

A Control View of the Free Energy Principle

Giovanni Russo

University of Salerno & Scuola Superiore Meridionale

Free-Energy Across Engineering, Complex Systems and Neuroscience
Scuola Superiore Meridionale

23, June, 2026

Acknowledgments



Francesco Bullo
UCSB



Veronica Centorrino
SSM → ETH



Karl Friston
UCL



Hozefa Jesawada
UniSA → NYUAD



Paula Monfort
SSM



Mike Richardson
Macquarie



Francesca Rossi
SSM



Arash Shafiei
UniSA → CTU




Cinzia Tomaselli
SSM



“I have seen superhuman performance”

Better behavior = reward/cost optimization



» Optimal policy $\neq$ adaptive policy

“I have seen superhuman performance”

Challenges

- » The environment
- » Task knowledge
- » Not everything is observable
- » Observations change



Waymo robotaxis stop in the streets during San Francisco power outage

B B C

22 December 2025

Share Save Add as preferred on Google

Grace Eliza Goodwin

Waypoint The official
Waymo blog

Autonomously navigating the real world: lessons from the PG&E outage

Optimizing Agent-Environment Interactions

Solve a control problem **redefined** by interactions

Challenges


- » The environment
- » Task knowledge
- » Not everything is observable
- » Observations change

Design Principle

Regulation: interoceptive vs exteroceptive tasks


Hypothesis: min-max VFE makes design principle operational

Example: autonomous navigation



Interoceptive


- » Gather information (**corrupted data**)
- » Build model



Exteroceptive

- » Navigate to destination
- » Avoid obstacles

Adaptive Behaviors from Free Energy Minimization



- » Perception → adjust beliefs
- » Learning → update internal models
- » Actions → plan to match predictions

(state estimation)
(internal model)
(optimal control)


Friston, "The free-energy principle: a rough guide to the brain?" Trends in cognitive sciences, 13(7):293–301, 2009

Parr, Pezzulo, Friston, "Active Inference", 2022

Diaz Monfort, Tomaselli, Richardson, GR, "Towards Decision-Making from Human Feedback Under the Free Energy Principle", 2026 ([preprint](#))

Shafiei, Jesawada, Friston, GR, "The Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, vol. 17, article number 707, 2025

Two Models Define Agent Behavior



$$\prod_{k=1}^N p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) p_k(\mathbf{u}_k | \mathbf{x}_{k-1}) p(\mathbf{x}_0)$$

Internal closed-loop model: $p_{0:N}$

$$\prod_{k=1}^N q_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) q_k(\mathbf{u}_k | \mathbf{x}_{k-1}) q(\mathbf{x}_0)$$

Generative model: $q_{0:N}$

Variational Free Energy

$$\underbrace{D_{\text{KL}}(p_{0:N} || q_{0:N})}_{\text{Statistical Complexity}} + \underbrace{\mathbb{E}_p \left[\sum_{k=1}^N c_k^{(x)}(\mathbf{x}_k) \right]}_{\text{Expected Loss}}$$


$$D_{\text{KL}}(p || q) := \int p(\mathbf{v}) \ln \frac{p(\mathbf{v})}{q(\mathbf{v})} d\mathbf{v}$$

Generative model:

- » **Time-series**: KL divergence → error signal
- » **Goal**: through cost function
- » **Combination**: both elements

Abbeel, Ballard, Friston, Jaynes, Jordan, Kappen, Murphy, Nemirovsky, Sinkhorn, Theodorou, Todorov, Yudin, ...


Example: autonomous navigation



Interoceptive

- » Gather information (**corrupted data**)
- » Build model

Internal model



Exteroceptive

- » Navigate to destination
- » Avoid obstacles

Generative model (and cost)


Outline

Adaptation through Variational Free Energy Min-Max:

- **Case Study: Distributionally Robust Free Energy Principle (DR-FREE)**
- **Embodying DR-FREE**

Conclusions and Future Work

FEP $\approx$ Adaptive + Optimal Control



$$\arg \min_{\{p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)\}} \text{Loss (Data)} \quad \longrightarrow \quad \arg \min_{\{p_k(\mathbf{u}_k | \mathbf{x}_{k-1})\}} D_{\text{KL}}(p_{0:N} || q_{0:N}) + \mathbb{E}_p \left[\sum_{k=1}^N c_k^{(x)}(\mathbf{x}_k) \right]$$


What About Closed-Loop Guarantees?

Exteroceptive:

- » Task → fully known
- » Generative model → defined


Interoceptive:

- » Model → best fit to data
- » Data → **corrupted**




Internal model:
parameter mismatch < 10%

Embedding Explicit Robustness Guarantees



Space of densities




$$\mathcal{B}_\eta \left(\bar{p}_{k|k-1}^{(x)} \right) := \left\{ D_{\text{KL}} \left(p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) \parallel \bar{p}_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) \right) \leq \eta_k(\mathbf{x}_{k-1}, \mathbf{u}_k) \right\}$$

DR-FREE Formulation

» Minimize free energy **across** ambiguity → **distributionally-robust** formulation

$$\left\{ p_k^\star (\mathbf{u}_k \mid \mathbf{x}_{k-1}) \right\} \in \arg \min_{\left\{ p_k(\mathbf{u}_k \mid \mathbf{x}_{k-1}) \right\}} \max_{p_k(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k)} D_{\text{KL}}(p_{0:N} \parallel q_{0:N}) + \mathbb{E} \left[\sum_{k=1}^N c_k^{(x)}(\mathbf{X}_k) \right]$$
$$\text{s.t. } \underbrace{p_k(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{u}_k) \in \mathcal{B}_\eta(\bar{p}_{k|k-1}^{(x)})}_{\text{Ambiguity}}$$


Finding DR-FREE Policy



Optimization Cost is Redefined Online

$$\begin{aligned} \max_{p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)} & D_{\text{KL}} \left(p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) || q_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) \right) + \mathbb{E}_{p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)} \left[\bar{c}_k(\mathbf{X}_k) \right] \\ \text{s.t.} & p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) \in \mathcal{B}_\eta \left(\bar{p}_{k|k-1}^{(x)} \right) \end{aligned}$$

Cost-to-go



$$\min_{p_k(\mathbf{u}_k | \mathbf{x}_{k-1})} D_{\text{KL}} \left(p_k(\mathbf{u}_k | \mathbf{x}_{k-1}) || q_k(\mathbf{u}_k | \mathbf{x}_{k-1}) \right) + \mathbb{E}_{p_k(\mathbf{u}_k | \mathbf{x}_{k-1})} [\text{Cost of Ambiguity}]$$

Computing The Cost Of Ambiguity

$$\begin{aligned} \max_{p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)} & D_{\text{KL}} \left(p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) \parallel q_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) \right) + \mathbb{E}_{p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)} \left[\bar{c}_k(\mathbf{X}_k) \right] \\ \text{s.t.} & p_k(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) \in \mathcal{B}_\eta \left(\bar{p}_{k|k-1}^{(x)} \right) \end{aligned}$$

Meta Theorem: the cost of ambiguity

$$\eta_k(\mathbf{x}_{k-1}, \mathbf{u}_k) + \tilde{c}(\mathbf{x}_{k-1}, \mathbf{u}_k)$$

▶ Ambiguity radius
 ▶ Convex scalar
▶ Global minimum
▶ Non-negative

$$\min_{\alpha \geq 0} \tilde{V}_\alpha(\mathbf{x}_{k-1}, \mathbf{u}_k)$$

The Optimal Policy





$$\min_{p_k(\mathbf{u}_k | \mathbf{x}_{k-1})} D_{KL} (p_k(\mathbf{u}_k | \mathbf{x}_{k-1}) || q_k(\mathbf{u}_k | \mathbf{x}_{k-1})) + \mathbb{E}_{p_k(\mathbf{u}_k | \mathbf{x}_{k-1})} [\eta_k(\mathbf{X}_{k-1}, \mathbf{U}_k) + \tilde{c}(\mathbf{X}_{k-1}, \mathbf{U}_k)]$$

» **Optimal policy**: balancing interoceptive/exteroceptive needs

$$p_k^\star(\mathbf{u}_k | \mathbf{x}_{k-1}) = \frac{q_k(\mathbf{u}_k | \mathbf{x}_{k-1}) \exp(-\eta_k(\mathbf{x}_{k-1}, \mathbf{u}_k) - \tilde{c}(\mathbf{x}_{k-1}, \mathbf{u}_k))}{\int q_k(\mathbf{u}_k | \mathbf{x}_{k-1}) \exp(-\eta_k(\mathbf{x}_{k-1}, \mathbf{u}_k) - \tilde{c}(\mathbf{x}_{k-1}, \mathbf{u}_k)) d\mathbf{u}_k}$$

» **Optimal cost**: smallest possible free energy

$$\hat{c}_k(\mathbf{x}_{k-1}) = -\ln \int q_k(\mathbf{u}_k | \mathbf{x}_{k-1}) \exp(-\eta_k(\mathbf{x}_{k-1}, \mathbf{u}_k) - \tilde{c}(\mathbf{x}_{k-1}, \mathbf{u}_k)) d\mathbf{u}_k$$



**Interoceptive vs
exteroceptive balance**


What About Closed-Loop Guarantees?

Exteroceptive:

- » Task → fully known
- » Generative model → defined


Interoceptive:

- » Model → best fit to data
- » Data → **corrupted**




Internal model:
parameter mismatch < 10%

DR-FREE in Action




A. Shafiei, H. Jesawada, K. Friston, GR, "The Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, vol. 17, article number 707, 2025

DR-FREE in Action





DR-FREE



Internal model:
parameter mismatch $< 10\%$

DR-FREE in Action (higher-dimensional task)



» **Model:** trained from corrupted data

» **Goal:** move forward




» **Comparison:** SoA methods

T. Berrueta, A. Pinosky, T. & Murphey, "Maximum diffusion reinforcement learning", *Nature Machine Intelligence*, vol. 6, 504–514, 2024

G. Williams, et al. "Information theoretic MPC for model-based reinforcement learning", IEEE ICRA, 2017

A. Shafiei, H. Jesawada, K. Friston, GR, "The Distributionally Robust Free Energy Principle for Decision-Making", *Nature Communications*, vol. 17, article number 707, 2025

Narrowing the Sim-to-Real Gap



H. Jesawada, GR, A. Swikir, F. Abu-Dakka, "Learning-Based Robust Control: Unifying Exploration and Distributional Robustness for Reliable Robotics via Free Energy", Pro. 2026 IEEE International Conference on Robotics and Automation, June 1 – 5, 2026 ([preprint](#))

Key References

- ▶ H. Jesawada, GR, A. Swikir, F. Abu-Dakka, "Learning-Based Robust Control: Unifying Exploration and Distributional Robustness for Reliable Robotics via Free Energy", IEEE International Conference on Robotics and Automation, 2026
- ▶ M. P. Diaz Monfort, C. Tomaselli, M. Richardson, GR, "Towards Decision-Making from Human Feedback Under the Free Energy Principle", 2026
- ▶ Shafiei, H. Jesawada, K. Friston, GR, "The Distributionally Robust Free Energy Principle for Decision-Making", Nature Communications, vol. 17, article number 707, 2025
- ▶ E. Garrabé, H. Jesawada, C. Del Vecchio, GR, "On Convex Data-Driven Inverse Optimal Control for Nonlinear Stochastic Non-stationary Systems", Automatica, vol. 173, 112015, 2025
- ▶ F. Rossi, V. Centorrino, F. Bullo, GR, "Neural Policy Composition from Free Energy Minimization", 2025
- ▶ F. Rossi, E. Garrabé, GR, "Free-Gate: Planning, Control And Policy Composition via Free Energy Gating", 6th Intl. Workshop on Active Inference, 2025
- ▶ E. Garrabé, M. Lamberti, GR, "Optimal Decision-Making for Autonomous Agents via Data Composition", IEEE Control Systems Letters, vol. 7, pp. 2557-2562, 2023
- ▶ E. Garrabé, GR, "Probabilistic design of optimal sequential decision-making algorithms in learning and control", Annual Reviews in Control, vol. 54, pp. 81-102, 2022
- ▶ D. Gagliardi, GR, "On a probabilistic approach to synthesize control policies from example datasets", Automatica, vol. 137, 110121, 2022

Thank you!

giovarusso@unisa.it

Open PhD positions at SSM!

Upcoming workshops on FEP in Reykjavik (7/Jul)

***Icons from flaticon.com**