Algorithmic Game Theory

Auction Games and Games with Dynamic population

Eva Tardos, Cornell

Valparaiso Summer School

Recall: Smoothness for auctions

Auction game is (λ,μ)-smooth if for some μ>1, λ>0 and some strategy s* and all s we have

$$\sum_i u_i(s_i^*, s_{-i}) \geq \lambda opt - \mu\, R(s)$$

R(s) = revenue at bid vector s (usually μ=1)

Theorem: [Syrgkanis-T'13] Price of anarchy for any (λ,μ)-auction game is at most $\mu\,/\lambda$ for full information games.

- If s_i^* depends all of $(v_1, \ldots, v_n)$ then applied also to as well as Bayesian game with independent priors.
- If s_i^* depends only on v_i then applied also with dependent priors.

Social welfare: $\sum_i u_i(s) + R(s)$

Topic 1: Auction for many items

Two cases:

1. Additive value $v_i(A_i) = \sum_{j \in A_i} v_{ij}$
   Separate analysis for each buyer
2. Unit demand: $v_i(A_i) = \max_{j \in A_i} v_{ij}$ (free disposal)
   First price item action each: yesterday (1/2, 1) –smooth

Extension theorem: smooth auctions when combined remain smooth

Individual item auctions….Christodoulou, Kovacs, Schapire ICALP'08

Truthful auction: VCG

Analog of second price:

- Compute $\max \sum_j v_i(A_i^*)$ over all valid allocation $(A_1^*, \ldots, A_n^*)$
- Charge player i the externality he poses to others:

$$p_i = \max \sum_{j \neq i} v_j(A_j) - \sum_{j \neq i} v_j(A_j^*)$$

Optimizes social welfare and is truthful

Centralized, and computation?

Valuations:

- Complements are a problem:

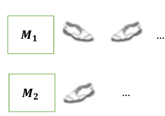

- Two bidders, with item set S of n items, k<<n
  - Bidders 1..k: $v_1(A) = \max_{j \in A} v_{ij}$ and v_{ij}=1 all j
  - Bidder k+1: $v_2(S) = n - k$ all other sets 0

No complements

Theorem. Valuation Hierarchy

Gross Substitutes ⊆ Decreasing marginal utility ⊆ XOS ≡ Fractionally Subadditive

Extending Lehmann et al'01 and Feige'06

XOS valuation: For any allocation $x_i = (x_i^1, \ldots, x_i^m)$

$$v_i(x_i) = \max_{k \in K} \sum_{j \in [m]} v_{ij}^k(x_i^j)$$

for some set of additively separable valuations K.

Example:

unit demand $v_i(x_i^1, \ldots, x_i^m) = \max_j v_i\left(x_i^j\right)$

No complements (cont)

- Decreasing marginal utility = Submodular:

  $v(A + j) - v(A) \geq v(B + j) - v(B)$ whenever $A \subset B$

Theorem. Submodular valuation can we written as XOS

Proof: For each ordering of the elements σ let v^{σ} be the marginal value in ordering σ

$v_j^{\sigma} = v(S_j^{\sigma}) - v(S_{-j}^{\sigma})$:

where S_j^{σ} is prefix till item j, and S_{-j}^{σ} prefix without j.

Note: max involves exponentially many

Simultaneous Composition

Theorem (Syrkganis-T'13): simultaneous item auctions where each is (λ,μ)-smooth and players have fractionally subadditive valuations, then composition is also (λ,μ)-smooth.

Note: with many items, learning is hard!

Simultaneous Composition

Theorem (Syrkganis-T'13): simultaneous item auctions where each is (λ,μ)-smooth and players have fractionally subadditive valuations, then composition is also (λ,μ)-smooth.

Proof: valuation $v(S)$ is fractionally subadditive ⇒ maximum of linear functions: $v(S) = \max_k \sum_{\{x \in S\}} v_j^k$

- optimal allocation $S_1^*, S_2^*, \ldots$ values $v_i(S_i^*) = \sum_{\{j \in S_i\}} v_j^{\{k_i^*\}}$
- Each auction smooth, so user i bids (not regrets) in auction j for value $v_j^{\{k_i^*\}}$
- Real value $v_i(A) \geq \sum_{\{j \in A\}} v_j^{\{k_i\}}$

Additional Topic: Dynamic Population

Main stability assumption in the literature:

- Game is repeated identically and nothing changes

Dynamic population model:

At each step t each player i

is replaced with an arbitrary new player with probability p

In a population of N players, each step, Np players replaced in expectation

Learning players can adapt….

Goal:

Bound average welfare assuming **adaptive** no-regret learners

$$PoA = \lim_{T \to \infty} \frac{\sum_{t=1} Cost(s^t; v^t)}{\sum_{t=1} Opt(v^t)}$$

where v^t is the vector of player types at time t

even when the rate of change is high, i.e. a large fraction can turn over at every step.

Need for adaptive learning

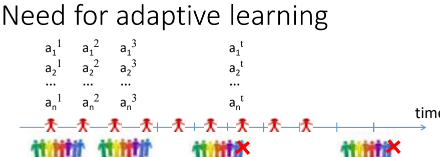

Example routing

- Strategy = path
- Best "fixed" strategy in hindsight very weak in changing environment
- Learners can adapt to the changing environment

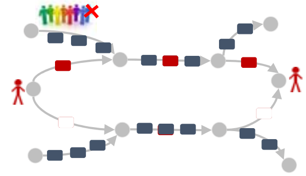

Need for adaptive learning

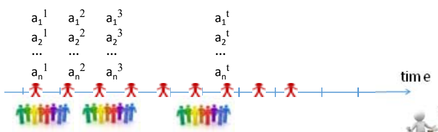

Example 2: matching (unit demand)

- Strategy = choose a mode
- Best "fixed" strategy in hindsight very weak in changing environment
- Learners can adapt to the changing environment

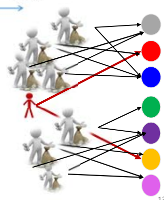

Adaptive Learning

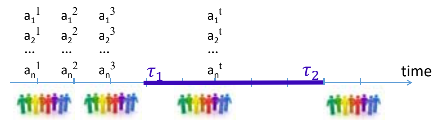

- Adaptive regret [Hazan-Seshadiri, Luo-Schapire]

  for all player i, strategy x and interval $[\tau_1, \tau_2]$

$$R_i(x, \tau_1, \tau_2) = \sum_{t=\tau_1}^{\tau_2} cost_i(s^t; v^t) - cost_i(x, s_{-i}^t; v^t) \le o(\tau_2 - \tau_1)$$

rates of $\sim\sqrt{\tau_2 - \tau_1}$

$\Rightarrow$ Regret with respect to a strategy that changes k times $\le \sim\sqrt{kT}$

Result (Lykouris, Syrgkanis, T'16) :

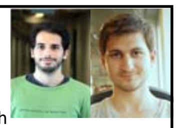

Bound average welfare close to Price of Anarchy for Nash

even when the rate of change is high, $p \approx \frac{1}{\log n}$ with n players

assuming **adaptive** no-regret learners

- Worst case change of player type $\Rightarrow$ need for adapting to changing environment
- Sudden large change is unlikely

No-regret and Price of Anarchy

Low regret:

$$R_i(x) = \sum_{t=1}^{T} cost_i(a^t; v^t) - cost_i(x, a_{-i}^t; v^t) \le o(T)$$

Best action varies with choices of others…

Consider Optimal Solution

Let $x = s_i^*$ be the choice in OPT

No regret for all players i:

$\sum_t cost_i(a^t) \le \sum_t cost_i(s_i^*, a_{-i})$

Players don't have to know s_i^*

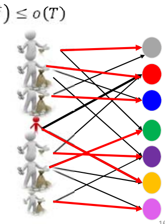

Adapting smoothness to dynamic populations

Inequality we "wish to have"

$$\sum_t cost_i(s^t; v^t) \le \sum_t cost_i(s_i^{*t}, s_{-i}^t; v^t)$$

where s_i^{*t} is the optimum strategy for the players at time t.

with stable population = no regret for s_i^*

Too much to hope for in dynamic case:

- sequence s^{*t} of optimal solutions changes too much.
- No hope of learners not to regret this!

Change in Optimum Solution

True optimum is too sensitive

- Example using matching
- The optimum solution
- One person leaving
- Can change the solution for everyone

- Np changes each step $\rightarrow$ No time to learn!! (we have p>>1/N)

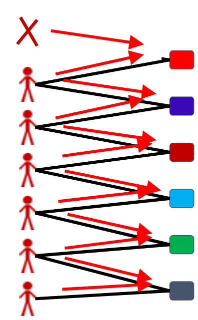

Theorem (high level)

If a game satisfies a "smoothness property" [Roughgarden'09]

The welfare optimization problem admits an approximation algorithm whose outcome $\tilde{s}^*$ is stable to changes in one player's type

Then any adaptive learning outcome is approximately efficient even when the rate of change is high.

Proof idea: use this approximate solution as $\tilde{s}^*$ in Price of Anarchy proof

With $\tilde{s}^*$ not changing much, learners have time to learn not to regret following $\tilde{s}^*$

Note: learner doesn't have to know $\tilde{s}^*$!!

19

Do Stable Solutions Exist?

- How close can we remain to the optimum, while being stable?
- How much change can we manage, while being stable?

Recall: Regret of adaptive learning is bounded by $\leq \sqrt{kT}$ with respect to any strategy that changes k times

Stable ≈ Optimum in Matching

True optimum is too sensitive

- Use greedy allocation: assign large values first (loss of factor of 2)
- Use coarse approximation of value, e.g., power of 2 only
- Potential function argument:
  increase in log value of allocation only m log v_{max}, decrease due to departures

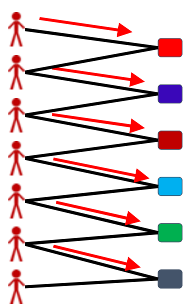

Use Differential Privacy → Stable Solutions

Joint privacy [Kearns et al. '14, Dwork et al. '06]

A randomized algorithm is jointly differentially private if

- when input from player **i** changes
- the probability of change in solution of players other than **i** is smaller than ϵ

- Turn a sequence of randomized solutions to a randomized sequence with small number of changes using Coupling Lemma
- and handling "failure probabilities" of private algorithms

22

Application 1: Large Congestion Games

- Using joint differentially private algorithm of Rogers et al EC'15,
- the (5/3,1/3)-smoothness congestion with affine cost:

Theorem. Atomic congestion game with m edges, and affine and increasing costs:

$$\frac{1}{T}\sum_t Cost(s^t; v^t) \leq 2.5(1+\epsilon)\frac{1}{T}\sum_t \mathrm{OPT}(v^t)$$

with $p = O\left(\frac{poly(\epsilon)}{poly(m)\ polylog(n)}\right)$ if each player controls only a 1/n fraction of the total flow.

Almost a constant fraction of change each step: dependence on number of players only polylog

23

Other Applications

Using joint differentially private algorithm of Hsu et al '14

Theorem 2. Matching markets if values are $[\rho,1]$

$\frac{1}{T}\sum_t W(s^t; v^t) \geq \frac{1}{4(1+\epsilon)}\frac{1}{T}\sum_t \mathrm{OPT}(v^t)$ with $p = O\left(\frac{\rho^2\epsilon^2}{polylog(m,1/\rho,1/\epsilon)}\right)$

Theorem 3. Large Combinatorial Markets with Gross-Substitutes

$\frac{1}{T}\sum_t W(s^t; v^t) \geq \frac{1}{2(1+\epsilon)}\frac{1}{T}\sum_t \mathrm{OPT}(v^t)$ with $p = O\left(\frac{\rho^5\epsilon^5}{m\ polylog(n)}\right)$

Each item in large supply $\Omega\left(polylog(n)\log(\frac{1}{\epsilon},\frac{1}{\rho})\right)$ and $\Theta(n)$ items

24

Conclusions/summary

- Analyzed quality of outcomes in games
  - Congestion games
  - Auction games

General technique: smoothness (Roughgarden'09, Syrgkanis-T'13)
  - Extension theorem to Bayesian games
  - Extension theorem to valuations to multiple items and valuations without complements
- Focus on learning outcomes in games:
  - Good way to adapt to opponents
  - No need for common prior
  - Takes advantage of opponent playing badly.
  - Learning players do well even in dynamic environments

25