ATTac-2000: An Adaptive Autonomous Bidding Agent

Greg Miller, Cody Musilek,
John Steinbach

Source Paper

Stone, P., M. L. Littman, S. Singh, and M. Kearns (2001). ATTac-2000: An Adaptive Autonomous Bidding Agent, *Journal of Artificial Intelligence Research*, **15**:189-206.

Outline

- Trading Agent Competition (TAC)
- ATTac-2000
  - Bidding Strategy
    - Flights, Hotels, Entertainment Tickets
  - Allocation Strategy
  - Adaptivity
- Results
  - Competition
  - Controlled Testing
- Our Critiques
- Q&A

Trading Agent Competition (TAC)

- Requires *autonomous bidding agents* to buy and sell multiple interacting goods in auctions of different types
- Preliminary round and many practice games before finals
  - Developers can change strategies in between
- Game instance pits 8 agents against one another
- Each agent is a travel agent with 8 clients
  - Clients want to travel during a common 5-day period
  - Clients characterized by random set of preferences
    - Arrival/Departure Dates, Hotel, Entertainment
- Must construct travel package for each client
- Agent's score in game instance is difference between sum of clients' utilities and agent's total expenditure

TAC continued

- Agents buy flights, hotel rooms and entertainment tickets via auctions

- Flights (8 auctions)
  - Inflights days 1-4, outflights days 2-5, separate auctions
  - Unlimited supply
  - Price varies from $150 to $600, randomly changes $0-10
  - If a bid is higher than ask price, ticket is immediately sold to that agent for ask price; no resale

TAC continued

- Hotels (8 auctions)
  - Boston Grand Hotel (BGH) or Le Fleabag Hotel (LFH)
    - Each has 16 rooms on each day 1-4
  - Sold in 16th price English auction
  - Must bid higher than current price, no withdraw/resale
  - Sold when auction closes, can close from inactivity
    - Prevents waiting until end of game to bid

TAC continued

- Entertainment Tickets (12 auctions)
  - Baseball, symphony and theater tickets
  - Continuous double auctions - agents can buy and sell
  - Sold immediately when bid at least as high as ask price
  - Sell price is ask price, not bid price
  - Each agent starts with a random endowment of tickets
  - Bid withdrawal and ticket resale permitted

TAC continued

- Clients have parameters for ideal arrival day, IAD (1-4); ideal departure day, IDD (2-5); grand hotel value, GHV ($50-150); and entertainment values, EV ($0-200) for each type of ticket
- The client's utility is defined by the following equation:
  utility = 1000 - travelpenalty + hotelbonus + funbonus
  - travelpenalty = 100(|AD - IAD| + |DD - IDD|)
  - hotelbonus = GHV if client is in BGH, 0 otherwise
  - funbonus = sum of EVs for each ticket type assigned

- Agent's final score is the sum of the clients' utilities minus the agent's expenditures

ATTac-2000

- Finished first in the Trading Agent Competition using a principled bidding strategy, which included several elements of adaptivity.

- Had the flexibility to cope with the wide variety of possible scenarios in competition.

ATTac Bidding Strategy

- Robust to parameter space and opponent strategies
- At every bidding opportunity, ATTac begins by computing the most profitable allocation of goods to clients (denoted by G*), given the goods that are currently owned and the current prices of hotels and flights.
- High-level bidding strategy based on two observations:
  - the expected change in price for airline auctions is $0
  - as the game proceeds, the hotel prices approach the eventual closing prices

Bidding Strategy contd..

- ATTac aims to delay most of its purchases, particularly airline purchases, until late in the game.
- Attempts to delay "committing" to the current G* for as long as possible.
- ATTac accomplishes this delay of commitment by bidding in two different modes: *passive* and *active*.
  - Starts out *passive*, switches to *active* when time is running out.

Flights

- Unlimited supply means no competition from other agents
- Passive mode: Does not bid
- Active mode: buys all currently unowned airline tickets needed for the current G*

Hotels

- Passive mode: Bids in hotel auctions to either win them cheaply or prevent the auctions from closing early
  - Tries to acquire *n* rooms where *n* depends on the number of rooms of a specific type needed for *G** and the prices
  - Bids $1 above current ask price, can risk $40-50 per room type for benefit of flexibility later in the game

- Active mode: Bids on rooms based on their marginal value
  - Bids a price of $V(G^*) - V(G^{*'}_c)$ for hotel rooms assigned to client *c* in *G**
    - $V(G^*)$ is income from all clients minus cost of yet-to-be-acquired goods
    - $V(G^{*'}_c)$ is value of optimal allocation should client *c* fail to get its hotel rooms

Entertainment Tickets

- Assumes opponent buy/sell price remains constant in a game
  - Gradually decreases/increases its bid over time
- On every iteration, ATTac places a buy bid for each type of ticket and a sell bid for each type of ticket it currently owns
  - In passive mode, for each owned ticket, sets sell bid at optimistic price and gradually lowers, but raises to 1¢ lower than current bid if current bid is higher than sell bid
  - In active mode, offers to sell any unused ticket for $30
  - Buy bids based on increased value derived from owning that type of ticket
    - Similar to active mode hotel purchasing but also includes a variable based on time remaining

Allocation Strategy

- Uses an integer linear programming approach (ILPA) instead of a greedy approach used by most other participants
  - Defines a set of variables, constraints on these variables, and an objective function
  - ATTac was able to compute optimal final allocations in every game instance during the finals
  - Switches to modified greedy solution as a fall-back if this approach takes longer than 6 seconds to compute

$$\sum_{c,f} u_P(c,f)P(c,f) + \sum_{c,e} u_E(c,e)E(c,e) - \sum_{d\in\{2,3,4,5\}} p_{DD}(d)B_{DD}(d) - \sum_{d\in\{1,2,3,4\},r\in\{\mathrm{BGH},\mathrm{LFI},AD\}} p_r(d)B_r(d)$$

Allocation Strategy contd.

- Many constraints were applied for this situation, including:
  - No client gets more than one travel package
  - Demand for resources from selected travel packages must not exceed sum of owned & bought resources
  - Total quantity of each entertainment ticket allocated does not exceed what is owned
  - An entertainment ticket can only be used if its day is between arrival and departure day of the selected package
  - Each client can only use one entertainment ticket per day
  - Each client can only use each type of entertainment ticket once
  - All variables are integers
- Solution to ILPA is value-maximizing allocation of owned resources plus list of resources needing to be purchased

Adaptivity

- In TAC game instance, only information available is ask prices
- Lack of within-game info precluded competitors from using detailed models of opponent strategies in decision making
- ATTac instead adapts its behavior in 3 different ways
  - Adaptable timing of bidding modes, adaptable allocation strategy, and adaptable hotel bidding
- Timing of bidding modes
  - Decides when to switch from passive to active based on observed server latency during game instance
- Allocation
  - Adapts allocation strategy based on amount of time it takes for ILPA to determine optimal allocations

Adaptibility contd.

- Hotel Bidding
  - Predicts closing prices of hotel auctions based on their closing prices in previous games
  - Divided 8 hotel rooms into 4 equivalence classes, exploiting symmetries (equal demand on days 1&4 and 2&3), assigned priors to expected closing prices, and adjusted these priors based on observed closing prices
  - When actual price was lower than predicted, it used the predicted values for computing allocation values
  - Looked for games with 3+ "high-bidders" to use predicted closing prices; 2 or less failed to cause prices to skyrocket
    - High-bidders bid their marginal utilities on hotel rooms
  - Extremely beneficial when prices escalate, no significantly degraded performance when they don't

Results - Competition

- Agents and conditions constantly changing - not a controlled testing environment
- Scores varied widely from -3000 to over 4500 (3000 to 4000 is considered good) with an average of 2700
  - Hadn't implemented adaptive timing of bidding modes
  - Occasionally failed to place bids in time due to lag
    - Fixed by implementing adaptive timing
- Adaptive allocation strategy never came into play, but adaptive hotel bidding did play a big role
  - Rivaled other best teams in early games where hotel prices stayed low, excelled in final games when hotel prices rose to high levels
- Ended up with highest average score and lowest standard deviation - consistently high scores

Results - Competition

Rank	Team	Avg. Score	Std. Dev.	Institution
1	ATTac-2000	3398	443	AT&T Labs – Research
2	RoxyBot	3283	545	Brown University, NASA Ames Research
3	aster	3068	493	STAR Lab, InterTrust Technologies
4	umbctac1	3051	1123	University of Maryland at Baltimore County
5	ALTA	2198	1328	Artificial Life, Inc.
6	m_rajatish	1873	1657	University of Tulsa
7	RiskPro	1570	1607	Royal Inst. Technology, Stockholm University
8	T1	1167	1593	Swedish Inst. Computer Science, Industilogik

Table 4: The scores of the 8 TAC finalists in the semi-finals and finals (13 games).

Results - Controlled Testing

- Ran several game instances against two variants of itself
  - High-bidder: always computed *G** with current hotel prices
  - Low-bidder: same as high-bidder, but only bid for hotel rooms at $50 over the current asking price
- Setup was same as in TAC - 8 agents competing
- With 7 high-bidders at least 1 hotel price skyrockets
- With 7 low-bidders hotel price never skyrocket
- Consistently beat all other agents in simulations
  - Many high-bidders results in many large negative scores
- Always used adaptive hotel price expectations, even when only 2 high-bidders were present

Results - Controlled Testing contd.

#high	agent 2	agent 3	agent 4	agent 5	agent 6	agent 7	agent 8
7 (14)	⟵	9526	⟶	⟶	⟶	⟶	⟶
6 (87)	⟵	10679	⟶	⟶	⟶	⟶	1389
5 (84)	⟵	10310	⟶	⟶	⟶	⟵	2650
4 (48)	⟵	10005	⟶	⟶	⟵	⟵	4015
3 (21)	⟵	5067	⟶	⟵	⟵	⟵	3639
2 (282)	⟵	*209*	⟵	⟵	⟵	⟵	2710

- First column shows number of high-bidders and in parentheses number of games simulated
- Main section shows difference between ATTac's score and average score of all high-bidders (left) and low-bidders (right)

Authors Conclusion

- While it was a successful event, some minor improvements would increase its interest from a multiagent learning perspective.
  - No incentive to buy airline tickets until the end of the game.
  - Impossible to observe the bidding patterns of individual agents
- ATTac plans to participate in future TACs

Our Critique - The Good

- ATTac's bidding strategy proved to be most dominant
- Potentially could use ATTac's strategy for real life travel agents
- Their integer linear programming approach was more reliable than the greedy methods of other agents.
- They had a greedy algorithm to fall back on if the linear programming approach was taking too long.

Our Critique - The Bad

- Strategies were tailored to the TAC, it might be difficult to apply them to other situations.

- Some threshold values were chosen arbitrarily.
  - Could have adjusted them during the rounds or done more pre-game experiments.

Questions?