CS344M
Autonomous Multiagent Systems

Todd Hester

Department of Computer Science
The University of Texas at Austin

Good Afternoon, Colleagues

Are there any questions?

Good Afternoon, Colleagues

Are there any questions?

- TAC currently
- Real-world TAC

Logistics

- FAI talk on Friday
  - Dr. Karthik Dantu (Fri, 11am, PAI 3.14)
  - Challenges in Building a Swarm of Robotic Bees

Logistics

- FAI talk on Friday
  - Dr. Karthik Dantu (Fri, 11am, PAI 3.14)
  - Challenges in Building a Swarm of Robotic Bees
- Final tournament: Monday 12/17, 2pm

Logistics

- FAI talk on Friday
  - Dr. Karthik Dantu (Fri, 11am, PAI 3.14)
  - Challenges in Building a Swarm of Robotic Bees
- Final tournament: Monday 12/17, 2pm
- Peer review process — thoughts?

Logistics

- FAI talk on Friday
  - Dr. Karthik Dantu (Fri, 11am, PAI 3.14)
  - Challenges in Building a Swarm of Robotic Bees
- Final tournament: Monday 12/17, 2pm
- Peer review process — thoughts?
- Progress reports coming back
  - Hand graded version in with your final reports

Logistics

- FAI talk on Friday
  - Dr. Karthik Dantu (Fri, 11am, PAI 3.14)
  - Challenges in Building a Swarm of Robotic Bees
- Final tournament: Monday 12/17, 2pm
- Peer review process — thoughts?
- Progress reports coming back
  - Hand graded version in with your final reports
- Final projects due in 3 weeks!

Your Progress Reports

- Overall quite good! (writing **and** content)

Your Progress Reports

- Overall quite good! (writing **and** content)
- Best ones motivate the problem before giving solutions

Your Progress Reports

- Overall quite good! (writing **and** content)
- Best ones motivate the problem before giving solutions
- Say not only what's done, but what's yet to do

Your Progress Reports

- Overall quite good! (writing **and** content)
- Best ones motivate the problem before giving solutions
- Say not only what's done, but what's yet to do
- More about what worked than what didn't

Your Progress Reports

- Overall quite good! (writing **and** content)
- Best ones motivate the problem before giving solutions
- Say not only what's done, but what's yet to do
- More about what worked than what didn't
- Clear enough for outsider to understand

Your Progress Reports

- Overall quite good! (writing **and** content)
- Best ones motivate the problem before giving solutions
- Say not only what's done, but what's yet to do
- More about what worked than what didn't
- Clear enough for outsider to understand
- Do not just paste in proposal text... modify/merge it in
  - Especially if your plans have changed
  - Report should not say what you plan to put in the report

Details

- Be specific - enough detail so that we could reimplement
  - Use pseudocode and/or diagrams

Details

- Be specific - enough detail so that we could reimplement
  - Use pseudocode and/or diagrams
- Break into sections

Details

- Be specific - enough detail so that we could reimplement
  - Use pseudocode and/or diagrams
- Break into sections
- Say up front specifically what you are doing

Details

- Be specific - enough detail so that we could reimplement
  - Use pseudocode and/or diagrams
- Break into sections
- Say up front specifically what you are doing
  - Not “working on passing”
  - But making pass decisions based on x, y, and z

Details

- Be specific - enough detail so that we could reimplement
  - Use pseudocode and/or diagrams
- Break into sections
- Say up front specifically what you are doing
  - Not "working on passing"
  - But making pass decisions based on x, y, and z
- It should not be left to the reader to figure it out

Details

- Be specific - enough detail so that we could reimplement
  - Use pseudocode and/or diagrams
- Break into sections
- Say up front specifically what you are doing
  - Not “working on passing”
  - But making pass decisions based on x, y, and z
- It should not be left to the reader to figure it out
- Can you say exactly how your work differs from baseline?

Style

- More about your approach, less about the process

Style

- More about your approach, less about the process
  - Not “What I did on summer vacation”

Style

- More about your approach, less about the process
  - Not "What I did on summer vacation"
  - Not just "we decided."
  - How? Why? What alternatives?

Style

- More about your approach, less about the process
  - Not “What I did on summer vacation”
  - Not just “we decided.”
  - How? Why? What alternatives?
  - Say where parameters came from

Style

- More about your approach, less about the process
  - Not “What I did on summer vacation”
  - Not just “we decided.”
  - How? Why? What alternatives?
  - Say where parameters came from
- Slides on resources page

Style

- More about your approach, less about the process
  - Not “What I did on summer vacation”
  - Not just “we decided.”
  - How? Why? What alternatives?
  - Say where parameters came from
- Slides on resources page
- Final projects: content matters more

Trading Agent Competition

- Put forth as a **benchmark problem** for e-marketplaces (Wellman, Wurman, et al., 2000)
- Autonomous agents act as **travel agents**

Trading Agent Competition

- Put forth as a **benchmark problem** for e-marketplaces (Wellman, Wurman, et al., 2000)

- Autonomous agents act as **travel agents**
  - **Game:** 8 *agents*, 12 min.
  - **Agent:** simulated travel agent with 8 *clients*
  - **Client:** TACtown $\leftrightarrow$ Tampa within 5-day period

Trading Agent Competition

- Put forth as a **benchmark problem** for e-marketplaces (Wellman, Wurman, et al., 2000)

- Autonomous agents act as **travel agents**
  - **Game:** 8 *agents*, 12 min.
  - **Agent:** simulated travel agent with 8 *clients*
  - **Client:** TACtown $\leftrightarrow$ Tampa within 5-day period

- **Auctions** for flights, hotels, entertainment tickets
  - **Server** maintains markets, sends prices to agents
  - Agent sends bids to server **over network**

28 Simultaneous Auctions

Flights: Inflight days 1-4, Outflight days 2-5 (8)

- Unlimited supply; prices tend to increase; immediate clear; no resale

28 Simultaneous Auctions

Flights: Inflight days 1-4, Outflight days 2-5 (8)

- Unlimited supply; prices tend to increase; immediate clear; no resale

Hotels: Tampa Towers/Shoreline Shanties days 1-4 (8)

- 16 rooms per auction; 16th-price ascending auction; quote is ask price; no resale
- Random auction closes minutes 4 – 11

28 Simultaneous Auctions

Flights: Inflight days 1-4, Outflight days 2-5 (8)

- Unlimited supply; prices tend to increase; immediate clear; no resale

Hotels: Tampa Towers/Shoreline Shanties days 1-4 (8)

- 16 rooms per auction; 16th-price ascending auction; quote is ask price; no resale
- Random auction closes minutes 4 – 11

Entertainment: Wrestling/Museum/Park days 1-4 (12)

- Continuous double auction; initial endowments; quote is bid-ask spread; resale allowed

Client Preferences and Utility

Preferences: randomly generated per client

- Ideal arrival, departure days
- Good Hotel Value
- Entertainment Values

Client Preferences and Utility

Preferences: randomly generated per client

- Ideal arrival, departure days
- Good Hotel Value
- Entertainment Values

Utility: 1000 (if valid) – travel penalty + hotel bonus + entertainment bonus

Client Preferences and Utility

Preferences: randomly generated per client

- Ideal arrival, departure days
- Good Hotel Value
- Entertainment Values

Utility: 1000 (if valid) – travel penalty + hotel bonus + entertainment bonus

Score: Sum of client utilities – expenditures

Allocation

$$
\begin{aligned}
G &\equiv \text{complete allocation of goods to clients} \\
v(G) &\equiv \text{utility of } G - \text{cost of needed goods} \\
G^* &\equiv \text{argmax } v(G)
\end{aligned}
$$

Allocation

$$
\begin{aligned}
G &\equiv \text{complete allocation of goods to clients} \\
v(G) &\equiv \text{utility of } G - \text{cost of needed goods} \\
G^* &\equiv \text{argmax } v(G)
\end{aligned}
$$

Given holdings and prices, find G^*

Allocation

$$\begin{aligned} G &\equiv \text{complete allocation of goods to clients} \\ v(G) &\equiv \text{utility of } G - \text{cost of needed goods} \\ G^* &\equiv \text{argmax } v(G) \end{aligned}$$

$$\boxed{\text{Given holdings and prices, find } G^*}$$

- General allocation NP-complete
  - Tractable in TAC: mixed-integer LP (ATTac-2000)
  - Estimate $v(G^*)$ quickly with LP relaxation

Allocation

$$G \equiv \text{complete allocation of goods to clients}$$
$$v(G) \equiv \text{utility of } G - \text{cost of needed goods}$$
$$G^* \equiv \text{argmax } v(G)$$

Given holdings and prices, find G^*

- General allocation NP-complete
  - Tractable in TAC: mixed-integer LP (ATTac-2000)
  - Estimate $v(G^*)$ quickly with LP relaxation

Prices known $\Rightarrow$ G^* known $\Rightarrow$ optimal bids known

High-Level Strategy

- Learn model of expected hotel price

High-Level Strategy

- Learn model of expected hotel price distributions

High-Level Strategy

- Learn model of expected hotel price distributions
- For each auction:
  - Repeatedly sample price vector from distributions

High-Level Strategy

- Learn model of expected hotel price distributions

- For each auction:
  - Repeatedly sample price vector from distributions
  - Bid avg marginal expected utility: $v(G_w^*) - v(G_l^*)$

High-Level Strategy

- Learn model of expected hotel price distributions

- For each auction:
  - Repeatedly sample price vector from distributions
  - Bid avg marginal expected utility: $v(G_w^*) - v(G_l^*)$

- Bid for all goods — not just those in G^*

High-Level Strategy

- Learn model of expected hotel price distributions
- For each auction:
  - Repeatedly sample price vector from distributions
  - Bid avg marginal expected utility: $v(G_w^*) - v(G_l^*)$
- Bid for all goods — not just those in G^*

Goal: analytically calculate optimal bids

Hotel Price Prediction

- **Features:**
  - Current hotel and flight prices
  - Current time in game
  - Hotel closing times
  - Agents in the game (when known)
  - Variations of the above

Hotel Price Prediction

- **Features:**
  - Current hotel and flight prices
  - Current time in game
  - Hotel closing times
  - Agents in the game (when known)
  - Variations of the above

- **Data:**
  - Hundreds of seeding round games

Hotel Price Prediction

- **Features:**
  - Current hotel and flight prices
  - Current time in game
  - Hotel closing times
  - Agents in the game (when known)
  - Variations of the above

- **Data:**
  - Hundreds of seeding round games
  - Assumption: similar economy

Hotel Price Prediction

- **Features:**
  - Current hotel and flight prices
  - Current time in game
  - Hotel closing times
  - Agents in the game (when known)
  - Variations of the above
- **Data:**
  - Hundreds of seeding round games
  - Assumption: similar economy
  - Features $\mapsto$ actual prices

The Learning Algorithm

- $X \equiv$ feature vector $\in \mathbb{R}^n$
- $Y \equiv$ closing price $-$ current price $\in \mathbb{R}$

The Learning Algorithm

- $X \equiv$ feature vector $\in \mathbb{R}^n$
- $Y \equiv$ closing price $-$ current price $\in \mathbb{R}$
- Break Y into $k \approx 50$ cut points $b_1 \leq \cdots \leq b_k$

The Learning Algorithm

- $X \equiv$ feature vector $\in \mathbb{R}^n$
- $Y \equiv$ closing price $-$ current price $\in \mathbb{R}$
- Break Y into $k \approx 50$ cut points $b_1 \leq \cdots \leq b_k$
- For each b_i, estimate probability $Y \geq b_i$, given X

The Learning Algorithm

- $X \equiv$ feature vector $\in \mathbb{R}^n$
- $Y \equiv$ closing price $-$ current price $\in \mathbb{R}$
- Break Y into $k \approx 50$ cut points $b_1 \leq \cdots \leq b_k$
- For each b_i, estimate probability $Y \geq b_i$, given X
  - Say X belongs to class C_i if $Y \geq b_i$

The Learning Algorithm

- $X \equiv$ feature vector $\in \mathbb{R}^n$
- $Y \equiv$ closing price $-$ current price $\in \mathbb{R}$
- Break Y into $k \approx 50$ cut points $b_1 \leq \cdots \leq b_k$
- For each b_i, estimate probability $Y \geq b_i$, given X
  - Say X belongs to class C_i if $Y \geq b_i$
  - k-class problem: each example in many classes

The Learning Algorithm

- $X \equiv$ feature vector $\in \mathbb{R}^n$
- $Y \equiv$ closing price $-$ current price $\in \mathbb{R}$
- Break Y into $k \approx 50$ cut points $b_1 \leq \cdots \leq b_k$
- For each b_i, estimate probability $Y \geq b_i$, given X
  - Say X belongs to class C_i if $Y \geq b_i$
  - k-class problem: each example in many classes
  - Use **BoosTexter** (boosting (Schapire, 1990))

The Learning Algorithm

- $X \equiv$ feature vector $\in \mathbb{R}^n$
- $Y \equiv$ closing price $-$ current price $\in \mathbb{R}$
- Break Y into $k \approx 50$ cut points $b_1 \leq \cdots \leq b_k$
- For each b_i, estimate probability $Y \geq b_i$, given X
  - Say X belongs to class C_i if $Y \geq b_i$
  - k-class problem: each example in many classes
  - Use **BoosTexter** (boosting (Schapire, 1990))
- Can convert to estimated distribution of $Y|X$

The Learning Algorithm

- $X \equiv$ feature vector $\in \mathbb{R}^n$
- $Y \equiv$ closing price $-$ current price $\in \mathbb{R}$
- Break Y into $k \approx 50$ cut points $b_1 \leq \cdots \leq b_k$
- For each b_i, estimate probability $Y \geq b_i$, given X
  - Say X belongs to class C_i if $Y \geq b_i$
  - k-class problem: each example in many classes
  - Use **BoosTexter** (boosting (Schapire, 1990))
- Can convert to estimated distribution of $Y|X$

$$\boxed{\textbf{New algorithm for conditional density estimation}}$$

Hotel Expected Values

- Repeat until time bound, for each hotel:

  1. Assume this hotel closes next

Hotel Expected Values

- Repeat until time bound, for each hotel:

1. Assume this hotel closes next
2. Sample prices from predicted price distributions

Hotel Expected Values

- Repeat until time bound, for each hotel:

1. Assume this hotel closes next
2. Sample prices from predicted price distributions
3. Given these prices compute $V_0, V_1, \ldots V_8$
   - $V_i = v(G^*)$ if own **exactly** i of the hotel
   - $V_0 \leq V_1 \leq \ldots \leq V_8$

Hotel Expected Values

- Repeat until time bound, for each hotel:
  1. Assume this hotel closes next
  2. Sample prices from predicted price distributions
  3. Given these prices compute $V_0, V_1, \ldots V_8$
     - V_i = $v(G^*)$if own **exactly** i of the hotel
     - $V_0 \leq V_1 \leq \ldots \leq V_8$
- Value of ith copy is avg($V_i - V_{i-1}$)

Other Uses of Sampling

Flights: Cost/benefit analysis for postponing commitment

Other Uses of Sampling

Flights: Cost/benefit analysis for postponing commitment

Cost: Price expected to rise over next n minutes

Benefit: More price info becomes known

- Compute expected marginal value of buying some different flight

Other Uses of Sampling

Flights: Cost/benefit analysis for postponing commitment

Cost: Price expected to rise over next n minutes

Benefit: More price info becomes known

- Compute expected marginal value of buying some different flight

Entertainment: Bid more (ask less) than expected value of having one more (fewer) ticket

Finals

Team	Avg.	Adj.	Institution
ATTac	3622	4154	AT&T
livingagents	3670	4094	Living Systems (Germ.)
whitebear	3513	3931	Cornell
Urlaub01	3421	3909	Penn State
Retsina	3352	3812	CMU
CaiserSose	3074	3766	Essex (UK)
Southampton	3253*	3679	Southampton (UK)
TacsMan	2859	3338	Stanford

- ATTac improves over time
- livingagents is an open-loop strategy

Controlled Experiments

- *ATTac$_s$*: "‘full-strength” agent based on boosting

Controlled Experiments

- *ATTac*$_s$: "`full-strength" agent based on boosting
- *SimpleMean*$_s$: sample from empirical distribution (previously played games)

Controlled Experiments

- *ATTac*$_s$: "`full-strength" agent based on boosting
- *SimpleMean*$_s$: sample from empirical distribution (previously played games)
- *ConditionalMean*$_s$: condition on closing time

Controlled Experiments

- *ATTac*$_s$: "`full-strength" agent based on boosting
- *SimpleMean*$_s$: sample from empirical distribution (previously played games)
- *ConditionalMean*$_s$: condition on closing time
- *ATTac*$_{ns}$, *ConditionalMean*$_{ns}$, *SimpleMean*$_{ns}$: predict expected value of the distribution

Controlled Experiments

- *ATTac*$_s$: "'full-strength" agent based on boosting
- *SimpleMean*$_s$: sample from empirical distribution (previously played games)
- *ConditionalMean*$_s$: condition on closing time
- *ATTac*$_{ns}$, *ConditionalMean*$_{ns}$, *SimpleMean*$_{ns}$: predict expected value of the distribution
- *CurrentPrice*: predict no change

Controlled Experiments

- *ATTac*$_s$: "`full-strength" agent based on boosting
- *SimpleMean*$_s$: sample from empirical distribution (previously played games)
- *ConditionalMean*$_s$: condition on closing time
- *ATTac*$_{ns}$, *ConditionalMean*$_{ns}$, *SimpleMean*$_{ns}$: predict expected value of the distribution
- *CurrentPrice*: predict no change
- *EarlyBidder*: motivated by TAC-01 entry livingagents

Controlled Experiments

- *ATTac*$_s$: "`full-strength" agent based on boosting
- *SimpleMean*$_s$: sample from empirical distribution (previously played games)
- *ConditionalMean*$_s$: condition on closing time
- *ATTac*$_{ns}$, *ConditionalMean*$_{ns}$, *SimpleMean*$_{ns}$: predict expected value of the distribution
- *CurrentPrice*: predict no change
- *EarlyBidder*: motivated by TAC-01 entry livingagents
  – Immediately bids high for G^* (with *SimpleMean*$_{ns}$)
  – Goes to sleep

Stability

- 7 *EarlyBidder*'s with 1 *ATTac*

Agent	Score	Utility
ATTac	2431 ± 464	8909 ± 264
EarlyBidder	-4880 ± 337	9870 ± 34

Stability

- 7 *EarlyBidder*'s with 1 *ATTac*

Agent	Score	Utility
ATTac	2431 ± 464	8909 ± 264
EarlyBidder	-4880 ± 337	9870 ± 34

- 7 *ATTac*'s with 1 *EarlyBidder*

Agent	Score	Utility
ATTac	2578 ± 25	9650 ± 21
EarlyBidder	2869 ± 69	10079 ± 55

Stability

- 7 *EarlyBidder*'s with 1 *ATTac*

Agent	Score	Utility
ATTac	2431 ± 464	8909 ± 264
EarlyBidder	-4880 ± 337	9870 ± 34

- 7 *ATTac*'s with 1 *EarlyBidder*

Agent	Score	Utility
ATTac	2578 ± 25	9650 ± 21
EarlyBidder	2869 ± 69	10079 ± 55

EarlyBidder gets more utility; *ATTac* pays less

Results

- *Phase I* : Training from TAC-01 (seeding round, finals)

Results

- *Phase I* : Training from TAC-01 (seeding round, finals)
- *Phase II* : Training from TAC-01, phases I, II

Results

- *Phase I* : Training from TAC-01 (seeding round, finals)
- *Phase II* : Training from TAC-01, phases I, II
- *Phase III* : Training from phases I – III

Results

- *Phase I* : Training from TAC-01 (seeding round, finals)
- *Phase II* : Training from TAC-01, phases I, II
- *Phase III* : Training from phases I – III

Agent	*Relative Score*	
	Phase I	*Phase III*
$ATTac_{ns}$	105.2 ± 49.5 (2)	166.2 ± 20.8 (1)
$ATTac_{s}$	27.8 ± 42.1 (3)	122.3 ± 19.4 (2)
EarlyBidder	140.3 ± 38.6 (1)	117.0 ± 18.0 (3)
$SimpleMean_{ns}$	-28.8 ± 45.1 (5)	-11.5 ± 21.7 (4)
$SimpleMean_{s}$	-72.0 ± 47.5 (7)	-44.1 ± 18.2 (5)
$ConditionalMean_{ns}$	8.6 ± 41.2 (4)	-60.1 ± 19.7 (6)
$ConditionalMean_{s}$	-147.5 ± 35.6 (8)	-91.1 ± 17.6 (7)
CurrentPrice	-33.7 ± 52.4 (6)	-198.8 ± 26.0 (8)

Other TAC competitions

- Supply Chain Management
- Ad Auctions
- Power

Discussion

- Are these agents useful for the real version of these tasks?

Discussion

- Are these agents useful for the real version of these tasks?
- What can we learn from these competitions?

Discussion

- Are these agents useful for the real version of these tasks?
- What can we learn from these competitions?
- General strategy that works well?

Last-minute bidding (R,O, 2001)

- eBay: first-price, ascending auction
- Amazon: auction extended if bid in last 10 minutes
- eBay: bots exist to incrementally raise your bid to a maximum

- Still people *snipe*. Why?
  - There's a risk that the bid might not make it
  - However, common-value $\Longrightarrow$ bid conveys info
  - Late-bidding can be seen as implicit collusion
  - Or . . . , lazy, unaware, etc. (Amazon and eBay)
- Finding: more late-bidding on eBay,
  - even more on antiques rather than computers

Small design-difference matters

Late Bidding as Best Response

- Good vs. incremental bidders
  - They start bidding low, plan to respond
  - Doesn't give them time to respond
- Good vs. other snipers
  - Implicit collusion
  - Both bid low, chance that one bid doesn't get in
- Good in common-value case
  - protects information

Overall, the analysis of multiple bids supports the hypothesis that last-minute bidding arises at least in part as a response by sophisticated bidders to unsophisticated incremental bidding.