Asymmetric Abstractions for Adversarial Settings

Nolan Bard, Michael Johanson, and Michael Bowling
{nolan, johanson, bowling}@cs.ualberta.ca

University of Alberta
Computer Poker Research Group

Motivation

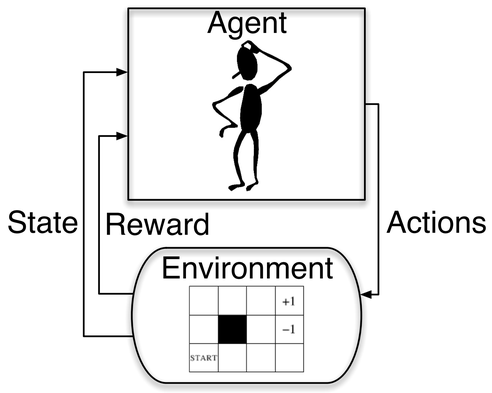

Problem: large state spaces
- **Abstraction** may be needed
- Considerable prior work on feature extraction but...

What granularity of abstraction?
- Trade-off between resources vs. fidelity
- Single agent: use finest tractable abstraction
- Multiagent: investigated in this work

Extensive-Form Games

- Game rules are common knowledge
- Representation of each agent's (unknown) behaviour is needed.
- Abstraction choices represent beliefs about other agents' **capabilities**
- Finer grained abstractions not necessarily better

Counterfactual Regret Minimization (CFR)
[Zinkevich *et al.*, NIPS 2008]
- Iterative self-play algorithm
- State-of-the-art algorithm for approximating Nash in 2-player zero-sum games

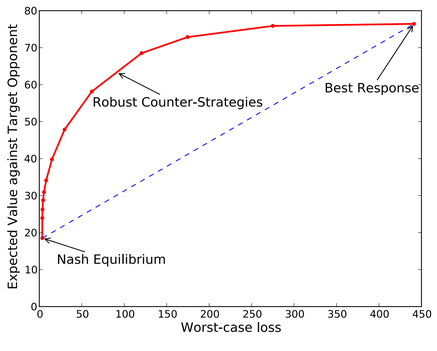

Allocating Abstraction Resources

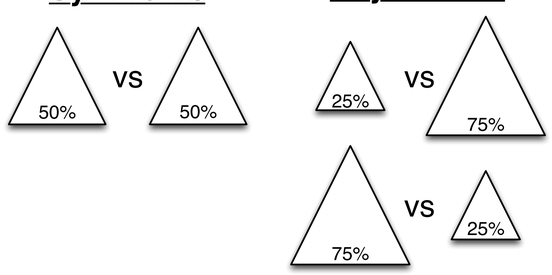

Prior asymmetric abstraction work

Pathologies
[Waugh *et al.*, AAMAS 2009]

CFR-BR
[Johanson *et al.*, AAAI 2012]

Restricted Nash Responses
[Johanson *et al.*, NIPS 2007]
- Good for known strategies

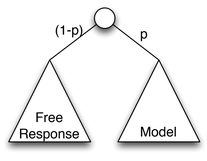

Data Biased Response
[Johanson *et al.*, AISTATS 2009]
- More effective for models based on data
- Model's abstraction impacts data **sparsity**

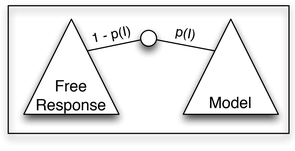

Limit Texas Hold'em

- Player actions: fold, call, raise
- Chance
  - 52 card deck
  - 2 private cards per player
  - 3/1/1 public on flop/turn/river
- Real game > 10^{14} infosets

Feature Extraction
- Player actions unabstracted
- Cards→E[HS^2] percentile buckets

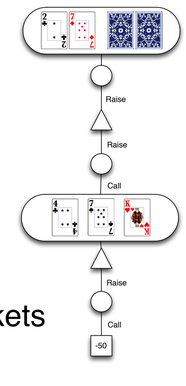

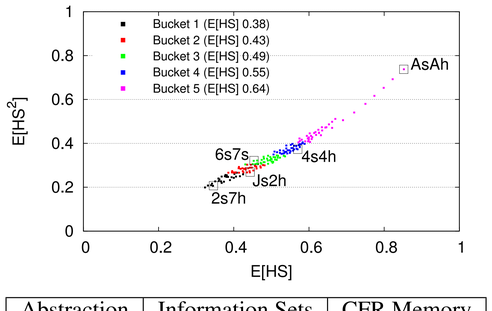

Abstraction	Information Sets	CFR Memory
5-Bucket	3624290	140 MB
8-Bucket	23551424	934 MB
12-Bucket	118671936	4708 MB

Empirical Results

Nash Equilibrium Approximation

	12-5	12-8	12-12	8-5	8-8	8-12	5-5	5-8	5-12	Mean	RGBR	Size
12-5	0	-3	-6	20	18	16	43	41	41	18.970 ± 0.128	435.757	2424 MB
12-8	3	0	-3	23	22	20	36	35	35	18.890 ± 0.143	378.919	2821 MB
12-12	6	3	0	16	16	14	29	28	30	15.842 ± 0.175	289.227	4708 MB
8-5	-20	-23	-16	0	-3	2	22	21	24	0.662 ± 0.121	379.659	537 MB
8-8	-18	-22	-16	3	0	4	16	15	20	0.276 ± 0.144	312.762	934 MB
8-12	-16	-20	-14	-2	-4	0	12	12	16	-1.985 ± 0.099	255.845	2821 MB
5-5	-43	-36	-29	-22	-16	-12	0	3	7	-16.189 ± 0.112	317.1	140 MB
5-8	-41	-35	-28	-21	-15	-12	-3	0	5	-16.751 ± 0.153	283.37	537 MB
5-12	-41	-35	-30	-24	-20	-16	-7	-5	0	-19.714 ± 0.190	234.351	2424 MB
12-FULL	-22	-22	-21	-14	-13	-11	-2	-1	2	-11.526 ± 0.221	87.2765	3450 MB
8-FULL	-36	-36	-32	-26	-24	-21	-14	-12	-7	-23.093 ± 0.070	101.256	1563 MB
5-FULL	-54	-50	-45	-42	-38	-35	-29	-26	-21	-37.585 ± 0.150	122.385	1166 MB

Observations
- 5-5→8-8→12-12: ↑ 1-vs-1, ↓ RGBR. Not guaranteed but often assumed anyways.
- 5-5→8-5→12-5: ↑ 1-vs-1, ↑ RGBR. First abstraction pathology in large game.
- 5-5→5-8→5-12→5-FULL: ↓ mean 1-vs-1, ↓ RGBR. Echoes CFR-BR results.
- 12-5→12-12→12-FULL: Symmetric abstractions optimize neither 1-vs-1 nor RGBR
- Demonstrates trade-off between 1-vs-1 and RGBR performance

Robust Counter-Strategies

Experimental Design
- Versus "tilted" 8-bucket strategy
- Observations gathered with probe
- Varied: abstractions, quantity of data, RNR/DBR confidence
- First evaluation using real game worst-case performance

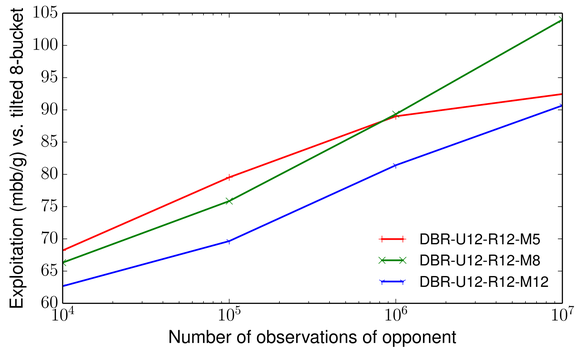

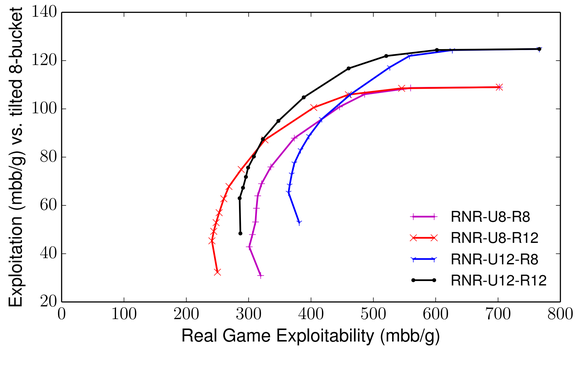

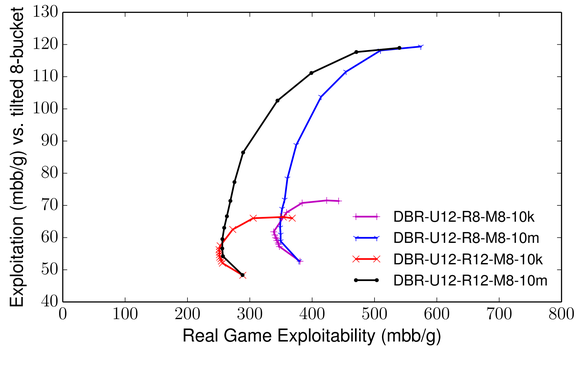

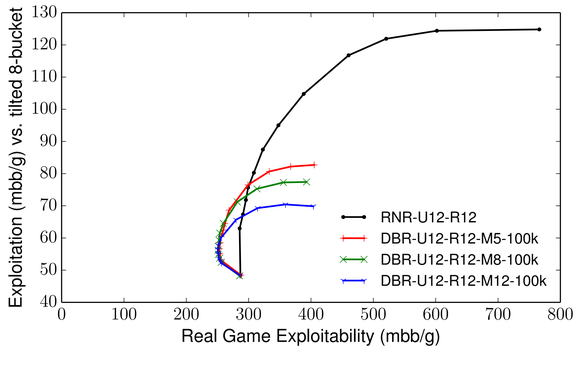

- Tailor model size to quantity of data
- RNR U8s-R8s→U12s-R8s: ↑1-vs-1, ↑RGBR. Similar performance trade-off.
- RNR U8s-R8s→U8s-R12s: ↓RGBR. Strictly dominant RNR performance.
- DBR trade-off depends on data
- Positive model weight: can be less exploitable than abstract Nash.
- Potentially "free" exploitation

Alberta Innovates
NSERC CRSNG
compute • calcul CANADA
Calcul Québec
WestGrid