

PREDICTIVE GAME THEORY

David H. Wolpert

NASA Ames Research Center

<http://ic.arc.nasa.gov/people/dhw/>

NASA-ARC-05-097

ROADMAP

1) *Review probability theory and game theory*



2) *Apply probability theory to games (as opposed to applying it within games)*



3) *E.g., Coupled players and Quantal Response Eq.*



4) *New mathematical tools: rationality functions, cost of computation, varying numbers of players, etc.*

ONLY IDEA IN THIS TALK:

Human beings are physical objects

REVIEW OF PROBABILITY THEORY

- 1) Probability theory is the only “calculus of uncertainty” that obeys Cox’s axioms
- 2) In particular obeying Cox forces *Bayes Theorem*:
$$P(\text{truth } z \mid \text{knowledge } \iota) \propto P(\iota \mid z) P(z)$$
- 3) Given a $P(z \mid \iota)$ and a *loss function* $L(\text{truth } z, \text{prediction } y)$, the *Bayes-optimal* prediction is $\operatorname{argmin}_y E_P[L(., y)]$ (Savage).
- 4) $\operatorname{argmax}_z P(z \mid \iota)$ is an approximation; the *MAP* prediction

Probability theory to reason about physical objects.

Minimize expected loss to distill $P(z)$ to a single z .

EXAMPLE OF PROBABILITY THEORY

1) Let the random variable we wish to predict itself be a probability distribution, $z = q(x)$.

2) Information theory tells us to use the *Entropic prior*

$$P(q) \propto \exp[\alpha S(q)]$$

where $S(q)$ is the Shannon entropy of q , and $\alpha \in \mathfrak{R}^+$

3) Let the knowledge ι about q be $E_q(H) = h$ for some $H(x)$:

$$P(q | \iota) \propto \exp[\alpha S(q)] \delta[E_q(H) - h]$$

EXAMPLE OF PROBABILITY THEORY - 2: *STATISTICAL PHYSICS*

4) So MAP q maximizes $S(q')$ over the q' obeying $E_{q'}(H) = h$:

5) Let x be phase space position of a physical system with $H(x)$ the Hamiltonian. The MAP q gives the **Canonical Ensemble**:

$$q(x) \propto \exp[\beta H(x)]$$

6) If the numbers of particles of various types also varies stochastically, the MAP q is the **Grand Canonical Ensemble**.

REVIEW OF GAME THEORY

- **N independent *players*, each with possible *moves*, $z_i \in Z_i$**
- **Each i has a distribution $q_i(z_i)$; $q(z) = \prod_i q_i(z_i)$**
- **N *utility functions* $u^i(z)$; player i wants maximal $E_q(u^i)$**
- **$E_q(u_i)$ depends on q — but i only sets q_i**

***Equilibrium concept*: mapping from $\{u^i\} \rightarrow q$**

E.g., Nash equilibrium: No $E_q(u^i)$ rises by changing (just) q_i

Hypothesis: Only equilibrium q can arise with humans.

“All we must do is find the right equilibrium concept.”

ONLY IDEA IN THIS TALK:

Human beings are physical objects

GAME THEORY AND LOSS FUNCTIONS

1) Humans are physical objects; to reason about the outcome of a game we *must* use distributions over outcomes:

Game theory hypothesis is wrong

- **N.b., bounded rationality automatic with using distributions.**

2) To distill a distribution over game outcomes to single outcome need a loss function L measuring the quality of the prediction:

“Equilibrium” of a game not meaningful without a loss function.

- **L associated with the external scientist, *not* with the players.**

COUPLED PLAYERS
(similar for uncoupled)

1) Say players are statistically coupled.

E.g., they have previously interacted.

2) Game outcome changes between game instances, but how “rational” the players are does not. How formalize that?

3) Define $U^i(x_i) = E(u_i | x_i)$, and require that for some function ε_i , all game instances obey

$$E_{q_i}(U^i) = \varepsilon^i(U^i)$$

4) Information theory: $\varepsilon^i(U^i) \propto \sum_{x'_i} \exp[\beta_i U^i(x'_i)] U^i(x'_i)$

E.g., $q_i(x_i) \propto \exp[\beta_i U^i(x_i)]$. Many other q_i as well.

QUANTAL RESPONSE EQUILIBRIUM

1) So Bayes theorem says that with the entropic prior over q ,

$$P(q | \iota) \propto \exp[\alpha S(q)] \prod_i \delta[E_{q_i}(U^i_{q_{-i}}) - \varepsilon^i(U^i_{q_{-i}})]$$

- All $\beta_i \rightarrow \infty$; the support of $P(q | \iota)$ is the Nash equilibria.

2) Locally MAP q 's - local maxima of $P(q | \iota)$ - are approximated by a set of coupled equations:

$$q_i(x_i) \propto \exp[\beta_i U^i_{q_{-i}}(x_i)]$$

- **Quantal Response Eq.** (QRE - McKelvey and Palfrey)

QRE and BAYES OPTIMALITY

1) Unimodal $P(q | \iota)$:

- The QRE approximates a q (the MAP),
which in turn approximates the Baye-optimal q .
- How good an approximation depends on loss function.

2) Multimodal $P(q | \iota)$. Say all $\beta_i \rightarrow \infty$ (full rationality):

**If the loss function $L(.,)$ is continuous,
the Bayes optimal prediction is not a Nash equilibrium.**

QUANTIFYING A PLAYER'S RATIONALITY

Want a way to quantify “how rational” an (arbitrary!) q_i is, for an (arbitrary) effective utility U^i .

Natural desiderata. **KL rationality** is one solution to them:

- 1) Use **Kullback-Leibler distance** $KL(p, p')$ to measure “distance” between distributions p and p' .
- 2) **KL rationality** is the β_i minimizing the KL distance from the associated Boltzmann distribution to q_i :

$$\rho_{KL}(U^i, q_i) = \operatorname{argmin}_{\beta_i} KL(q_i, \exp(\beta_i U^i))$$

GAMES WITH VARIABLE NUMBERS OF PLAYERS

- 1) Recall: The MAP q for physical systems where the numbers of particles of various types varies stochastically is the Grand Canonical Ensemble (GCE).**

Intuition: Players with “types” = particles with types

- 2) So MAP q for a game with varying numbers of players is governed by the GCE:**
 - i) Corrections to replicator dynamics,**
 - ii) New ways to analyze firms (varying numbers of employees of various types), etc.**

FUTURE WORK

- 1) *Apply to cooperative game theory - issue of what equilibrium concept to use rendered moot.*
- 2) *Apply to mechanism design - bounded rational mechanism design, corrections to incentive compatibility criterion, etc.*
- 3) *Extend (1, 2) to games with varying numbers of players.*
- 4) *Investigate alternative choices of $P(t | q)$ and $P(q)$, e.g., to reflect Allais' paradox.*
- 5) *Integrate (predictive) game theory with the field of user modeling (i.e., with modeling real people as Bayes nets).*

CONCLUSION

- 1) *Probability theory governs outcome of a game; there is a distribution over mixed strat.'s, not a single “equilibrium”.*
- 2) *To predict a single mixed strategy must use our loss function (external to the game's players).*
- 3) *Provides a quantification of any strategy's rationality.*
- 4) *Prove rationality falls as cost of computation rises (for players who have not previously interacted).*
- 5) *All extends to games with varying numbers of players.*