Value Theory of Meta-Learning

Abraham Bagherjeiran
University of Houston

AAAI 2006 Doctoral Consortium

Example

	L1		L2	
	A1	A2	B1	B2
D1	96%	96%	94%	97%
D2	70%	73%	69%	72%
Complexity	5	1	1	2

Hypothetical Experiment

Better for Datasets

Better for Learners (Not in Meta-Learning)

What is Meta-Learning

Learning to Learn

Bias Learning

Transfer Learning

Ensemble Classifiers

Game Theory

What is Meta-Learning

Choose the right learner to apply to a new task.

Choose: Select one from a set of learners.

Right: Maximize the value for task and learner.

Learner to apply: Reuse an existing algorithm.

New task: Generalize to new problems as they come.

Contributions and Progress

- Game Theory for Meta-Learning ✓
- Learning Transaction Model
- Algorithm
- Applications

Game Theory

- Players
- Payoffs
  - Maximize
  - Better for One Worse for Others
- Equilibrium
  - Strategies
  - Would not Change
  - Compromise

Transformation

- Consumers
  - Datasets Select Learners
  - Do not Pick Parameters
  - Average Accuracy using Learner
- Producers
  - Learners Select Parameters
  - $$\pi_p(\mathcal{H}(p), s_{-p}) = C(p) r_p(\mathcal{H}(p))$$

Number of Consumers that Select this Producer

Preference for a Parameter

Consumers

- Different Values Same Dataset
  - Accuracy
  - TP/FP
  - Cost Matrix
- Different
  - Learner
  - Parameters

Producers

- Preference over Parameters
  - Pruned v. Not Pruned
- Consumer Demand

$$\pi_p(\mathcal{H}(p), s_{-p}) \quad = \quad C(p) r_p\left(\mathcal{H}(p)\right)$$

Game Theory for Meta-Learning

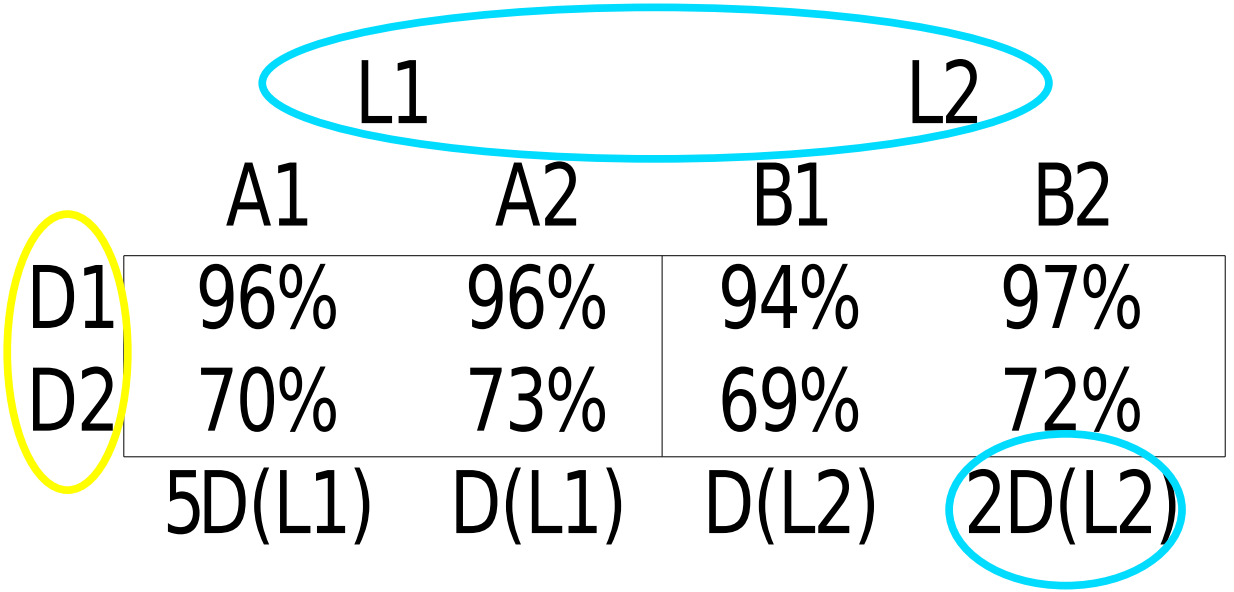

	L1		L2	
	A1	A2	B1	B2
D1	96%	96%	94%	97%
D2	70%	73%	69%	72%
	5D(L1)	D(L1)	D(L2)	2D(L2)

Producers

Consumers

Depend on Consumers

Depend on Producers

Use L2 with B2 on D1 and D2
Payoff(D1,D2,L1,L2)=(97,72,0,4)

What to Produce

What to Consume

Hard Issues

- What is value?
  - Consumer: Accuracy, Time
  - Learner: Simplicity, Popularity
- What if there is no game?

Value Theory of Meta-Learning

- Learning Transaction
  - (Dataset, Value), (Learner, Value)
  - Dataset Consumes Learner
  - Learners Have Value
- Solution
  - Estimate Payoff Function
  - Map Datasets to Learners
  - Stream of Transactions

Learn Game

- Estimate Payoffs
- Players
- Strategies

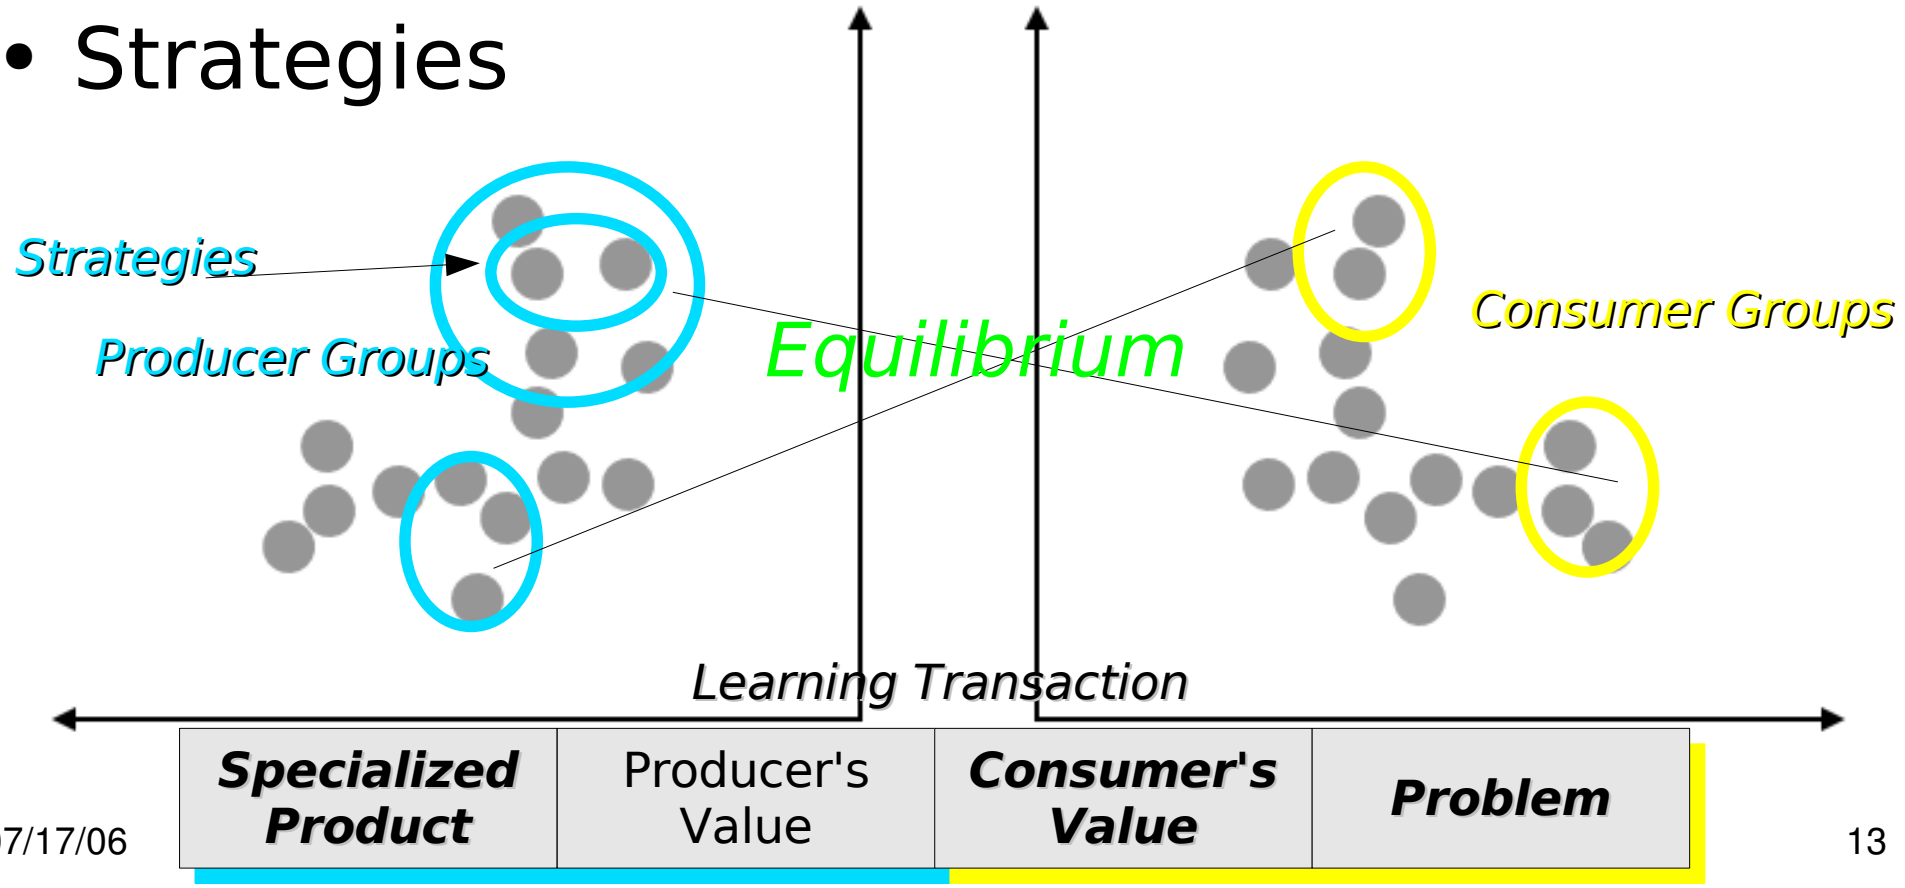

Evaluation

- Error
  - Game
  - Payoffs
  - Generalization
- Approximate Equilibrium

Contributions

- Model Meta-Learning Algorithms
- Learning Transactions
- General Problem, Many Applications

Open Questions

- Representation
  - *How to represent both task and value?*
  - How does the equilibrium handle imperfect information?
  - Are there more algorithms to try?
- Applications
  - *Can I evaluate other meta-learning with the framework?*
  - Are there real-world datasets (not meta-learning) for this?
- Theoretical
  - *Can I prove convergence to optimal game and solution?*
  - *How to handle multiple equilibria?*
  - *Is game theory the best solution to the problem?*
  - Can this help explain why a learning algorithm works?

Other People

- Ricardo Vilalta
- Christoph F. Eick
- Albert M. K. Cheng
- Venkat Subramaniam
- Nick Feltovich

Applications

- *Meta-Learning*
  - Learners, Datasets
- Recommender Systems
  - Product Lines, Niche Markets
- Information Retrieval
  - Specialized Query Service, User Groups
- Other
  - Network Control, P2P Distributed Caching

Learning Transaction Model

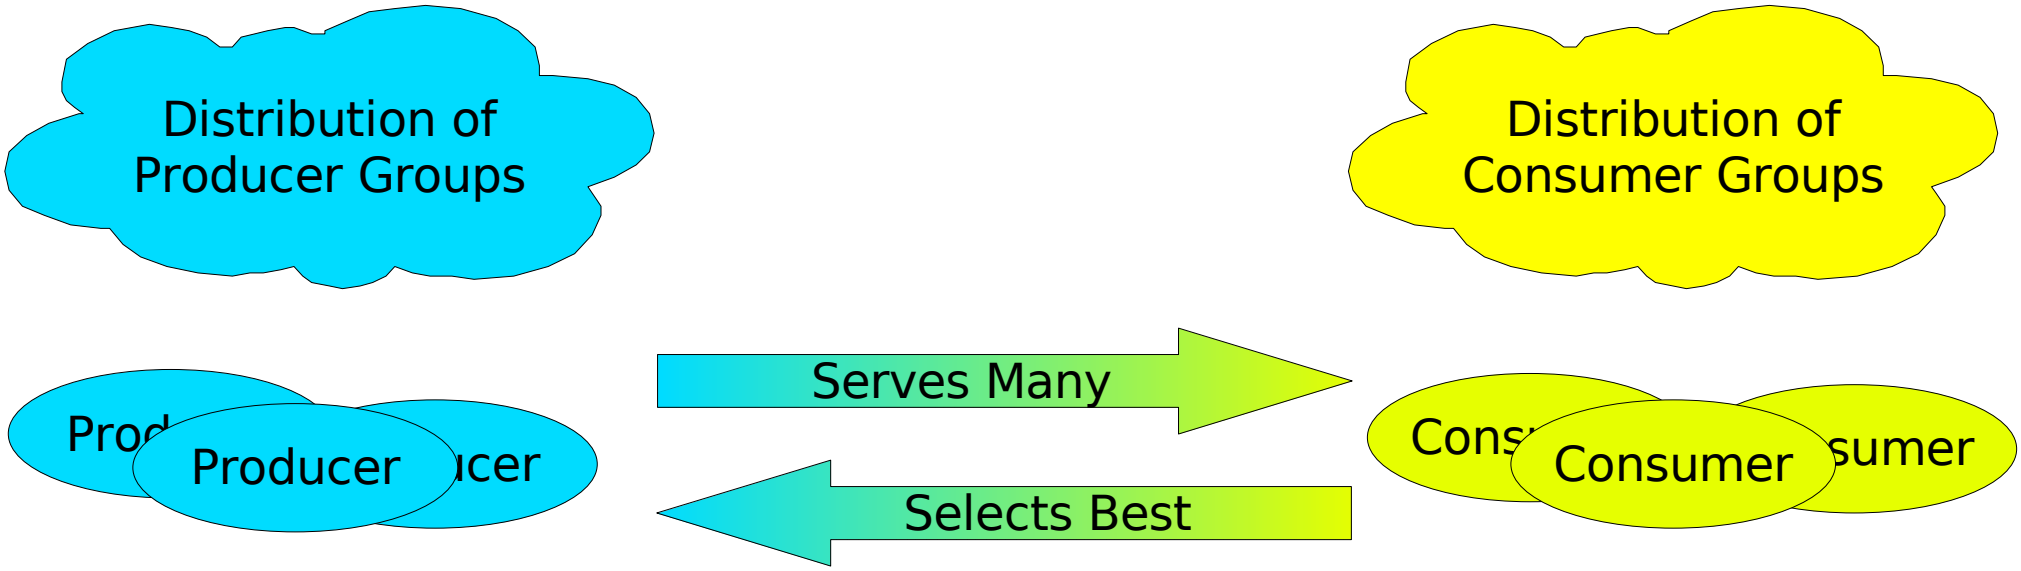

Learning Transaction

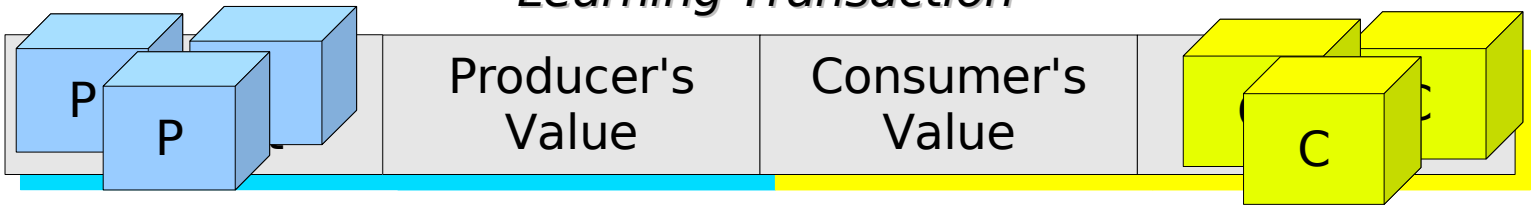

Example Outcome of Game

Learn the Game from Stream of Outcomes