The Pseudodimension of Near-Optimal Auctions

Jamie Morgenstern (Penn)
Tim Roughgarden (Stanford)

What Is…Simple?

What makes a mechanism "simple" or "complex"?

What Is...Simple?

What makes a mechanism "simple" or "complex"?

Simple vs. Optimal Theorem [Hartline/Roughgarden 09] (extending [Chawla/Hartline/Kleinberg 07]): in single-parameter settings, independent but not identical private valuations:

expected revenue of VCG with monopoly reserves $\geq$ ½ •(OPT expected revenue)

What Is…Simple?

[Babaioff/Immorlica/Lucier/Weinberg 14] for a single buyer, k items, additive and independent valuations:

better of selling the grand bundle or selling items separately $\geq$ constant •(OPT expected revenue)

- [Yao 15] extends to multiple buyers
- [Rubinstein/Weinberg 15] extends to subadditive valuations.

Quantifying Simplicity

Goal: quantitative definition of "mechanism simplicity."

Some example research directions:

- upper and lower bounds on best-possible performance guarantees of simple mechanisms
  - e.g., identify settings where only complex mechanisms can be approximately optimal
- automatic consequences of simplicity
  - formal justification for pursuit of simple mechanisms
  - e.g., to learning near-optimal auctions from data

Simplicity Has Many Forms

- will consider only direct-revelation DSIC mechanisms
  - randomized mechanisms OK
- not discussed: distinctions between DSIC, "obviously" DSIC [Li 15], deferred acceptance [Milgrom/Segal 15]
- not discussed: indirect mechanisms, e.g. with message space << type space
  - useful simplicity measure = number of actions/dimension of message space [Roughgarden 14]
- not discussed: polynomial communication/computation
  - not very relevant in our motivating examples

Related Work

- menu complexity [Hart/Nisan 13]
  - measures complexity of a single deterministic mechanism
  - maximum number of distinct options (allocations/prices) available to a player (ranging over others' bids)
  - selling items separately = maximum-possible menu complexity (exponential in the number of items)

Related Work

- menu complexity [Hart/Nisan 13]
  - measures complexity of a single deterministic mechanism
  - maximum number of distinct options (allocations/prices) available to a player (ranging over others' bids)
  - selling items separately = maximum-possible menu complexity (exponential in the number of items)
- mechanism design via machine learning [Balcan/Blum/Hartline/Mansour 08]
  - covering number measures complexity of a family of auctions
  - prior-free setting (benchmarks instead of unknown distributions)
  - near-optimal mechanisms for unlimited-supply settings

Pseudodimension: Examples

Proposed simplicity measure of a class C of mechanisms: *pseudodimension* of the real valued functions (from valuation profiles to revenue) induced by C.

Pseudodimension: Examples

Proposed simplicity measure of a class C of mechanisms: *pseudodimension* of the real valued functions (from valuation profiles to revenue) induced by C.

Examples:

- Vickrey auction, anonymous reserve — $O(1)$
- Vickrey auction, bidder-specific reserves — $O(n \log n)$
- grand bundling/selling items separately — $O(k \log k)$
- virtual welfare maximizers — unbounded

Pseudodimension: Implications

Theorem: [Haussler 92], [Anthony/Bartlett 99] if C has low pseudodimension, then it is easy to learn from data the best mechanism in C.

Pseudodimension: Implications

Theorem: [Haussler 92], [Anthony/Bartlett 99] if C has low pseudodimension, then it is easy to learn from data the best mechanism in C.

- "data" = samples from unknown valuation distribution F
  - Yahoo! example: [Ostrovsky/Schwarz 09]

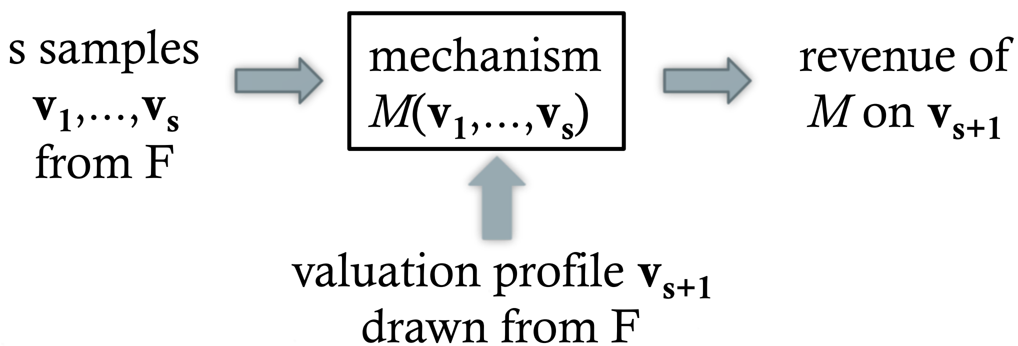

Pseudodimension: Implications

Theorem: [Haussler 92], [Anthony/Bartlett 99] if C has low pseudodimension, then it is easy to learn from data the best mechanism in C.

- "data" = samples from unknown valuation distribution F
  - Yahoo! example: [Ostrovsky/Schwarz 09]
- theoretical work: [Elkind 07], [Dhangwatnotai/Roughgarden/Yan 10], [Cole/Roughgarden 14], [Chawla/Hartline/Nekipelov 14], [Medina/Mohri 14], [Cesa-Bianchi/Gentile/Mansour 15], [Dughmi/Han/Nisan 15], [Huang/Mansour/Roughgarden 15], [Devanur/Huang/Psomas 15], ...

Pseudodimension: Implications

Theorem: [Haussler 92], [Anthony/Bartlett 99] if C has low pseudodimension, then it is easy to learn from data the best mechanism in C.

- obtain $s = \tilde{\Omega}(H^2 \varepsilon^{-2} d)$ samples $\mathbf{v}_1,\ldots,\mathbf{v}_s$ from F, where d = pseudodimension of C, valuations in [0,H]
- let M^* = mechanism of C with maximum total revenue on the samples

Guarantee: with high probability, expected revenue of M^* (w.r.t. F) within ε of optimal mechanism in C.

Pseudodimension: Definition

[Pollard 84]

Let F = set of real-valued functions on domain X.

(for us, X = valuation profiles, F = mechanisms, range = revenue)

F *shatters* a finite subset $S=\{\mathbf{v_1},\ldots,\mathbf{v_s}\}$ of X if:

- there exist real-valued thresholds $t_1,\ldots,t_s$ such that:
- for every subset T of S
- there exists a function f in F such that:

$$f(\mathbf{v_i}) \geq t_i \iff \mathbf{v_i} \text{ in } T$$

Pseudodimension: Example

Let C = second-price single-item auctions with bidder-specific reserves.

Claim: C can only shatter a subset $S=\{\mathbf{v}_1,\ldots,\mathbf{v}_s\}$ if $s = O(n \log n)$. (hence pseudodimension $O(n \log n)$)

Pseudodimension: Example

Let C = second-price single-item auctions with bidder-specific reserves.

Claim: C can only shatter a subset $S=\{\mathbf{v}_1,\ldots,\mathbf{v}_s\}$ if $s = O(n \log n)$. (hence pseudodimension $O(n \log n)$)

Proof sketch: Fix S.

- Bucket auctions of C according to relative ordering of the n reserve prices with the ns numbers in S. (#buckets $\approx (ns)^n$)
- Within a bucket, allocation is constant, revenue varies in simple way => at most s^n distinct "labellings" of S.
- Since need 2^s labellings to shatter S, $s = O(n \log n)$.

Consequences

Meta-theorem: simple vs. optimal results automatically extend from known distributions to unknown distributions with a polynomial number of samples.

Examples:

- Vickrey auction, anonymous reserve — O(1)
- Vickrey auction, bidder-specific reserves — O(n log n)
- grand bundling/selling items separately — O(k log k)

Guarantee: with $s = \tilde{\Omega}(H^2 \varepsilon^{-2} d)$, with high probability, expected revenue of M^* (w.r.t. F) within ε of optimal mechanism in C.

Simplicity-Optimality Trade-Offs

Simple vs. Optimal Theorem: in single-parameter settings, independent but not identical private valuations:

expected revenue of VCG with monopoly reserves $\geq$ ½ •(OPT expected revenue)

Simplicity-Optimality Trade-Offs

Simple vs. Optimal Theorem: in single-parameter settings, independent but not identical private valuations:

expected revenue of VCG with monopoly reserves $\geq$ ½ •(OPT expected revenue)

t-Level Auctions: can use t reserves per bidder.

- winner = bidder clearing max # of reserves, tiebreak by value

Simplicity-Optimality Trade-Offs

Simple vs. Optimal Theorem: in single-parameter settings, independent but not identical private valuations:

expected revenue of VCG with monopoly reserves $\geq$ ½ •(OPT expected revenue)

t-Level Auctions: can use t reserves per bidder.

- winner = bidder clearing max # of reserves, tiebreak by value

Theorem: (i) pseudodimension = $O(nt \log nt)$;

(ii) to get a $(1-\varepsilon)$-approximation, enough to take $t \approx H/\varepsilon$ [for matroids] and $t \approx Hn^2/\varepsilon$ [in general]

Summary

- pseudodimension = classical definition from statistical learning theory, appealing way to quantify the "simplicity" of a family of mechanisms
- analytically tractable to upper bound in many cases
- simple vs. optimal results extend from known distributions to unknown distributions with a polynomial number of samples

Wide open: incorporate computational complexity issues (cf., computational learning theory).