An Adversarial Interpretation of Information-Theoretic Bounded Rationality

GRASP Laboratory

Penn Engineering

Pedro A. Ortega and Daniel D. Lee

Introduction

Typically, exact planning in complex and uncertain environments is intractable.

Understanding bounded rationality is an important challenge in artificial intelligence.

Using the free energy rather than expected utility as the objective function, we can model bounded rational planning.

Free energy has been mainly justified through statistical mechanics and information theory, but its relation to game theory is unclear.

Our contribution: new game-theoretic interpretation of free energy.

Expected Utility

Many problems in AI can be expressed as maximizations of *expected utilities*.

Sequential decisions can always be rephrased as equivalent single-step decisions (even MDPs).

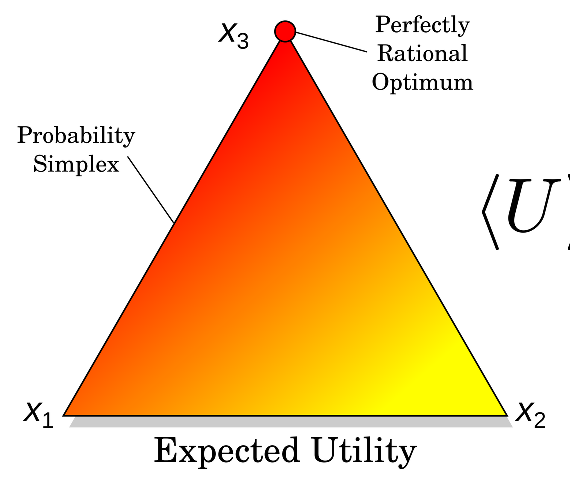

$$\langle U \rangle = \sum_x p(x) U(x)$$

Probabilities (Policy) — Utilities

With full control of the (expected) outcome, there exists always a deterministic optimal policy.

Free Energy

Free energy models planning with resource limitations by adding a regularization term. A bounded rational agent is thought of *as if* it is optimizing the free energy.

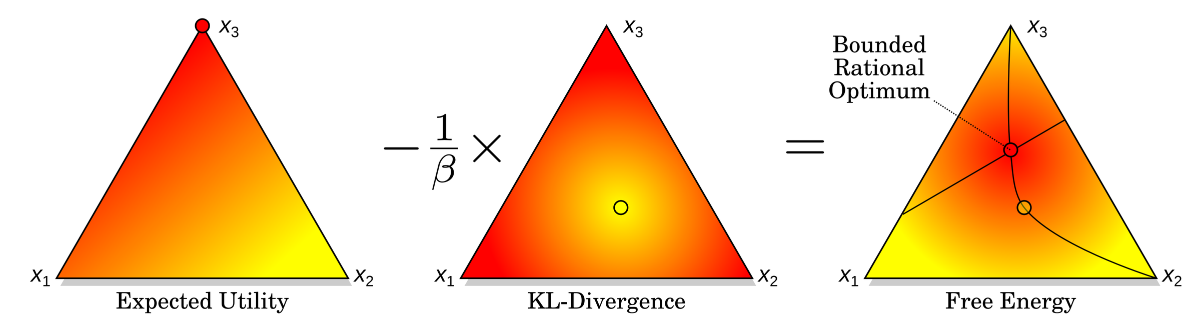

$$F = \sum_x p(x)U(x) - \frac{1}{\beta}\sum_x p(x)\log\frac{p(x)}{q(x)}$$

Expected Utility — Regularization; $p(x)$: Posterior Policy (after planning); $q(x)$: Prior Policy (before planning)

Bounded Rational Solution

The new optimal policy is

$$p^*(x) = \frac{1}{Z} q(x) e^{\beta U(x)}$$

where

$$Z = \sum_x q(x) e^{\beta U(x)}$$

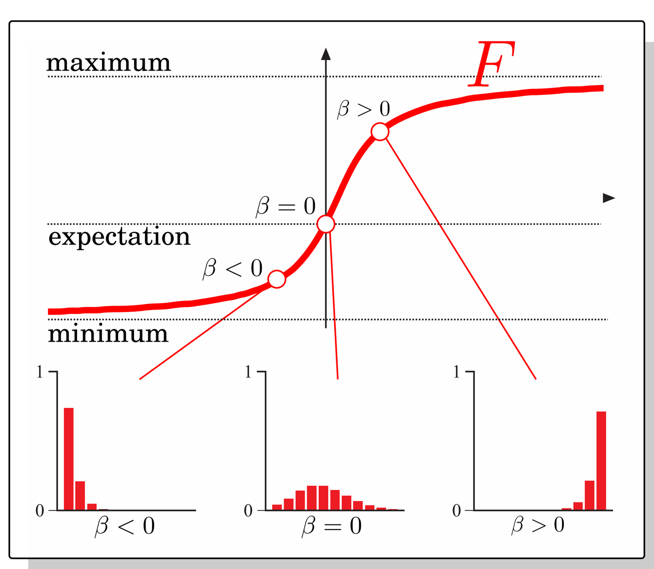

Different values of β parameterize different degrees of control of the outcome.

Main Result: Adversarial Interpretation of Free Energy

Single-agent free energy optimization is equivalent to a game between the agent and an *imaginary* adversary.

Adversary can choose costs C that diminish the agent's payoffs.

Costs are not arbitrary: adversary must pay a penalty that is exponential in the costs.

Limited Resources ↔ Adversary

$$\max_p \sum_x \left\{ p(x)U(x) - \frac{1}{\beta} p(x) \log\frac{p(x)}{q(x)} \right\}$$

$$= \max_p \min_C \sum_x \left\{ p(x)[U(x) - C(x)] - q(x)e^{\beta C(x)} \right\} + \text{const.}$$

Result obtained by applying Legendre-Fenchel transformation only on the regularizer.

Minimax game against an imaginary adversary.

Costs: diminish agent's payoffs

Penalties: The price the adversary pays for the costs.

Indifference

The imaginary adversary's best strategy consists in choosing costs such that the agent's payoffs are *uniform*.

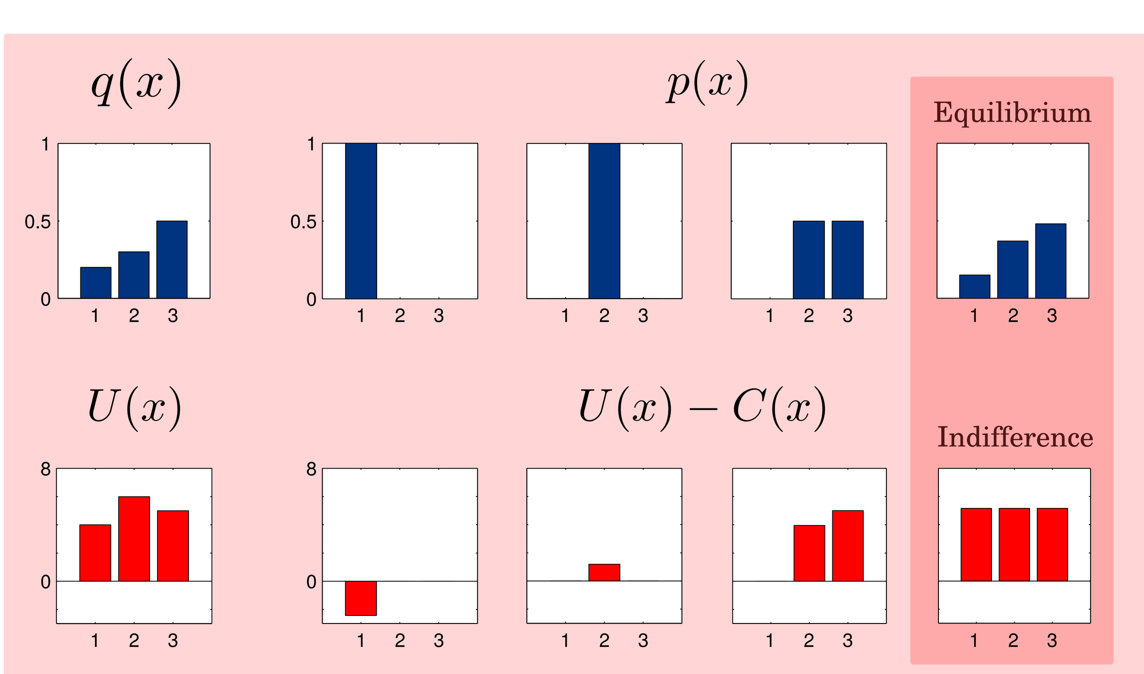

Other Regularizers

Using the same idea, one can unveil the imaginary advesary implicit in other regularizers.

Expected Utility (no regularization):

$$\max_p \min_C \sum_x p(x)[U(x) - C(x)] - \delta(C(x))$$

Power Function:

$$\max_p \min_C \sum_x p(x)[U(x) - C(x)] - |C(x)|^\alpha$$

Portfolio Theory:

$$\max_p \min_C \sum_x p(x)[U(x) - C(x)] - \frac{1}{2}\lambda C^T \Sigma C$$

Take-Home Message

Regularization encodes assumptions about deviations from expected utility.

Legendre-Fenchel transformation reinterprets the free energy as a game against an imaginary adversary.

Each form of regularization is associated to a penalization scheme for adversarial costs.

Stochastic policies guard against unexpected, adversarial costs, which can happen when planning is constrained.

GRASP Laboratory, School of Engineering and Applied Sciences, University of Pennsylvania, Philadelphia, U.S.A.