

Estimating Bidders' Valuation Distributions in Online Auctions

Albert Xin Jiang, Kevin Leyton-Brown
Department of Computer Science
University of British Columbia

Bidding Agents

- Given a valuation function, compute a bidding strategy that maximizes EU
 - notwithstanding “Wilson Doctrine”: mechanisms should be *detail-free*
 - Motivating example: how should agents behave in **a sequence of eBay auctions?**
- **Game Theoretic Approach** [Milgrom & Weber, 1982], much subsequent work from econ.
 - model the situation as a Bayesian game
 - compute and then play a Bayes-Nash equilibrium of the game
 - when other bidders’ valuations are not known, estimate them from history
 - drawbacks:
 - rationality of other agents may be in doubt
 - intractability of computing equilibrium
 - multiple equilibria
- **Decision Theoretic Approach** [Boutilier et al. 1999; Byde 2002; Stone et al. 2003; Greenwald & Boyan 2004; MacKie-Mason et al. 2004; Osepayshvili et al. 2005]
 - learn the behavior of other bidders from historical data
 - treat other bidders as part of the environment
 - play an optimal strategy in the resulting single-agent decision problem

Learning Valuation/Price Distributions

- Whether the GT or DT approach is taken, a shared subproblem is using historical data to **estimate distribution of bidders' bid amounts** or valuations
- [Athey & Haile, 2000], various other papers in econometrics:
 - assume that bidders are perfectly rational and follow **equilibrium strategies**
 - estimation of valuation distributions in various auction types given observed bids
- [Byde, 2002], [Stone *et al.* 2003], [Greenwald & Boyan, 2004], [MacKie-Mason *et al.* 2004], [Osepayshvili *et al.* 2005]:
 - estimate the **distribution of the final prices** in (e.g.) English auctions based on selling price and number of agents
- [Boutilier *et al.* 1999]:
 - a decision-theoretic MDP approach to bidding in sequential first-price auctions for complementary goods
 - for the case where these sequential auctions are repeated, discusses learning a distribution of other agents' highest bid for each good, based on winning bids
 - **uses EM**: the agent's own bid wins, hiding the highest bid by other agents

Talk Outline

1. Background

2. Online Auction Model and Learning Problem


3. Bidding in Sequential Auctions

4. Experimental Evaluation

Online Auction Model

- A (possibly repeated) online English auction such as eBay
 - m potential bidders, with m drawn from a distribution $g(m)$
 - let n denote the number of bidders who place (accepted) bids in the auction
 - each bidder i has an independent private valuation drawn from distribution $f(v)$
- **Bidding dynamics**
 - start with reserve price of zero
 - bidders sequentially place proxy bids (each bidder gets only one bid)
 - auctioneer maintains current price: second-highest proxy amount declared so far
 - if a new bid is less than the current price, it is dropped
- **Bidding history**
 - some bidders' proxy bid amounts will be **perfectly observed** (denote this set of bids x_o)
 - any bidder who placed a proxy bid and was outbid ($n-1$ such bidders)
 - however, some bids will be **hidden** (denote this set x_h)
 - highest bid (one bidder)
 - revealed only up to the second-highest bidder's proxy amount
 - any bid which was lower than the current price when it was placed ($m - n$ bidders)
 - either the bidder leaves or the bid is rejected

Bidding Example




highest
price

Learning the Distributions $f(v)$ and $g(m)$

- Data: a set of **auction histories**
 - number of bidders and bids distributed identically in each auction
- **Simple technique** for estimating $f(v)$ and $g(m)$:
 - ignore hidden bids, considering only x_o and n from each auction
 - use any standard density estimation technique to learn the distributions
 - essentially this is the straightforward price estimation technique described earlier
- Problem:
 - the simple technique **ignores the hidden bids** and so introduces bias
 - $g(m)$ will be skewed towards small values because $n \leq m$
 - $f(v)$ may be
 - skewed towards small values because it ignores the winning bid
 - skewed towards large values because ignores dropped, losing bids

EM Algorithm

- Solution: use EM to account for hidden bids
 - similar in spirit to the approach described above by Boutilier *et al.* (1999)
 - however, in our setting some losing bids are also hidden; the number of bidders is uncertain; expected number of hidden bids depends on x_o and $f(v)$
- E step: generate the missing data given estimates of f' , g' and bidding model
 - for each observation x_o , repeat until N samples of x_h have been generated:
 - sample m from $g'(m | m \geq n)$
 - simulate bidding process until $m - n + 1$ bids have been generated:
 - Draw a sample from $f'(v)$ to represent a new bid
 - If the sampled bid exceeds the next bid in x_o , replace the bid with the next bid from x_o .
Otherwise, add the sampled bid to x_h
 - if x_h does not contain exactly one bid that exceeds the highest bid in x_o , reject sample
- M step:
 - update $f'(v)$ and $g'(m)$ to maximize the likelihood of the bids $x_o \cup x_h$
 - depends on functional form of f' , g' ; either analytic or using e.g. simulated annealing

Learning $f(v)$ and $g(m)$ in a Game Theoretic Setting

- The approach described above is decision-theoretic
- What if we want to take a **game-theoretic approach**?
 - Athey & Haile, (2000) discuss estimation in the game theoretic setting
 - however, they generally assume that number of bidders is known
 - brief discussion of unknown number of bidders, but not relevant to our online auction setting
 - let $f(v)$ be the distribution of bidder's valuations (instead of bid amounts)
 - $g(m)$ remains the distribution of number of bidders, as before
 - given a bidder's valuation v , what is his bid amount?
 - solve for Bayes-Nash equilibrium of the auction game: bid function $b(v | f, g)$
- **EM algorithm** to estimate f and g in a GT setting:
 - E step: for each sample given observation x_o :
 - sample m from $g'(m | m \geq n)$
 - compute observed bidders' valuations v_o from x_o by inverting the bid function
 - generate new bidders with valuations v_h who place hidden bids $x_h = b(v_h | f', g')$
 - simulate the auction until $m - n + 1$ bids are generated, where exactly one hidden bid is higher than the highest observed bid
 - M step: update f' and g' to maximize likelihood of the valuations $v_o \cup v_h$

Talk Outline

1. Background

2. Online Auction Model and Learning Problem

3. Building an Agent

4. Experimental Evaluation

Building an Agent

- Consider the construction of a decision-theoretic agent to participate in a finite **sequence of auctions** (under our online auction model)
 - given estimates $f'(v)$ and $g'(m)$, what are the optimal bidding strategies?
- **Auction environment**
 - k sequential, single-good online auctions for possibly non-identical goods
 - we want only one item
 - e.g. buying a Playstation 2 from eBay, where such auctions are held regularly
 - denote our valuation for the item in auction j as v_j and our bid as b_j
 - let U_j denote expected payoff at time j , conditional on not having won already
 - a function of our valuations for the goods in the auctions j, \dots, k
- Greenwald & Boyan (2004) and Arora *et al.* (2003) analyzed similar domains
 - using similar reasoning, we derive the **optimal bidding strategy** for our model

Computing the Optimal Strategy

- **Optimal bidding:** $b_j^* = v_j - U_{j+1}^*(v_{j+1}, \dots, v_k)$
 - U_{j+1}^* is the EU of the bidding strategy that maximizes U_{j+1} (derived in the paper)
$$U_{j+1}(b_{j+1}, \dots, b_k, v_{j+1}, \dots, v_k) = \int_{-\infty}^{b_{j+1}} (v_{j+1} - x) f_{j+1}^1(x) dx$$
$$+ (1 - F_{j+1}^1(b_{j+1})) U_{j+2}(b_{j+2}, \dots, b_k, v_{j+2}, \dots, v_k)$$
 - first term: payoff from current auction; second term: payoff from future auctions
 - note that U_{j+1} depends on the distribution of the highest bid:
$$F_j^1(x) = \sum_{m=2}^{\infty} g_j(m) (F_j(x))^m$$
 - ...and that F_j^1 depends in turn on $f(v)$, $g(m)$
 - thus we must estimate $f(v)$, $g(m)$ to build a decision theoretic agent in this setting
- Our agent computes U_{j+1}^* by approximating an integral using Monte Carlo sampling, again relying on our model of the auction

Elaborations

- Auctions that **overlap in time**
 - note that while the optimal bid in auction j does not depend on f_j^1 , it does depend on f_l^1 for $l > j$
 - If an auction l receives a set of (observed) bids b_l before auction j has ended, we can compute a posterior estimate of $f_l^1(v)$, and thus a better bid for auction j
 - sample from $f_l^1(v)$ by simulating auction l according to our auction model
- What about the **game theoretic approach**?
 - If each bidder (other than our agent) only participates in **one auction**:
 - dominant strategy is to bid truthfully: $b(v) = v$
 - we can use the decision-theoretic approach
 - If other bidders participate in **more than one auction** [Milgrom & Weber, 1982]
 - equilibrium strategy gets more complex (both strategically and computationally)
 - depends on entry, exit policies of other agents
 - If we have to estimate f and g , presumably other agents do too.
How should we account for the possibility that they will learn incorrect distributions?
 - success in this domain is much harder to benchmark experimentally
 - do we believe that all agents will follow an equilibrium strategy on eBay?

Talk Outline

1. Background

2. Online Auction Model and Learning Problem

3. Building an Agent

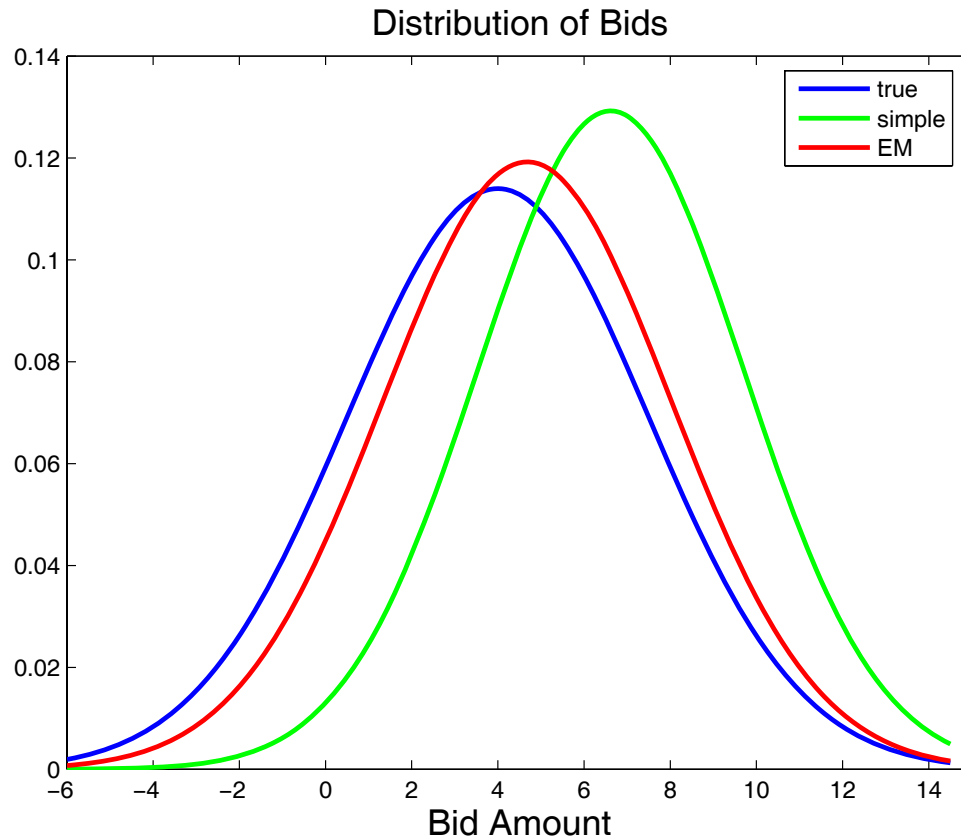
4. Experimental Evaluation

Experiments

- We compared our EM approach against the simple approach
 - I. Synthetic data: sequence of auctions for identical items, known distribution families
 - II. Synthetic data: sequence of auctions for non-identical items, known distribution families
 - III. Synthetic data: sequence of auctions for identical items, unknown distribution families
 - IV. eBay data: auctions for Playstation 2, March 2005.
- For each dataset, we ask two questions:
 1. Which approach gives better estimates of the distributions $f(v)$, $g(m)$, $f^1(v)$?
 2. Which approach gives better expected payoffs under the decision-theoretic bidding model?

Data Set I: Identical Items

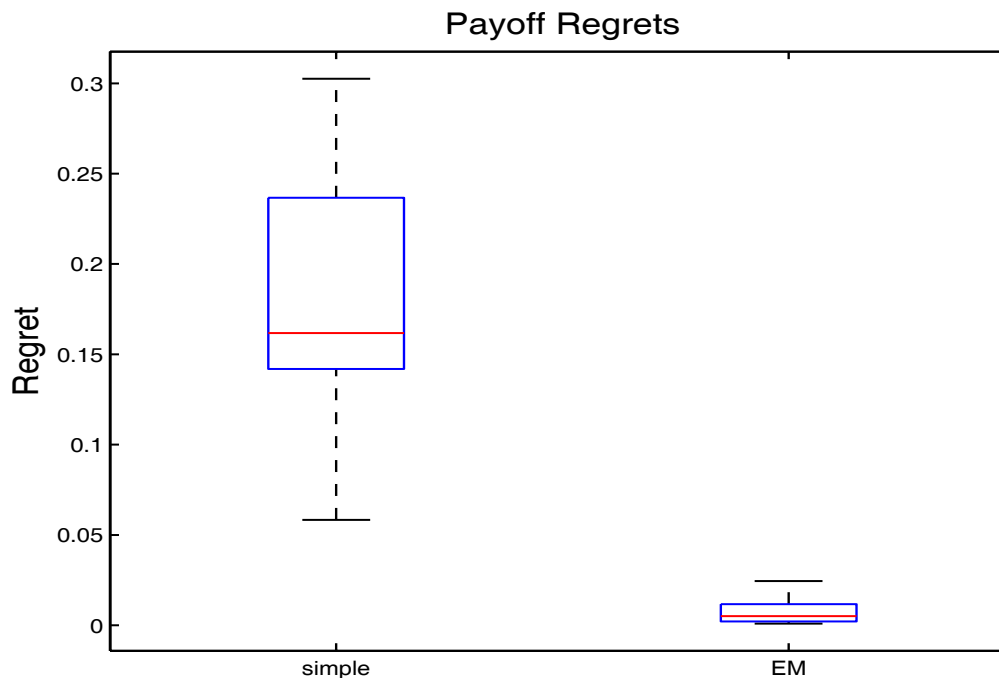
- Synthetic Data: $f(v)$ is a normal distribution; $g(m)$ is a Poisson distribution
- Bidding history of 40 auctions is generated for each instance.
- Both learning approaches use the correct (normal & Poisson) families of distributions to estimate $f(v)$ and $g(m)$
- Question 1: which approach made a **better estimate** of $f(v)$, $g(m)$, $f^1(v)$?



- EM approach consistently has **lower KL divergence** than the simple approach
- statistically significant difference: Wilcoxon sign-rank test (non-parametric)


Data Set I: Comparing Expected Payoffs

- Sequence of eight new auctions, after learning from the 40-auction history
 - in the new auctions, we still use the true $g(m)$ and $f(v)$
- Question 2: following the optimal strategy with the EM estimates gives **higher expected payoffs** than following this strategy with the simple approach's estimates



Data Set II: Non-identical Items


- The mean of $f(v)$ depends **linearly** on some unknown parameter a
- Both approaches use linear regression to estimate the linear coefficients
- Question 1: EM approach gives (stat. significantly) **better estimates**



- Question 2: EM approach achieves significantly **better expected payoffs**

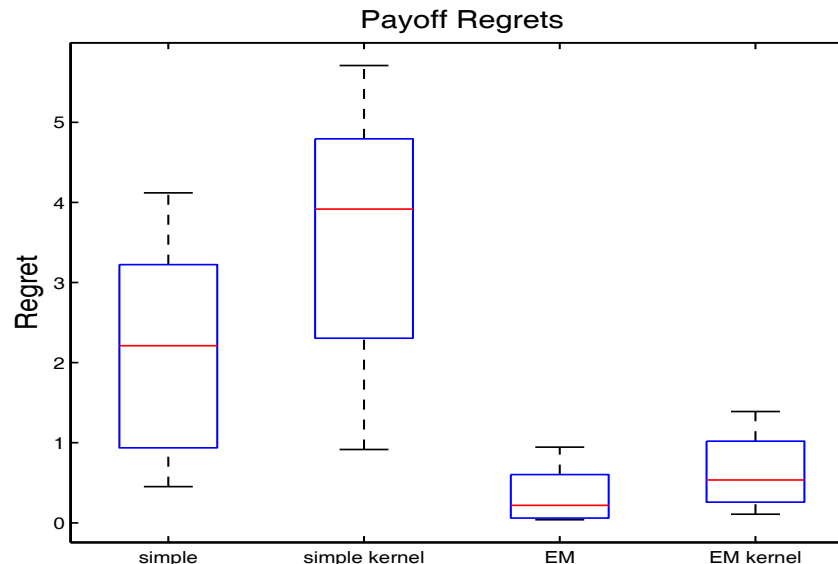
Data Set III: Unknown distributions

- Identical items. Distribution families for $f(v)$ and $g(m)$ are unknown
 - ground truth: $f(v)$ is Gamma distributed; $g(m)$ is a mixture of two Poissons
- Use kernel density estimation to estimate $f(v)$ and $g(m)$
- Result: the EM approach gives **better estimates** (significantly lower KL divergence); both approaches achieved **similar payoffs** (difference not significant)



Data Set IV: eBay Data

- 60 Sony **Playstation-2 auctions from eBay**, March 2005
 - considered only one-day auctions with at least 3 bidders
- Problem: highest bids not available
- Workaround: “pretend” second-highest bid is the highest bid
 - justification: this “shifted” data set should have similar characteristics to the actual bidding history
- Compared four approaches:
 - EM, simple approaches estimating normal and Poisson distributions
 - EM, simple approaches using kernel density estimation
- Question 1: **no ground truth** for this data set—dropped bids are *really* dropped, etc.
- Question 2: the EM approaches achieve **significantly higher expected payoffs** than the simple approaches.



Conclusion & Future Work

- **Bidding agents in online auction settings** face the problem of estimating
 - distribution of bid amounts;
 - distribution of number of biddersfrom incomplete auction data
- We proposed a **learning approach based on EM**
- We considered the application of **building a decision theoretic agent** for sequences of online auctions
- We showed in experiments that our EM approach **never did worse** and **usually did better** than the straightforward approach, on both synthetic and real-world data
- Thank you for your attention!