} \frac{1}{|\mathcal{C}_k^i|} \sum_{j \in \mathcal{C}_k^i} g(\|\mathbf{p}_k^i - \mathbf{p}_k^j\|) \frac{\mathbf{p}_k^i - \mathbf{p}_k^j}{\|\mathbf{p}_k^i - \mathbf{p}_k^j\|}, & \text{if } |\mathcal{C}_k^i| > 0, \\ 0, & \text{otherwise,} \end{cases}$$

where $\mathcal{C}_k^i$ is the set of neighbors visible to boid i in the cohesion zone.

The three forces depend on the function g ; this is a function of the distance between agents. According to the literature [36] this is set to

$$g(d) = \begin{cases} u_{\text{max}} \left(1 - \frac{d}{r_{\text{sep}}} \right) & \text{if } d \in [0, r_{\text{sep}}] \\ u_{\text{max}} & \text{if } d \in [r_{\text{sep}}, r_{\text{ali}}] \\ u_{\text{max}} \left(\frac{r_{\text{coh}} - d}{r_{\text{coh}} - r_{\text{ali}}} \right) & \text{if } d \in [r_{\text{ali}}, r_{\text{coh}}]. \end{cases} \quad (20)$$

The function g modulates the magnitude of the social forces. The function is decreasing inside the separation zone. Then, it becomes constant in the alignment zone and finally decreases again in the cohesion zone. Note that, according to the above social forces, when there is no neighbor in the alignment zone, the boid is encouraged to maintain its current heading [33].

In Fig. 3C-D, the polarization order parameter from the main text [11, 58] is shown at each simulation step $t = k \, dt$ and is computed as:

$$p_k = \left\| \frac{1}{N} \sum_{i=1}^N \frac{\mathbf{v}_k^i}{\|\mathbf{v}_k^i\|} \right\|.$$

Fig. 3C also shows how the weights computed with GateMod change over time for one boid. Fig. SI-1 shows the evolution of the weights for all the boids in Fig. 3C. The figure confirms the phenomenon reported in the main text: weights tend to become uniform as polarization arises. As reported in the main text, the generative model in Fig. 3C is inspired by [27]. Specifically, $q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ is a Gaussian centered in

$$\bar{\mathbf{x}}_{\mathbf{q},k}^i = \begin{bmatrix} \frac{1}{|\mathcal{C}_k^i|} \sum_{j \in \mathcal{C}_k^i} \mathbf{p}_k^j \\ \frac{1}{|\mathcal{A}_k^i|} \sum_{j \in \mathcal{A}_k^i} \mathbf{v}_k^j \end{bmatrix}. \quad (21)$$

The covariance $\Sigma_{\mathbf{q}}$ of the Gaussian is in Tab. 1. In Fig. 3C, the generative model captures the tendency of a boid to stay close to its neighbors and align with their velocity in a free environment [48, 46, 34, 56]. In the main text we also reported that, when the generative model includes a collision avoidance cost, then GateMod yields a milling collective behavior consistently with the literature [27]; see Fig. SI-2. The experiments in Fig. 3D from the main text show that GateMod enables the group to achieve goal-directed behavior with loss of cohesion when a small portion of boids (10%) is goal-informed and seeks to reach the goal destination. In these experiments, the generative model for follower, non-informed, boids is left unchanged while informed bodes are equipped with a model encoding a goal-directed behavior. This is set as a Gaussian centered in

$$\bar{\mathbf{x}}_{\mathbf{q},k}^i = \begin{bmatrix} \mathbf{p}_g \\ \frac{\mathbf{p}_g - \mathbf{p}_k^i}{\|\mathbf{p}_g - \mathbf{p}_k^i\|} \end{bmatrix}$$

and covariance $\Sigma_{\mathbf{q}}$ given in Tab. 1. In the model, the velocity points towards the goal position, $\mathbf{p}_g$, set to $\mathbf{p}_g = [-15, -15]^\top$ in the experiments. Fig. SI-3A and B report experiments result in the same setting from the main text but with 100% and 20% of goal-directed boids. The figures, further supporting the results from main experiments, confirm the onset of a collective goal-directed behavior for the group. Moreover, in the settings of Fig. 3D, we also evaluate how temperature affects the collective behavior. Fig. SI-3C reveals that low values of ϵ allow informed agents to prioritize goal-directed motion, resulting in successful guidance of the flock to the goal. In contrast, higher values of ϵ promote more uniform (high-entropy) primitive weights, which can jeopardize the goal. Finally, to further evaluate GateMod ability to soft-control a group shown in Fig. 3C, we conduct a supplementary study where informed agents still use GateMod (same settings as the main text) but this time followers use three different remarkable models from [11, 13, 60]. As these models are from the literature without any variation, they are not described here. Results from this ablation study are shown in Fig. SI-4A, B, and C. In these figures, the group consists of 40 boids with 6 being informed. Collectively, these experiments confirm GateMod ability to soft-control the group towards the goal position, while maintaining cohesiveness. Remarkably, across all the models, GateMod primitives exhibit temporal variations that suggest, again, a flexible – adaptive – behavior for the leaders.

5.2 Multi-armed bandits experiments

In the experiments from the main text we use the same data and numerical methods from [21].

Data are collected from participants completing a two-armed bandit task organized into 20 blocks of 10 trials each. On trial k , each participant selects an action $u_k \in \{1, 2\}$ and observes a potentially stochastic reward r_k specific to the chosen arm. Belief over rewards is updated via a Kalman filter. As in [21], for the chosen arm $i = u_k$, the posterior mean $\mu_{i,k}$ and posterior standard deviation $s_{i,k}$ are updated according to

$$\mu_{i,k+1} = \mu_{i,k} + \alpha_k (r_k - \mu_{i,k}), \quad (22)$$

$$s_{i,k+1} = s_{i,k} - \alpha_k s_{i,k}, \quad (23)$$

with learning rate

$$\alpha_k = \frac{s_{i,k}^2}{s_{i,k}^2 + \tau_i^2}, \quad (24)$$

where τ_i^2 is the observation noise variance for arm i , and $s_{i,0}^2$ is the prior variance, representing initial uncertainty about the arm's reward. In Experiment 1, one arm is stochastic ($s_0^2 = 10$, $\tau^2 = 10$) and the other yields a fixed reward of zero. In Experiment 2, both arms are stochastic ($s_0^2 = 100$, $\tau^2 = 10$). Then, the hybrid model in [21] is

$$P(u_k = 1) = \Phi\left(\gamma (s_{1,k} - s_{2,k}) + \beta \frac{\mu_{1,k} - \mu_{2,k}}{\sqrt{s_{1,k}^2 + s_{2,k}^2}}\right), \quad (25)$$

where $\Phi(\cdot)$ denotes the standard normal cumulative distribution function, γ is the weight for directed exploration, and β is the weight for random exploration. The weights are inferred from the data. In our GateMod experiments, the state is the vector of posterior mean and variances updated as in [21] using Eq. [22]. At each trial $k = 1, \dots, T$, the agent selects an action $\mathbf{u}_{i,k}, i \in \{1, 2\}$, corresponding to the choice of arm i . The exploitation, uncertainty-seeking and risk aversion primitives from the main text are, respectively:

$$\begin{aligned} \pi^1(\mathbf{u}_{i,k} \mid \mathbf{x}_{k-1}) &= \begin{cases} 0.99 & \text{if } i = \arg \max_{j \in \{1, \dots, N\}} \mu_{j,k} \\ 0.01 & \text{otherwise} \end{cases}, \\ \pi^2(\mathbf{u}_{i,k} \mid \mathbf{x}_{k-1}) &= \frac{s_{i,k}}{\sum_{j=1}^N s_{j,k}}, \\ \pi^3(\mathbf{u}_{i,k} \mid \mathbf{x}_{k-1}) &= \frac{\exp(-s_{i,k})}{\sum_{j=1}^N \exp(-s_{j,k})}. \end{aligned} \quad (26)$$

Since we do not have available the transition kernel $p(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k)$ and we do not want to learn additional parameters over [21], we set this model equal to $q(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k)$. We note that, even if $q(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k)$ also not available, if the two models are the same then they disappear from GateMod formulation – the KL divergence between $p(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k)$ and $q(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k)$ in Eq. [10] is 0, and we can omit this term. In the results from the main text the temperature is $\varepsilon = 0.01$.

In Fig. SI-5 A-C we present supplementary evaluation of GateMod in the same settings considered in the main text but using an alternative fitting via logistic (softmax) regression. Fig. SI-5 D and E show the Predictive Log-Likelihood (PLL) from a 3-fold cross validation among models, both in the case of probit (Fig. SI-5 D) and of logit (Fig. SI-5 E) regression. These findings confirm the results from the

main text. Finally, Fig. SI-6 shows the PXP comparison $\varepsilon \in [0.001, 1]$, using both regression approaches. The figure indicates that low temperatures can be beneficial, whereas excessively large values of ε degrade performance. This deterioration is however expected and it arises because for high temperatures the entropic term in GateFrame dominates, driving the weights towards a uniform distribution.

6 Proving the Statements

The following result on the product of symmetric matrices is used for our convergence analysis of the softmax gradient flow.

Lemma 7 (Product of symmetric matrices). *Let $A, B \in \mathbb{R}^{n \times n}$, $A = A^\top$ and $B = B^\top \succeq 0$. If*

1. $A \succeq 0$, then $\lambda_{\max} \left(\frac{AB + BA}{2} \right) \geq 0$;
2. $A \preceq 0$, then $\lambda_{\max} \left(\frac{AB + BA}{2} \right) \leq 0$.

Proof. First note that $\frac{AB+BA}{2}$ is symmetric and real, so its eigenvalues are real. By the Rayleigh quotient characterization, we have

$$\begin{aligned} \lambda_{\max} \left(\frac{AB + BA}{2} \right) &= \sup_{\|\mathbf{x}\|=1} \mathbf{x}^\top \left(\frac{AB + BA}{2} \right) \mathbf{x} \\ &= \sup_{\|\mathbf{x}\|=1} \frac{1}{2} \left(\mathbf{x}^\top AB\mathbf{x} + \mathbf{x}^\top BA\mathbf{x} \right) = \sup_{\|\mathbf{x}\|=1} \frac{1}{2} \left(\mathbf{x}^\top AB\mathbf{x} + \mathbf{x}^\top A^\top B^\top \mathbf{x} \right) = \sup_{\|\mathbf{x}\|=1} \mathbf{x}^\top AB\mathbf{x} \in \mathbb{R}, \end{aligned} \quad (27)$$

where in the last equality we used the symmetry of A and B . To handle potential singularities of B , define the perturbed matrix $B_t := B + tI_n$, for $t > 0$, which is strictly positive definite. Then $B_t^{1/2}AB_t^{1/2}$ is similar to AB_t , so the spectrum of AB_t is real. Additionally, $B_t^{1/2}AB_t^{1/2}$ is congruent to A and Sylvester's law of inertia implies AB_t has the same number of negative, zero, and positive eigenvalues as A . Moreover, for each $t > 0$, we have:

$$\lambda_{\max} \left(\frac{AB_t + B_tA}{2} \right) = \sup_{\|\mathbf{x}\|=1} \mathbf{x}^\top AB_t\mathbf{x} = \sup_{\|\mathbf{x}\|=1} \left(\mathbf{x}^\top AB\mathbf{x} + t\mathbf{x}^\top A\mathbf{x} \right).$$

Since the Rayleigh quotient is continuous in t and eigenvalues depend continuously on matrix entries, we have

$$\lim_{t \rightarrow 0^+} \lambda_{\max} \left(\frac{AB_t + B_tA}{2} \right) = \sup_{\|\mathbf{x}\|=1} \mathbf{x}^\top AB\mathbf{x} = \lambda_{\max} \left(\frac{AB + BA}{2} \right).$$

Now, assume $A \succeq 0$. Then statement 1 follows being the left-hand side a limit of nonnegative quantities. Similar reasonings apply for statement 2. This concludes the proof. $\square$

Proof of Lemma 4

For any $\mathbf{x} \in \mathbb{R}^n$, we have

$$\begin{aligned} \text{prox}_{\mathbf{H}_{\text{barrier}} - \frac{\|\cdot\|^2}{2}}(\mathbf{x}) &:= \arg \min_{\mathbf{z} \in \mathbb{R}^n} \left(\mathbf{H}_{\text{barrier}}(\mathbf{z}) - \frac{\|\mathbf{z}\|^2}{2} + \frac{1}{2} \|\mathbf{x} - \mathbf{z}\|_2^2 \right) \\ &= \arg \min_{\mathbf{z} \in \Delta_n} \left(\mathbf{H}_{\text{barrier}}(\mathbf{z}) - \frac{\|\mathbf{z}\|^2}{2} + \frac{1}{2} \|\mathbf{x}\|^2 - \mathbf{x}^\top \mathbf{z} + \frac{1}{2} \|\mathbf{z}\|^2 \right) \end{aligned} \quad (28)$$

$$= \arg \min_{\mathbf{z} \in \Delta_n} \left(\sum_{i=1}^n z_i \ln z_i - \mathbf{x}^\top \mathbf{z} \right), \quad (29)$$

where in the equality [28] we used the fact that, by definition of entropic barrier function, the minimum is obtained in Δ_n . Now, to solve the constrained problem [29], consider Lagrangian function

$$\mathcal{L}(\mathbf{z}, \lambda, \mu) = -\mathbf{x}^\top \mathbf{z} + \sum_{i=1}^n z_i \ln z_i + \lambda \left(\sum_{i=1}^n z_i - 1 \right) - \sum_{i=1}^n \mu_i z_i,$$

where $\lambda \in \mathbb{R}$ is the multiplier associated with the equality constraint $\sum_{i=1}^n z_i = 1$, and $\mu_i \geq 0, i \in \{1, \dots, n\}$, are the multipliers associated with the inequality constraints $-z_i \leq 0$. For all $i \in \{1, \dots, n\}$, the optimality KKT conditions are: (i) *stationarity*: $\frac{\partial \mathcal{L}}{\partial z_i} = -x_i + (1 + \ln z_i) + \lambda - \mu_i = 0$; (ii) *primal feasibility*: $\sum_{i=1}^n z_i = 1$; (iii) *dual feasibility*: $\mu_i \geq 0$; and (iv) *complementary slackness*: $\mu_i z_i = 0$.

Now, since the term $(1 + \ln z_i)$ tends to $-\infty$ as any $z_i \rightarrow 0^+$, for the stationarity conditions to hold, the minimizer must lie strictly in the interior of the simplex. By complementary slackness, this implies that $\mu_i = 0$. Thus, the stationarity condition simplifies to

$$-x_i + (1 + \ln z_i) + \lambda = 0 \quad \Longleftrightarrow \quad z_i = \exp(x_i - \lambda - 1).$$

Then, by primal feasibility we have

$$\sum_{i=1}^n z_i = 1 \quad \Longleftrightarrow \quad \sum_{i=1}^n \exp(x_i - \lambda - 1) = 1 \quad \Longleftrightarrow \quad \exp(-\lambda - 1) = \frac{1}{\sum_{j=1}^n \exp(x_j)}.$$

Finally, the unique optimal solution is:

$$z_i^* = \frac{\exp(x_i)}{\sum_{j=1}^n \exp(x_j)} =: (\text{softmax}(\mathbf{x}))_i, \quad i \in \{1, \dots, n\}.$$

This concludes the proof. $\square$

Before giving the proof for Lemma 5 – where we establish the connections between problem [4] and the softmax gradient flow [6] – we note here that, rather interestingly, equation [29] in the proof of Lemma 4 reveals that the smooth best response map is the proximal operator of the function $-\mathbf{H} - \frac{\|\cdot\|^2}{2}$. This immediately follows noticing that Eq. [29] can be equivalently written as $\arg \max_{\mathbf{z} \in \Delta_n} \left(\mathbf{x}^\top \mathbf{z} - \sum_{i=1}^n z_i \ln z_i \right) = \arg \max_{\mathbf{z} \in \Delta_n} \left(\mathbf{x}^\top \mathbf{z} + \mathbf{H}(\mathbf{z}) \right)$. this observation – which to the best of our knowledge is not reported elsewhere – suggests that smooth best-response dynamics and softmax flows both arise naturally from composite optimization problems. We can now give the proof of Lemma 5.

Proof of Lemma 5

First, note that problem [4] is equivalent to the following composite optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) + g(\mathbf{x}) \quad (30)$$

where $f(\mathbf{x}) = F(\mathbf{x}) + \varepsilon \frac{\|\mathbf{x}\|^2}{2}$ is continuously differentiable and $g(\mathbf{x}) = \varepsilon \left(H_{\text{barrier}}(\mathbf{x}) - \frac{\|\mathbf{x}\|^2}{2} \right)$ is closed, proper and convex. The fact that $g(\mathbf{x})$ is CCP follows from its effective domain being Δ_n , and on Δ_n we have $g(\mathbf{x}) = \varepsilon \left(\sum_{i=1}^n x_i \ln x_i - \frac{1}{2} \|\mathbf{x}\|^2 \right)$, whose Hessian is a diagonal matrix with entries $(\nabla^2 g(\mathbf{x}))_i = \varepsilon \left(\frac{1}{x_i} - 1 \right)$ and thus is positive semidefinite for $\mathbf{x} \in \Delta_n$.

Next, let $\mathbf{x}^*$ be the minimizer of [30]. Then, by first-order necessary and sufficient optimality conditions for convex optimization problems [49] we have $0_n \in \nabla f(\mathbf{x}^*) + \partial g(\mathbf{x}^*)$. Multiplying by ε^{-1} and adding and subtracting $\mathbf{x}^*$ to the right-hand side of the above inclusion yields

$$\begin{aligned} 0_n \in [I_n + \varepsilon^{-1} \partial g](\mathbf{x}^*) + \varepsilon^{-1} \nabla f(\mathbf{x}^*) - \mathbf{x}^* &\iff (I_n + \varepsilon^{-1} \partial g)(\mathbf{x}^*) \in \mathbf{x}^* - \varepsilon^{-1} \nabla f(\mathbf{x}^*) \\ &\iff \mathbf{x}^* \in (I_n + \varepsilon^{-1} \partial g)^{-1} (\mathbf{x}^* - \varepsilon^{-1} (\nabla F(\mathbf{x}^*) + \varepsilon \mathbf{x}^*)). \end{aligned}$$

Recalling that $\text{prox}_{\varepsilon^{-1}g} = (I_n + \varepsilon^{-1} \partial g)^{-1}$ and, being by assumption g convex, closed, and proper, then $\text{prox}_{\varepsilon^{-1}g}$ is single-valued [43] and, by Lemma 4, we have $\text{prox}_{\varepsilon^{-1}g} = \text{softmax}$. Therefore,

$$\mathbf{x}^* = \text{softmax}(-\varepsilon^{-1} \nabla F(\mathbf{x}^*)),$$

that is, $\mathbf{x}^*$ is an equilibrium point of the gradient flow dynamics [6]. Since all results are equivalence, the proof of items 1 and 2 is complete. Finally, item 3 follows directly from [22, Theorem 2]. $\square$

Proof of Lemma 6.

First note that $\Delta_n = \mathbb{R}_{\geq 0}^n \cap \{\mathbf{x} \in \mathbb{R}^n \mid \sum_{i=1}^n x_i = 1\}$. Then, by Nagumo's Theorem [41] (see also [5, Exercise 3.12]), to show that Δ_n is forward invariant for the dynamics [6] we need to show that (i) the positive orthant is forward invariant, (ii) the dynamics preserve the total mass constraint, that is $\sum_{i=1}^n \dot{x}_i = 0$ for all $\mathbf{x} \in \mathbb{R}^n$ such that $\sum_{i=1}^n x_i = 1$.

By Nagumo's Theorem [41], the positive orthant is forward invariant for a vector field f if and only if $f_i(\mathbf{x}) \geq 0$, for all $\mathbf{x} \in \mathbb{R}_{\geq 0}^n$ such that $x_i = 0$. Let us consider the softmax gradient flow [6] written in components

$$\tau \dot{x}_i = -x_i + \text{softmax}(-\varepsilon^{-1} \nabla F(\mathbf{x}))_i = -x_i + \frac{\exp(-(\varepsilon^{-1} \nabla F(\mathbf{x}))_i)}{\sum_{j=1}^n \exp(-(\varepsilon^{-1} \nabla F(\mathbf{x}))_j)} := F_{s,i}(\mathbf{x}), \quad i \in \{1, \dots, n\}.$$

Then, for all $\mathbf{x} \in \mathbb{R}_{\geq 0}^n$ such that $x_i = 0$ we have $F_{s,i}(\mathbf{x}) = \frac{\exp(-(\varepsilon^{-1} \nabla F(\mathbf{x}))_i)}{\sum_{j=1}^n \exp(-(\varepsilon^{-1} \nabla F(\mathbf{x}))_j)} > 0$, for each i . Next, let $\mathbf{x} \in \mathbb{R}^n$ such that $\sum_{i=1}^n x_i = 1$. We have

$$\sum_{i=1}^n \tau \dot{x}_i = -\sum_{i=1}^n x_i + \sum_{i=1}^n \text{softmax}(-\varepsilon^{-1} \nabla F(\mathbf{x}))_i = -1 + 1 = 0.$$

This concludes the proof. $\square$

Proof of Theorem 1

For simplicity of notation, let $\mathbf{y} := -\varepsilon^{-1}\nabla\mathbf{F}(\mathbf{x})$ define $G(\mathbf{y}) := D\text{softmax}(\mathbf{y})$. Recall that, being $G(\mathbf{y})$ the Jacobian of a proximal operator, it is symmetric and satisfies $0 \preceq G(\mathbf{y}) \preceq I_n$, for all $\mathbf{y} \in \mathbb{R}^n$ [2, Proposition 12.28]. The Jacobian of $\mathbf{F}_{\text{sm}}$ is therefore $D\mathbf{F}_{\text{sm}}(\mathbf{x}) = \frac{1}{\tau}(-I_n - \varepsilon^{-1}G(\mathbf{y})\nabla^2\mathbf{F}(\mathbf{x}))$. To prove our statement, we have to show that $\mu(\mathbf{F}_{\text{sm}}(\mathbf{x})) \leq -c$ for all $\mathbf{x} \in \mathbb{R}^n$. By applying the log-norm translation (for all $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}$, for any log-norm μ it holds $\mu(A + bI_n) = \mu(A) + b$) and positive homogeneity (for all $A \in \mathbb{R}^{n \times n}$ and $a \in \mathbb{R}$, for any log-norm μ it holds $\mu(aA) = a\mu(\text{sign}(a)A)$) properties, we have

$$\sup_{\mathbf{x}} \mu_2(D\mathbf{F}_{\text{sm}}(\mathbf{x})) \leq -\frac{1}{\tau} + \frac{1}{\tau\varepsilon} \max_{\substack{0 \preceq G \preceq I_n \\ B \succeq 0}} \mu_2((-G)B) =: -\frac{1}{\tau} + (\tau\varepsilon)^{-1} \max_{\substack{0 \preceq G \preceq I_n \\ B \succeq 0}} \lambda_{\max}\left(\frac{(-G)B + B(-G)}{2}\right) \leq -\frac{1}{\tau},$$

where in the first inequality we have used the fact that $\nabla^2\mathbf{F}(\mathbf{x}) \succeq 0$, being $\mathbf{F}$ convex, while the last inequality follows by applying Lemma 7. This concludes the proof. $\square$

Proof of Corollary 1

The statements follow directly from the infinitesimal contractivity of the flow. Specifically, exponential convergence (statement 1) is established in [35], while entrainment to periodic input (statement 2) in [50].

For a more detailed discussion, we refer to [5, Chapter 3]. Finally, to prove item 3, we first compute the Lipschitz constant L_{sm} of the vector field $\mathbf{F}_{\text{sm}}$. We have:

$$\begin{aligned} \|\nabla\mathbf{F}_{\text{sm}}(\mathbf{x})\| &= \tau^{-1} \left\| \nabla \left(-\mathbf{x} + \text{softmax}(-\varepsilon^{-1}\nabla\mathbf{F}(\mathbf{x})) \right) \right\| = \tau^{-1} \left\| -I_n - \varepsilon^{-1}D\text{softmax}(-\varepsilon^{-1}\nabla\mathbf{F}(\mathbf{x}))\nabla^2\mathbf{F}(\mathbf{x}) \right\| \\ &\leq \tau^{-1} + (\tau\varepsilon)^{-1} \left\| D\text{softmax}(-\varepsilon^{-1}\nabla\mathbf{F}(\mathbf{x})) \right\| \left\| \nabla^2\mathbf{F}(\mathbf{x}) \right\| \leq \tau^{-1} (1 + \varepsilon^{-1}L_{\mathbf{F}}) := L_{\text{sm}}. \end{aligned}$$

The conclusion then follows from [6, Theorem 4]. $\square$

6.1 GateMod Proofs

Proof of Lemma 1

We begin with item 1. By Lemma 2 and Remark 1, GateFrame objective becomes:

$$D_{\text{KL}}(p(\mathbf{u}_k | \mathbf{x}_{k-1}) || q(\mathbf{u}_k | \mathbf{x}_{k-1})) + \mathbb{E}_{p(\mathbf{u}_k | \mathbf{x}_{k-1})} [D_{\text{KL}}(p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k))] - \varepsilon\mathbf{H}(\mathbf{w}_k). \quad (31)$$

Substituting the constraint $p(\mathbf{u}_k | \mathbf{x}_{k-1}) = \sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})$ into the function [31] yields, for the first term:

$$\begin{aligned} D_{\text{KL}}(p(\mathbf{u}_k | \mathbf{x}_{k-1}) || q(\mathbf{u}_k | \mathbf{x}_{k-1})) &:= \int \sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1}) \left(\ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) - \ln q(\mathbf{u}_k | \mathbf{x}_{k-1}) \right) d\mathbf{u}_k \\ &= \sum_{\alpha=1}^{n_\pi} w_k^\alpha \mathbb{E}_{\pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})} \left[\ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) - \ln q(\mathbf{u}_k | \mathbf{x}_{k-1}) \right]. \end{aligned}$$

For the second term we have

$$\sum_{\alpha=1}^{n_\pi} w_k^\alpha \mathbb{E}_{\pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})} [D_{\text{KL}}(p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k))].$$

Embedding the GateFrame simplex constraint in the cost yields the desired reformulation.

To prove item 2, first note that Assumption 2 implies that the function $F(\mathbf{w}_k)$ is smooth in $\mathbf{w}_k$. Indeed, the only potential source of discontinuity is the term $\ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1})$. However, since this sum is always non-negative for $\mathbf{w}_k \in \Delta_{n_\pi}$, the function $F(\mathbf{w}_k)$ remains well-defined. Next, we compute the gradient of $F(\mathbf{w}_k)$ w.r.t. a generic weight, say it w_k^γ . We have

$$\begin{aligned} \frac{\partial}{\partial w_k^\gamma} F(\mathbf{w}_k) &= \frac{\partial}{\partial w_k^\gamma} \left(\sum_{\alpha=1}^{n_\pi} w_k^\alpha \mathbb{E}_{\pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})} \left[D_{\text{KL}}(p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)) - \ln q(\mathbf{u}_k | \mathbf{x}_{k-1}) \right] \right) \\ &\quad + \frac{\partial}{\partial w_k^\gamma} \left(\sum_{\alpha=1}^{n_\pi} w_k^\alpha \mathbb{E}_{\pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})} \left[\ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) \right] \right) \\ &= \mathbb{E}_{\pi^\gamma(\mathbf{u}_k | \mathbf{x}_{k-1})} \left[D_{\text{KL}}(p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)) - \ln q(\mathbf{u}_k | \mathbf{x}_{k-1}) + \ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) \right] + \end{aligned}$$

The last terms of the above equality follows from the following computations:

$$\begin{aligned} \frac{\partial}{\partial w_k^\gamma} \left[\sum_{\alpha=1}^{n_\pi} w_k^\alpha \mathbb{E}_{\pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})} \left[\ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) \right] \right] &= \frac{\partial}{\partial w_k^\gamma} \left[\int \sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1}) \ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) d\mathbf{u}_k \right] \\ &= \int \frac{\partial}{\partial w_k^\gamma} \left[\sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1}) \ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) d\mathbf{u}_k \right] \\ &= \int \pi^\gamma(\mathbf{u}_k | \mathbf{x}_{k-1}) \ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) + \sum_{\alpha=1}^{n_\pi} w_k^\alpha \pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1}) \frac{\pi^\gamma(\mathbf{u}_k | \mathbf{x}_{k-1})}{\sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1})} d\mathbf{u}_k \\ &= \int \pi^\gamma(\mathbf{u}_k | \mathbf{x}_{k-1}) \left(\ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) + 1 \right) d\mathbf{u}_k \\ &= \mathbb{E}_{\pi^\gamma(\mathbf{u}_k | \mathbf{x}_{k-1})} \left[\ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) \right] + 1. \end{aligned}$$

To prove item 3 we compute the Hessian of $F(\mathbf{w})$ and use the characterizations F is (i) convex if and only if $\nabla^2 F(\mathbf{w}) \succeq \mathbb{0}_{n_\pi}$, for all $\mathbf{w} \in \mathbb{R}^{n_\pi}$; and (ii) L_F -smooth if and only if $\nabla^2 F(\mathbf{w}) \preceq L_F I_{n_\pi}$, for all $\mathbf{w} \in \mathbb{R}^{n_\pi}$.

We compute

$$\begin{aligned}
(\nabla^2 F(\mathbf{w}_k))_{\alpha\gamma} &= \frac{\partial}{\partial w_k^\alpha} \left(\mathbb{E}_{\pi^\gamma(\mathbf{u}_k | \mathbf{x}_{k-1})} \left[\ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) \right] \right) \\
&= \int \pi^\gamma(\mathbf{u}_k | \mathbf{x}_{k-1}) \frac{\partial}{\partial w_k^\alpha} \left(\ln \sum_{\beta=1}^{n_\pi} w_k^\beta \pi^\beta(\mathbf{u}_k | \mathbf{x}_{k-1}) \right) d\mathbf{u}_k \\
&= \int \pi^\gamma(\mathbf{u}_k | \mathbf{x}_{k-1}) \frac{\pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})}{\mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1})} d\mathbf{u}_k \\
&= \mathbb{E}_{\pi^\gamma(\mathbf{u}_k | \mathbf{x}_{k-1})} \left[\frac{\pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})}{\mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1})} \right].
\end{aligned}$$

Note that, for any $\mathbf{z} \in \mathbb{R}^{n_\pi}$, $\mathbf{z} \neq \mathbf{0}_{n_\pi}$, it holds

$$\begin{aligned}
\mathbf{z}^\top \nabla^2 F(\mathbf{w}_k) \mathbf{z} &= \sum_{\alpha} \sum_{\gamma} z_\alpha z_\gamma \mathbb{E}_{\pi^\gamma(\mathbf{u}_k | \mathbf{x}_{k-1})} \left[\frac{\pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})}{\mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1})} \right] \\
&= \sum_{\alpha} \sum_{\gamma} z_\alpha z_\gamma \int \frac{\pi^\gamma(\mathbf{u}_k | \mathbf{x}_{k-1}) \pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})}{\mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1})} d\mathbf{u}_k \\
&= \int \sum_{\alpha} \sum_{\gamma} z_\alpha z_\gamma \frac{\pi^\gamma(\mathbf{u}_k | \mathbf{x}_{k-1}) \pi^\alpha(\mathbf{u}_k | \mathbf{x}_{k-1})}{\mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1})} d\mathbf{u}_k \\
&= \int \frac{\mathbf{z}^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1}) \pi(\mathbf{u}_k | \mathbf{x}_{k-1})^\top \mathbf{z}}{\mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1})} d\mathbf{u}_k \\
&= \int \frac{(\mathbf{z}^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1}))^2}{\mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1})} d\mathbf{u}_k \geq 0,
\end{aligned}$$

where the last inequality follows being $\mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1}) > 0$ for all $\mathbf{w}_k \in \Delta_{n_\pi}$. Therefore, the Hessian $\nabla^2 F(\mathbf{w}_k)$ is a semi-positive definite matrix, which in turn ensures that the map $F(\mathbf{w}_k)$ is convex. Next, recall that for any symmetric matrix, the largest eigenvalue is given by the maximum of the Rayleigh quotient, that is

$$\lambda_{\max}(\nabla^2 F(\mathbf{w}_k)) = \max_{\|\mathbf{y}\|_2=1} \mathbf{y}^\top \nabla^2 F(\mathbf{w}_k) \mathbf{y}.$$

Then, to prove our statement, it suffices to show $\mathbf{y}^\top \nabla^2 F(\mathbf{w}_k) \mathbf{y} \leq L_F$, for all $\mathbf{y} \in \mathbb{R}^{n_\pi}$ such that $\|\mathbf{y}\|_2 = 1$. We have

$$\begin{aligned}
\mathbf{y}^\top \nabla^2 F(\mathbf{w}_k) \mathbf{y} &= \int \frac{(\mathbf{y}^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1}))^2}{\mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1})} d\mathbf{u}_k \leq \int \frac{(\mathbf{y}^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1}))^2}{\pi_{\min}} d\mathbf{u}_k \\
&\leq \int \frac{\|\mathbf{y}\|^2 \|\pi(\mathbf{u}_k | \mathbf{x}_{k-1})\|^2}{\pi_{\min}} d\mathbf{u}_k = n_\pi \frac{\pi_{\max}^2}{\pi_{\min}},
\end{aligned}$$

where in the first inequality we used the fact that $\mathbf{w}_k^\top \pi(\mathbf{u}_k | \mathbf{x}_{k-1})$ is a convex combination of $\pi(\mathbf{u}_k | \mathbf{x}_{k-1})$, while in the second inequality we applied Cauchy Schwarz inequality. This concludes the proof. $\square$

7 Supplementary Tables

Table 1: Values for the parameters used in the experiments from the main text. Parameters are assumed in accordance with [28, 36, 11]. The initial conditions for the simulations are randomly selected and the numerical solver `solve_ivp` in Python is used.

Parameter	Description	Value
dt	Time step	0.05
N	Number of boids	40
d_u	Action space	30
$v_{\max}$	Maximum velocity norm	1
$u_{\max}$	Maximum acceleration norm	3
α	Vision angle	320°
r_{sep}	Separation radius	1
r_{ali}	Alignment radius	3
r_{coh}	Cohesion radius	12
Σ_{sep}	Covariance of separation primitive	$0.1\mathbb{I}_2$
Σ_{ali}	Covariance of alignment primitive	$0.1\mathbb{I}_2$
Σ_{coh}	Covariance of cohesion primitive	$0.1\mathbb{I}_2$
Σ_p	Covariance of $p(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k)$	$0.01\mathbb{I}_4$
Σ_q	Covariance of $q(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k)$	$0.01\mathbb{I}_4$
ϵ	Entropy regularization parameter	0.5
τ	GateFlow time-scale	1

8 Supplementary Figures

Supplementary References

References

- [1] B. Abbas and H. Attouch. Dynamical systems and forward-backward algorithms associated with the sum of a convex subdifferential and a monotone cocoercive operator. *Optimization*, 64(10):2223–2252, 2014.
- [2] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2 edition, 2017.
- [3] A. Beck. *First-Order Methods in Optimization*. SIAM, 2017.
- [4] A. Beck and M. Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- [5] F. Bullo. *Contraction Theory for Dynamical Systems*. Kindle Direct Publishing, 1.2 edition, 2024.
- [6] F. Bullo, P. Cisneros-Velarde, A. Davydov, and S. Jafarpour. From contraction theory to fixed point algorithms on Riemannian and non-Euclidean spaces. In *IEEE Conf. on Decision and Control*, December 2021.
- [7] V. Centorrino, A. Gokhale, A. Davydov, G. Russo, and F. Bullo. Positive competitive networks for sparse reconstruction. *Neural Computation*, 36(6):1163–1197, 2024.
- [8] P. L. Combettes and J.-C. Pesquet. Deep neural network structures solving variational inequalities. *Set-Valued and Variational Analysis*, 28(3):491–518, 2020.
- [9] R. Cominetti, E. Melo, and S. Sorin. A payoff-based learning procedure and its application to traffic games. *Games and Economic Behavior*, 70(1):71–83, September 2010.
- [10] P. Coucheney, B. Gaujal, and P. Mertikopoulos. Penalty-regulated dynamics and robust learning procedures in games. *Mathematics of Operations Research*, 40(3):611–633, August 2015.
- [11] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective Memory and Spatial Sorting in Animal Groups. *Journal of Theoretical Biology*, 218(1):1–11, 2002.
- [12] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, USA, 2006.
- [13] F. Cucker and S. Smale. Emergent behavior in flocks. *IEEE Transactions on Automatic Control*, 52(5):852–862, 2007.
- [14] A. Davydov, V. Centorrino, A. Gokhale, G. Russo, and F. Bullo. Time-varying convex optimization: A contraction and equilibrium tracking approach. *IEEE Transactions on Automatic Control*, 70(11):7446–7460, 2025.
- [15] A. Davydov, S. Jafarpour, and F. Bullo. Non-Euclidean contraction theory for robust nonlinear stability. *IEEE Transactions on Automatic Control*, 67(12):6667–6681, 2022.
- [16] A. Davydov, A. V. Proskurnikov, and F. Bullo. Non-Euclidean contraction analysis of continuous-time neural networks. *IEEE Transactions on Automatic Control*, 70(1):235–250, 2025.

- [17] I. M. Elfadel and J. L. Wyatt Jr. The” softmax” nonlinearity: Derivation using statistical mechanics and useful properties as a multiterminal analog circuit element. *Advances in neural information processing systems*, 6, 1993.
- [18] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1–2):95–110, March 1956.
- [19] B. Gao and L. Pavel. On the properties of the softmax function with application in game theory and reinforcement learning. *arXiv preprint arXiv:1704.00805*, 2017.
- [20] E. Garrabé and Giovanni Russo. Probabilistic design of optimal sequential decision-making algorithms in learning and control. *Annual Reviews in Control*, 54:81–102, 2022.
- [21] S. J. Gershman. Deconstructing the human algorithms for exploration. *Cognition*, 173:34–42, 2018.
- [22] A. Gokhale, A. Davydov, and F. Bullo. Proximal gradient dynamics: Monotonicity, exponential convergence, and applications. *IEEE Control Systems Letters*, 8:2853–2858, 2024.
- [23] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.
- [24] P. Guan, M. Raginsky, and R. M. Willett. Online markov decision processes with kullback–leibler control cost. *IEEE Transactions on Automatic Control*, 59(6):1423–1438, June 2014.
- [25] S. Hassan-Moghaddam and M. R. Jovanović. Proximal gradient flow and Douglas-Rachford splitting dynamics: Global exponential stability via integral quadratic constraints. *Automatica*, 123:109311, 2021.
- [26] H. Hazimeh, Z. Zhao, A. Chowdhery, M. Sathiamoorthy, Y. Chen, R. Mazumder, L. Hong, and E. Chi. DSelect-k: Differentiable Selection in the Mixture of Experts with Applications to Multi-Task Learning. In *Advances in Neural Information Processing Systems*, volume 34, pages 29335–29347, 2021.
- [27] C. Heins, B. Millidge, L. Da Costa, R. P. Mann, K. J. Friston, and I. D. Couzin. Collective behavior from surprise minimization. *Proceedings of the National Academy of Sciences*, 121(17):e2320239121, 2024.
- [28] C. K. Hemelrijk and H. Hildenbrandt. Self-organized shape and frontal density of fish schools. *Ethology*, 114(3):245–254, 2008.
- [29] E. Jang, S. Gu, and B. Poole. Categorical Reparameterization with Gumbel-Softmax. In *International Conference on Learning Representations*, 2017.
- [30] L. Kozachkov, K. V. Kastanenka, and D. Krotov. Building transformers from neurons and astrocytes. *Proceedings of the National Academy of Sciences*, 120(34), 2023.
- [31] S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–86, 1951.
- [32] D. S. Leslie and E. J. Collins. Individual q-learning in normal form games. *SIAM Journal on Control and Optimization*, 44(2):495–514, January 2005.

- [33] H. Levine, W.J. Rappel, and I. Cohen. Self-organization in systems of self-propelled particles. *Phys. Rev. E*, 63:017101, Dec 2000.
- [34] H. Ling, G. E. McIvor, J. Westley, K. van der Vaart, R. T. Vaughan, A. Thornton, and N. T. Ouellette. Behavioural plasticity and the transition to order in jackdaw flocks. *Nature Communications*, 10(1):5174, 2019.
- [35] W. Lohmiller and J.-J. E. Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683–696, 1998.
- [36] R. Lukeman, Y. Li, and L. Edelstein-Keshet. Inferring individual rules from collective behavior. *Proceedings of the National Academy of Sciences*, 107(28):12576–12580, June 2010.
- [37] R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, July 1995.
- [38] P. Mertikopoulos and W. H. Sandholm. Learning in games via reinforcement and regularization. *Mathematics of Operations Research*, 41(4):1297–1324, November 2016.
- [39] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1928–1937. PMLR, Jun 2016.
- [40] Kevin P. Murphy. *Probabilistic Machine Learning: Advanced Topics*. MIT Press, 2023.
- [41] M. Nagumo. Über die Lage der Integralkurven gewöhnlicher Differentialgleichungen. *Proceedings of the Physico-Mathematical Society of Japan. 3rd Series*, 24:551–559, 1942.
- [42] R. Olfati-Saber. Flocking for multi-agent dynamic systems: Algorithms and theory. *IEEE Transactions on Automatic Control*, 51(3):401–420, 2006.
- [43] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):127–239, 2014.
- [44] J. Peters, K. Mulling, and Y. Altun. Relative entropy policy search. *Proceedings of the AAAI Conference on Artificial Intelligence*, 24(1):1607–1612, July 2010.
- [45] G. Peyré and M. Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019.
- [46] A. M. Reynolds, G. E. McIvor, A. Thornton, P. Yang, and N. T. Ouellette. Stochastic modelling of bird flocks: accounting for the cohesiveness of collective motion. *Journal of the Royal Society Interface*, 19(189):20210745, 2022.
- [47] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, page 25–34, 1987.
- [48] C. W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, 1987.

- [49] R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- [50] G. Russo, M. Di Bernardo, and E. D. Sontag. Global entrainment of transcriptional systems to periodic inputs. *PLoS Computational Biology*, 6(4):e1000739, 2010.
- [51] G. Russo, M. Di Bernardo, and E. D. Sontag. A contraction approach to the hierarchical analysis and design of networked systems. *IEEE Transactions on Automatic Control*, 58(5):1328–1331, 2013.
- [52] W. H. Sandholm. *Population Games and Evolutionary Dynamics*. MIT Press, 2010.
- [53] A. Shafiei, H. Jesawada, K. Friston, and G. Russo. Distributionally robust free energy principle for decision-making. *Arxiv report*, 2025.
- [54] M. Snow and J. Orchard. Biological softmax: Demonstrated in modern Hopfield networks. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 44, 2022.
- [55] E. D. Sontag. Contractive systems with inputs. In J. C. Willems, S. Hara, Y. Ohta, and H. Fujioka, editors, *Perspectives in Mathematical System Theory, Control, and Signal Processing*, pages 217–228. Springer, 2010.
- [56] D. J. T. Sumpter. The principles of collective animal behaviour. *Philosophical Transactions of The Royal Society B: Biological Sciences*, 361(1465):5–22, 2006.
- [57] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
- [58] K. Tunstrøm, Y. Katz, C. C. Ioannou, C. Huepe, M. J. Lutz, and I. D. Couzin. Collective States, Multistability and Transitional Behavior in Schooling Fish. *PLOS Computational Biology*, 9(2):1–11, 02 2013.
- [59] A. Ullah. Entropy, divergence and distance measures with econometric applications. *Journal of Statistical Planning and Inference*, 49(1):137–162, 1996.
- [60] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6-7):1226–1229, 1995.
- [61] S. Xie, G. Russo, and R. H. Middleton. Scalability in nonlinear network systems affected by delays and disturbances. *IEEE Transactions on Control of Network Systems*, 8(3):1128–1138, 2021.
- [62] A. L. Yuille and D. Geiger. Winner-take-all mechanisms. In *The Handbook of Brain Theory and Neural Networks*, 1995.




Figure SI-1: Weights evolution of all the 40 boids in Fig. 3C. In the main text we show the evolution of GateMod weights for boid 1. The weights evolution for all the other boids confirms that, after an initial transient phase where the boids form a cohesive and aligned group, the weights tend to become uniform.




Figure SI-2: **A** GateMod yields milling when the generative model is equipped with a collision avoidance term [42]. The cost is a penalty for actions that point toward neighbors within the separation radius. Plots are obtained for $N = 60$ boids; all other parameters are the same as the ones used in the main text. **B** evolution of the polarization (top) and milling (middle) order parameters. See [11, 58] for the standard definitions of these parameters. A value close to 1 of the milling parameter suggests that agents rotate coherently. Higher milling implies lower polarization order parameter. The bottom panel shows the small distance between boids across the group. The collision avoidance term promotes a non-zero booi-to-boid distance. **C** Weights evolution during milling. GateMod suggests that, at the onset of milling, the alignment and cohesion primitives have higher weights (and hence are used more) over separation.




Figure SI-3: **(A)** (Left) Group behavior when 100% of the boids are informed. (Middle) The distance from the target decreases and the group exhibits polarization. This indicates a goal-directed yet polarized emerging behavior. (Right) Time evolution of the optimal primitives' weights from GateMod. The weights are shown for one boid representative of all the (informed) boids in the group. The evolution of the weights confirms a flexible, adaptive, behavior. **(B)** (Left) Group behavior when 20% of the boids are informed. (Middle) Confirming the results from the main text, the distance from the target decreases (bottom) and the group exhibits polarization with low milling. This indicates a goal-directed yet polarized emerging behavior. (Right) Time evolution of the optimal primitives' weights from GateMod for representative informed and uninformed boids. The weights of the informed boids are more variable than the ones for uninformed agents. This suggests a more adaptive behavior for leaders. **(C)** (Left) Mean and standard deviation of the group metrics across 30 simulations in the setting of Fig. 3D from main text, with different values of the temperature ϵ . While increasing the temperature does not affect polarization, it may hinder the navigation of boids to the goal destination; note the evolution of the distance for higher temperatures. (Right) Mean and standard deviation of GateMod weights for an informed agent (red) and an uninformed agent (blue). Increasing the temperature leads to more uniform weights and, as a result, the informed agents end up orchestrating their primitives according to a policy that is similar to the one of the uninformed agents.





Figure SI-4: **(A)** (Left) Trajectories of N boids from random initial conditions. The 6 goal-informed boids implement GateMod, guiding the 34 uninformed boids, which evolve according to the three-zone model from [11], toward the goal. (Middle) Temporal evolution of group-level metrics. The distance from the goal decreases with time. (Right) Time evolution of the optimal primitives' weights of an informed agent on the simplex. **(B)** (Left) Trajectories of N boids from random initial conditions. The 6 goal-informed boids implement GateMod, guiding the 34 Cucker-Smale boids [13] toward the goal. (Middle) Temporal evolution of group-level metrics. (Right) Time evolution of the optimal primitives' weights of an informed agent on the simplex. **(C)** (Left) Trajectories of N boids from random initial conditions. The 6 goal-informed boids implement GateMod, guiding the 34 uninformed Vicsek boids [60] toward the goal. (Middle) Temporal evolution of group-level metrics. (Right) Time evolution of the optimal primitives' weights of an informed agent on the simplex.






Figure SI-6: Value of PXP for different values of regularization gain ϵ in the case of probit regression (panel (A) for experiment 1 and (B) for experiment 2) and logistic regression (panel (C) for experiment 1 and (D) for experiment 2).