Evaluating State-Space Abstractions in Extensive-Form Games

Michael Johanson, Neil Burch, Richard Valenzano and Michael Bowling
University of Alberta, Canada

University of Alberta
Computer Poker Research Group

Creating Agents for Large Domains

Large multiagent domains, like Texas Hold'em Poker, are too big to calculate an optimal strategy. Instead, we derive a smaller **abstract domain**, calculate an **abstract strategy**, and use it to choose actions in the real domain.

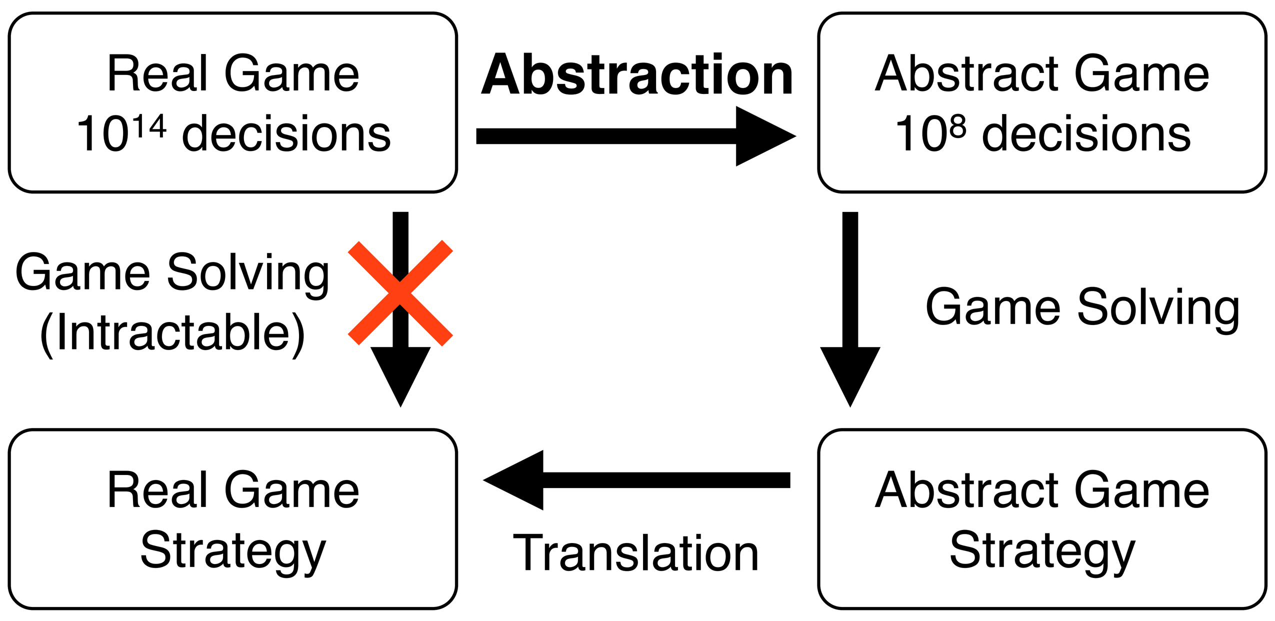

Abstraction Example: Percentile Hand Strength (PHS)

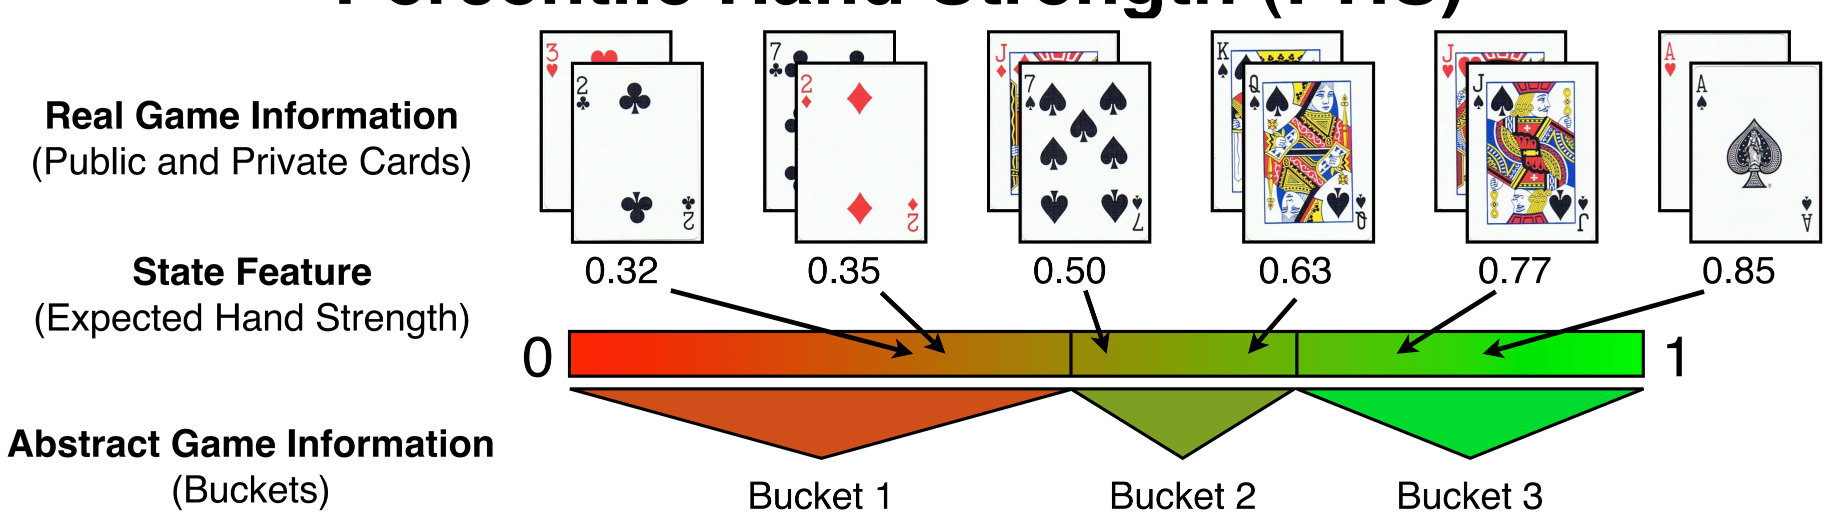

How can we evaluate an abstraction?

♦ **In-game performance against a Nash equilibrium**
Can't compute an optimal policy in large games!

♣ **In-game performance against other strategies**
Many suboptimal strategies tie, intransitivities are possible

♥ **Exploitability (or suboptimality) in the real game**
Abstract equilibria are rarely the abstract strategies with the lowest real-game exploitability
Does not correspond well with in-game performance

Game Solving Algorithms: CFR and CFR-BR

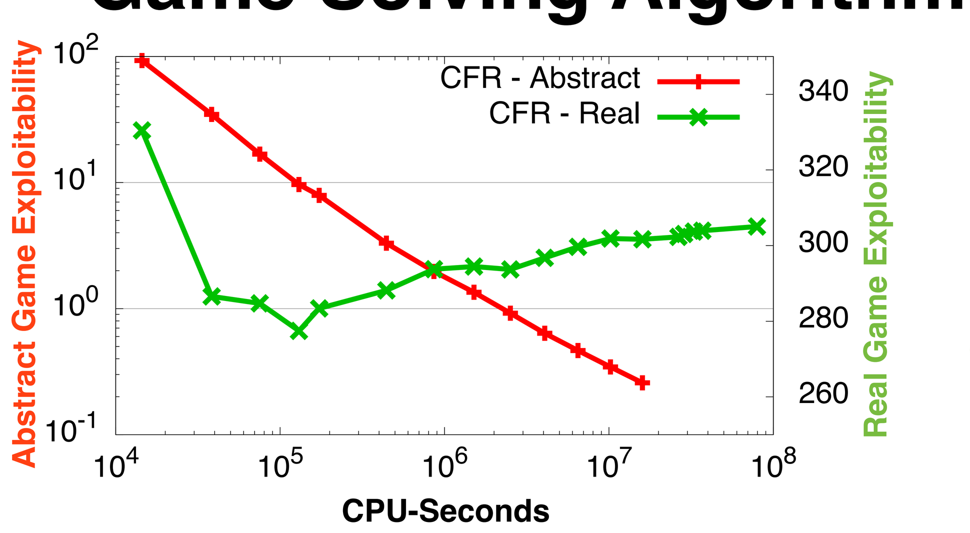

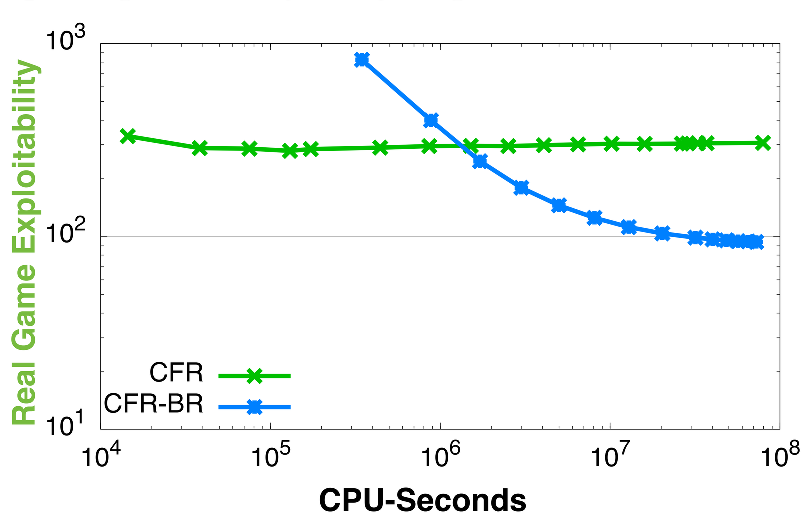

CFR is a state-of-the-art algorithm for approximating Nash equilibria. **CFR-BR** is a new variant that works well with abstraction: it finds the abstract strategy with the lowest real-game exploitability. This gives a new, fourth way to evaluate an abstraction: **by its ability to represent a real game Nash equilibrium**.

Three Objectives of This Paper:

♦ **Demonstrate the use of CFR-BR for evaluating abstractions.**

♣ **Evaluate state-of-the-art poker abstraction techniques.**
Expectation-Based versus Potential Aware.
New state features used by a world-class poker agent.

♥ **Evaluate the use of Imperfect Recall abstractions.**
Popular technique in the Annual Computer Poker Competition.
No proof of convergence, but effective in practice.

Perfect Recall

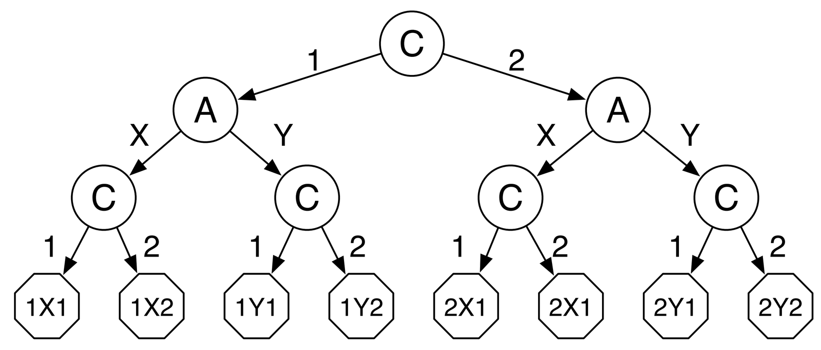

Imperfect Recall

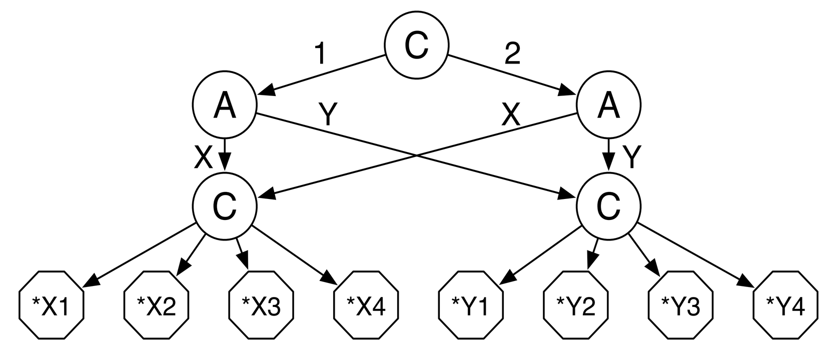

Imperfect Recall discards chosen old state information to focus the abstraction on more recent and important state information.

New Poker Distance Metrics

Earthmover's Distance over Hand Strength Distributions
In the first three rounds, we can compare two poker hands' distribution over future strength.

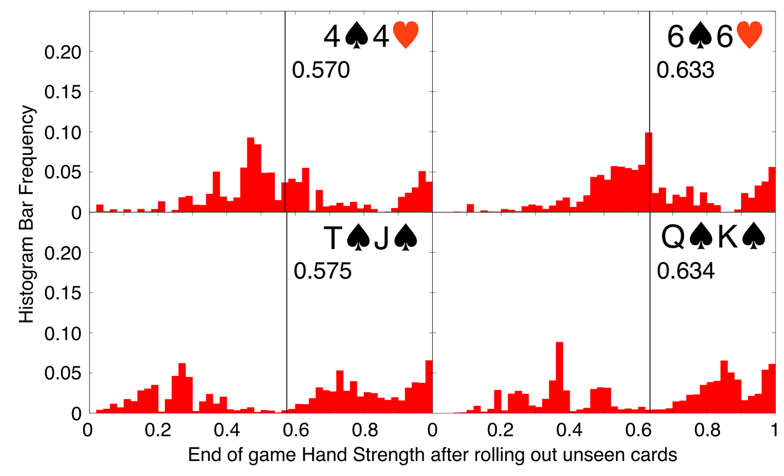

Opponent Cluster Hand Strength (OCHS)
In any round, measure a hand's probability of winning against clusters of opponent hands. Compare vectors with L^2.

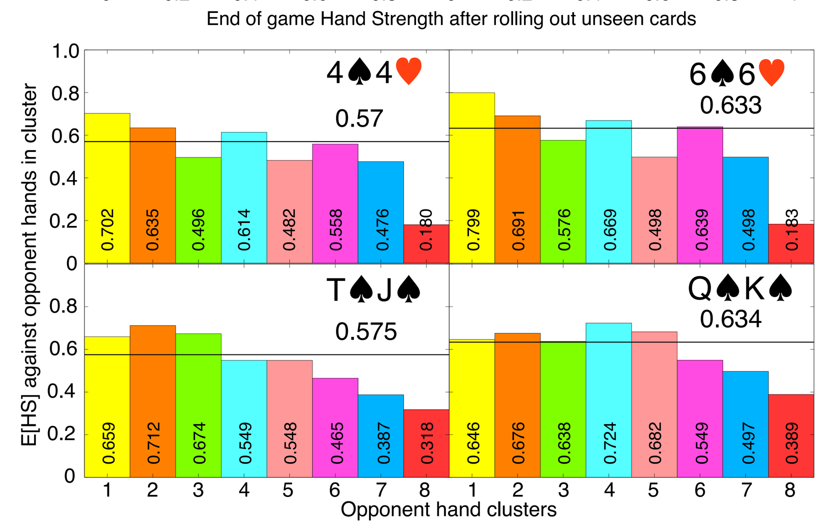

	Hand Strength Difference				Earthmover Distance				OCHS L^2			
	4♠4♥	6♠6♥	T♠J♠	Q♠K♠	4♠4♥	6♠6♥	T♠J♠	Q♠K♠	4♠4♥	6♠6♥	T♠J♠	Q♠K♠
4♠4♥	--	0.063	0.005	0.064	--	3.101	6.212	5.143	--	0.066	0.095	0.104
6♠6♥	0.063	--	0.015	0.001	3.101	--	6.462	5.286	0.066	--	0.117	0.105
T♠J♠	0.005	0.015	--	0.059	6.212	6.462	--	3.103	0.095	0.117	--	0.098
Q♠K♠	0.064	0.001	0.059	--	5.143	5.286	3.103	--	0.104	0.105	0.098	--

Forming Abstractions with *k*-Means Clustering

After choosing a distance metric (Earthmover or OCHS), we use ***k*-Means clustering** to group poker hands into buckets to form an abstraction for each round. Either Perfect Recall (PR) or Imperfect Recall (IR) can be used. This choice determines *k*, the number of buckets, as listed in the table.

Combining a **recall type** (PR or IR), an **early-game abstraction** (PHS, KE or KO) and an **end-game abstraction** (PHS or KO) gives a full game abstraction, which can be solved with CFR or CFR-BR to generate strategies.

Evaluating the strategies allows us to evaluate the abstraction technique.

	Perfect Recall	**Imperfect Recall**
Preflop Buckets	10	169
Flop Buckets	10*10	9000
Turn Buckets	10^2*10	9000
River Buckets	10^3*10	9000
Game Size (infosets)	57,330,780	57,331,352

In-game Performance (CFR)

		Perfect Recall (PR)						Imperfect Recall (IR)						Mean
		PHS-PHS	PHS-KO	KO-PHS	KO-KO	KE-PHS	KE-KO	PHS-PHS	PHS-KO	KO-PHS	KO-KO	KE-PHS	KE-KO	
PR	PHS-PHS		-0.9	14.3	14.5	-1.9	-2.7	-8.8	-14.8	-19.1	-22.4	-21.1	-25.7	-8.1
	PHS-KO	0.9		14.2	13.5	-0.3	-2.3	-8.7	-14.7	-17.6	-22.2	-21.0	-26.3	-7.7
	KO-PHS	-14.3	-14.2		-1.69	-15.2	-16.9	-25.7	-30.3	-32.6	-37.3	-34.9	-39.9	-23.9
	KO-KO	-14.5	-13.5	1.69		-14.5	-16.3	-25.1	-30.2	-31.8	-37.2	-34.8	-39.6	-23.3
	KE-PHS	1.9	0.3	15.2	14.5		-2.4	-6.5	-13.0	-16.1	-21.6	-18.0	-24.2	-6.4
	KE-KO	2.7	2.3	16.9	16.3	2.4		-6.6	-12.5	-15.8	-21.0	-18.6	-24.8	-5.3
IR	PHS-PHS	8.8	8.7	25.7	25.1	6.5	6.6		-5.5	-10.8	-15.0	-15.5	-18.9	1.4
	PHS-KO	14.8	14.7	30.3	30.2	13.0	12.5	5.5		-6.0	-10.4	-9.6	-14.6	7.3
	KO-PHS	19.1	17.6	32.6	31.8	16.1	15.8	10.8	6.0		-7.2	-2.6	-9.3	11.9
	KO-KO	22.4	22.2	37.3	37.2	21.6	21.0	15.0	10.4	7.2		3.8	-2.8	17.8
	KE-PHS	21.1	21.0	34.9	34.8	18.0	18.6	15.5	9.6	2.6	-3.8		-6.7	15.1
	KE-KO	25.7	26.3	39.9	39.6	24.2	24.8	18.9	14.6	9.3	2.8	6.7		21.2

Results are in milli-big-blinds/game and are accurate to 1.1 mbb/g

CFR and CFR-BR Exploitability

	CFR		CFR-BR	
	PR	IR	PR	IR
PHS-PHS	288.942	358.188	89.632	94.841
PHS-KO	289.152	318.197	90.371	85.275
KO-PHS	335.902	355.881	105.389	88.547
KO-KO	330.319	291.574	105.523	73.091
KE-PHS	**281.63**	339.872	90.720	80.557
KE-KO	282.856	**282.395**	84.039	**64.820**

In perfect recall abstractions, CFR-BR measures the ability of an abstraction to represent a Nash equilibrium. The imperfect recall strategies found by CFR-BR are often even less exploitable.

Varying the Abstraction Size

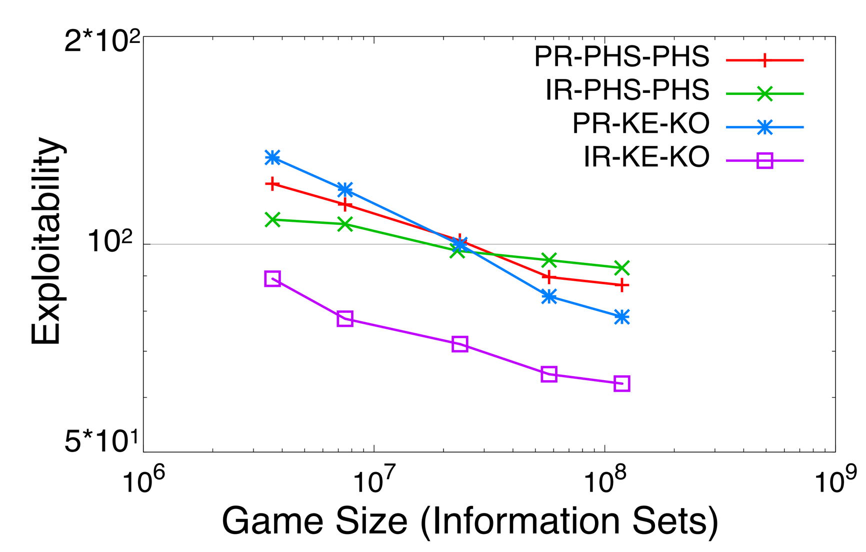

The Imperfect Recall KE-KO abstraction is less exploitable at every abstraction size measured.

University of Alberta · Alberta Innovates · compute • calcul Canada · WestGrid · NSERC CRSNG · Making IT happen Computing Science