When Algorithmic Game Theory met Machine Learning

Christos H. Papadimitriou
Simons Institute, Berkeley

Outline

- I Learning and equilibria
- II How to classify points that will respond strategically to your classifier (joint work with Hardt, Megiddo, Wootters)
- III How to optimize statistical learning through mechanism design – by incentivizing your data providers (joint work with Cai and Daskalakis)

Not covered

- IV Learning the priors in auctions, or the clickthrough rates
- V Reducing mechanism design to algorithm design via learning
- VI…
- VII…
- …

ca 1950: Fictitious play for zero-sum games

- G. W. Brown 1949: In repeated play, suppose that both players interpret the empirical probabilities of the other's plays as a mixed strategy, and best respond to it
- J. Robinson 1951: It converges to min-max
- S. Karlin 1959: Conjectured quadratic time
- Daskalakis and Pan 2014: Exponential counterexample

Non-zero sum games?

- Shapley [1964]: Fictitious play does not converge to Nash even in 3 × 3 games
- [DFPPV 2010]: Even multiplicative weight dynamics does not converge to Nash, even in 3 × 3 games
- (But remember: [Hart – Mas Colel, others]: No-regret dynamics do approach correlated equilibria…)

Two other interpretations

1. This dynamics is problematic, too weak to achieve the Nash equilibrium

1. This solution concept (Nash equilibrium) must be wrong, because the natural dynamics never gets there

A propos the second interpretaion

- For which games does the multiplicative weights dynamics converge to a Nash equilibrium?
  - Two-player zero sum games
  - Potential games
  - Coordination games
  - Zero-sum multimatrix games [CDOP 2014]

Complexity interpretation

- We know that Nash is intractable. So, these are lower bounds *without complexity assumptions such as PPAD ≠ P,* showing that certain simple classes of algorithms fail (we have done this for NP-complete problems in the 1970s and 1980s)
- btw: The important theorem here remains to be proved

Interesting fact:

- Every 2-player game is the sum of a zero-sum game and a potential game:

$$(A, B) = \tfrac{1}{2}(A - B, B - A) + \tfrac{1}{2}(A + B, A + B)$$

- OK, interesting, but what does it mean?

A more meaningful decomposition

- The mw dynamics eventually converge to cyclical behavior (= equilibrium)
- Therefore, the mixed strategy space of all games can be decomposed into

Transient potential game + Equilibrium behavior

- ***All games are potential games!***
- Requires serious dynamical systems machinery (ongoing work with G. Piliouras)

Changing the subject:
II Classification
when the data are strategic

- (Joint work with Hardt, Megiddo, Wootters)
- Fact: The number of books in the household is an excellent predictor of success at school – actually, a bit better than grades
- So, an admissions classifier should use this

Problem:

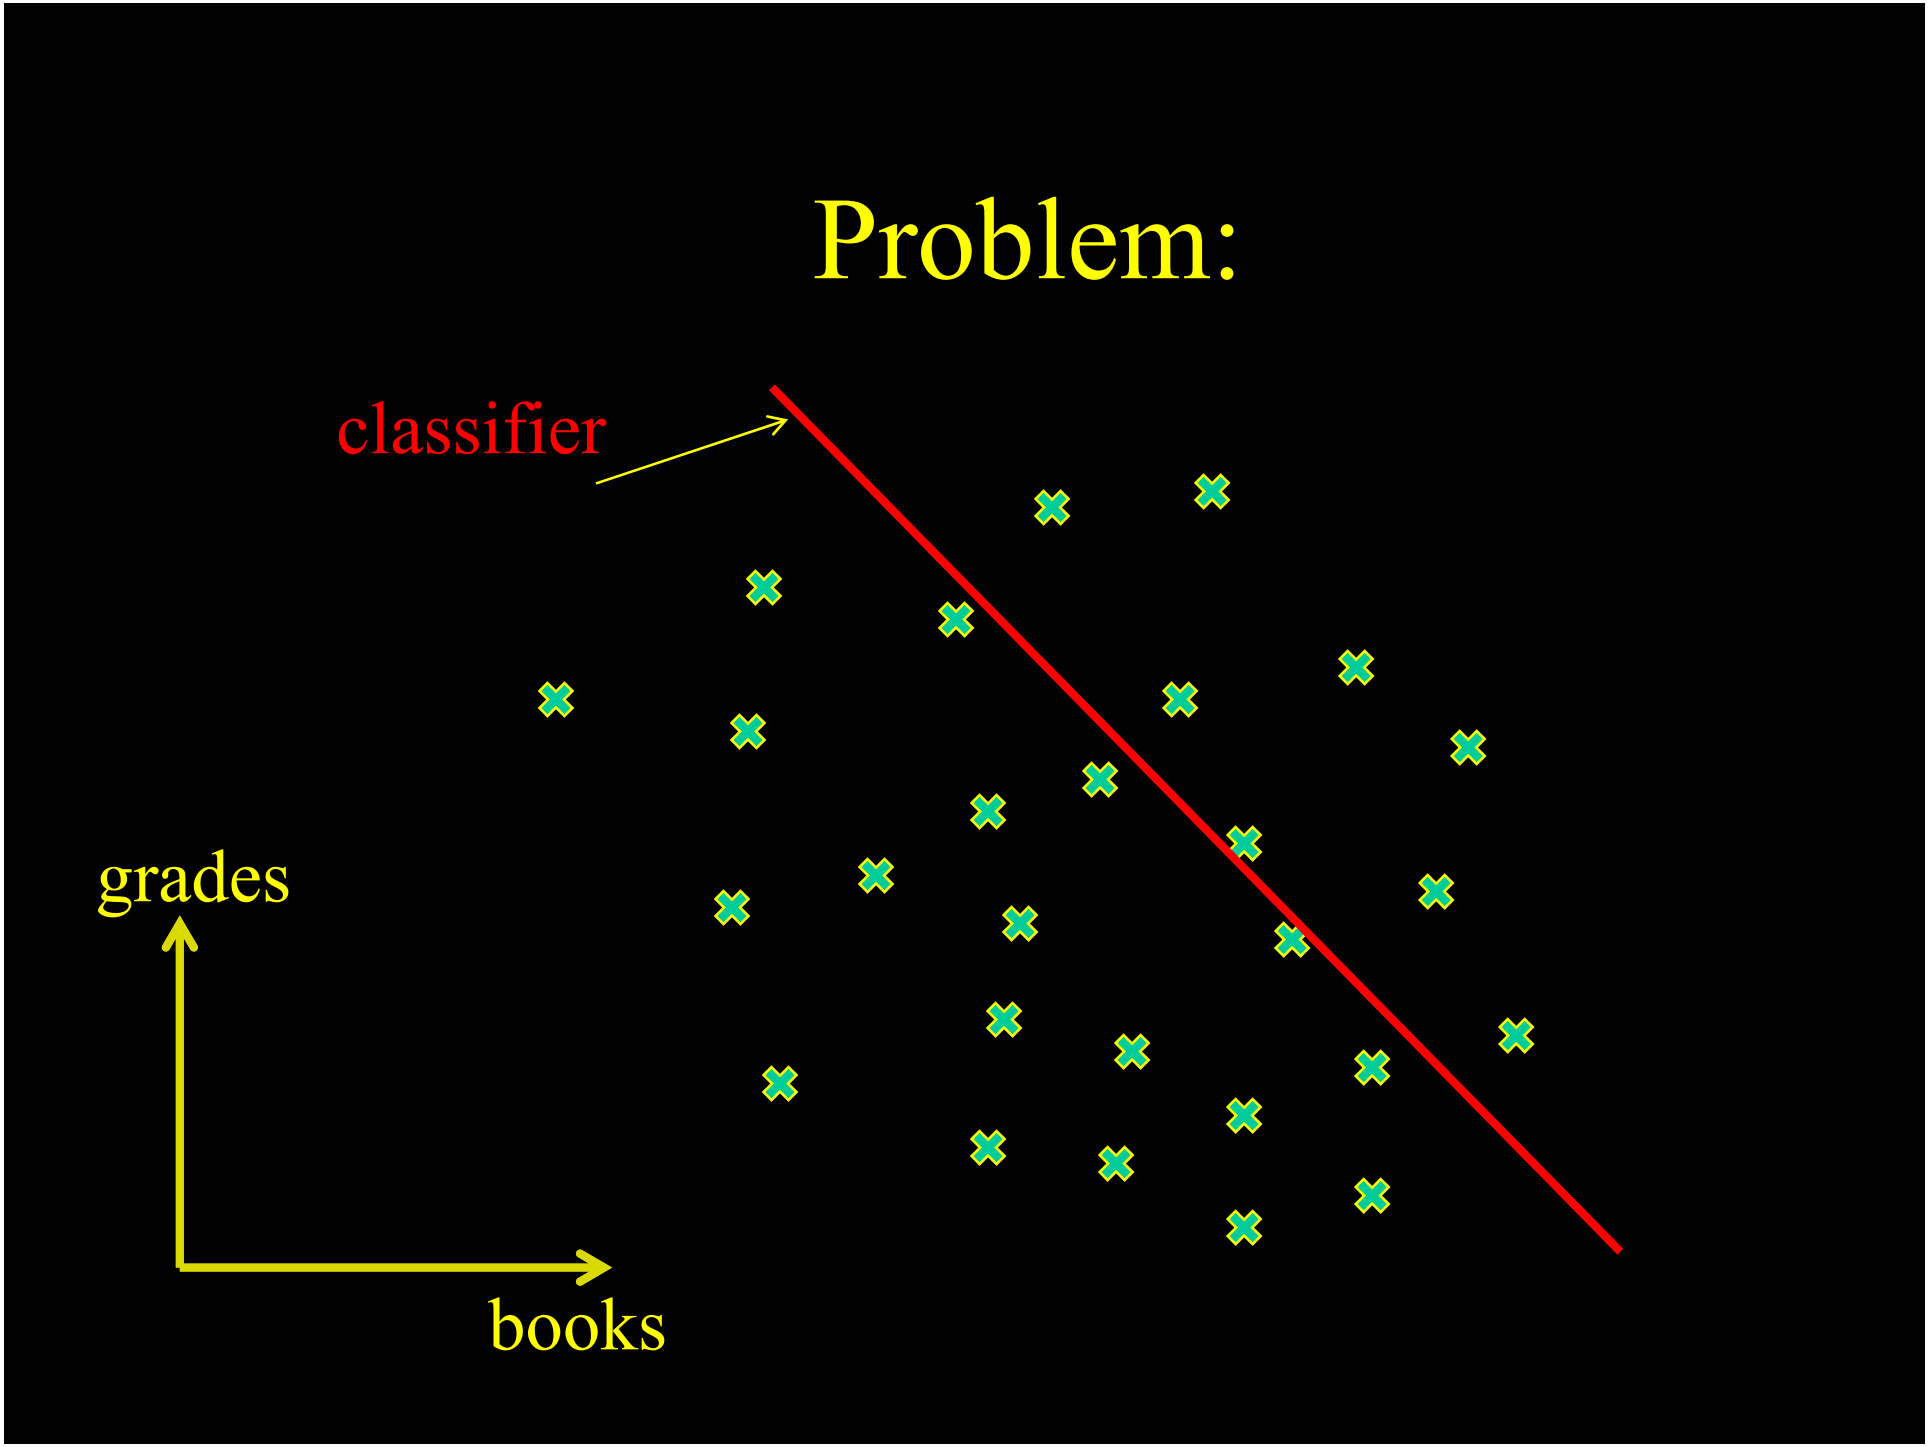

The data are strategic!

- Applicants will buy enough books to get in
- (As long as their utility of admission covers this expense)

So, you have to…

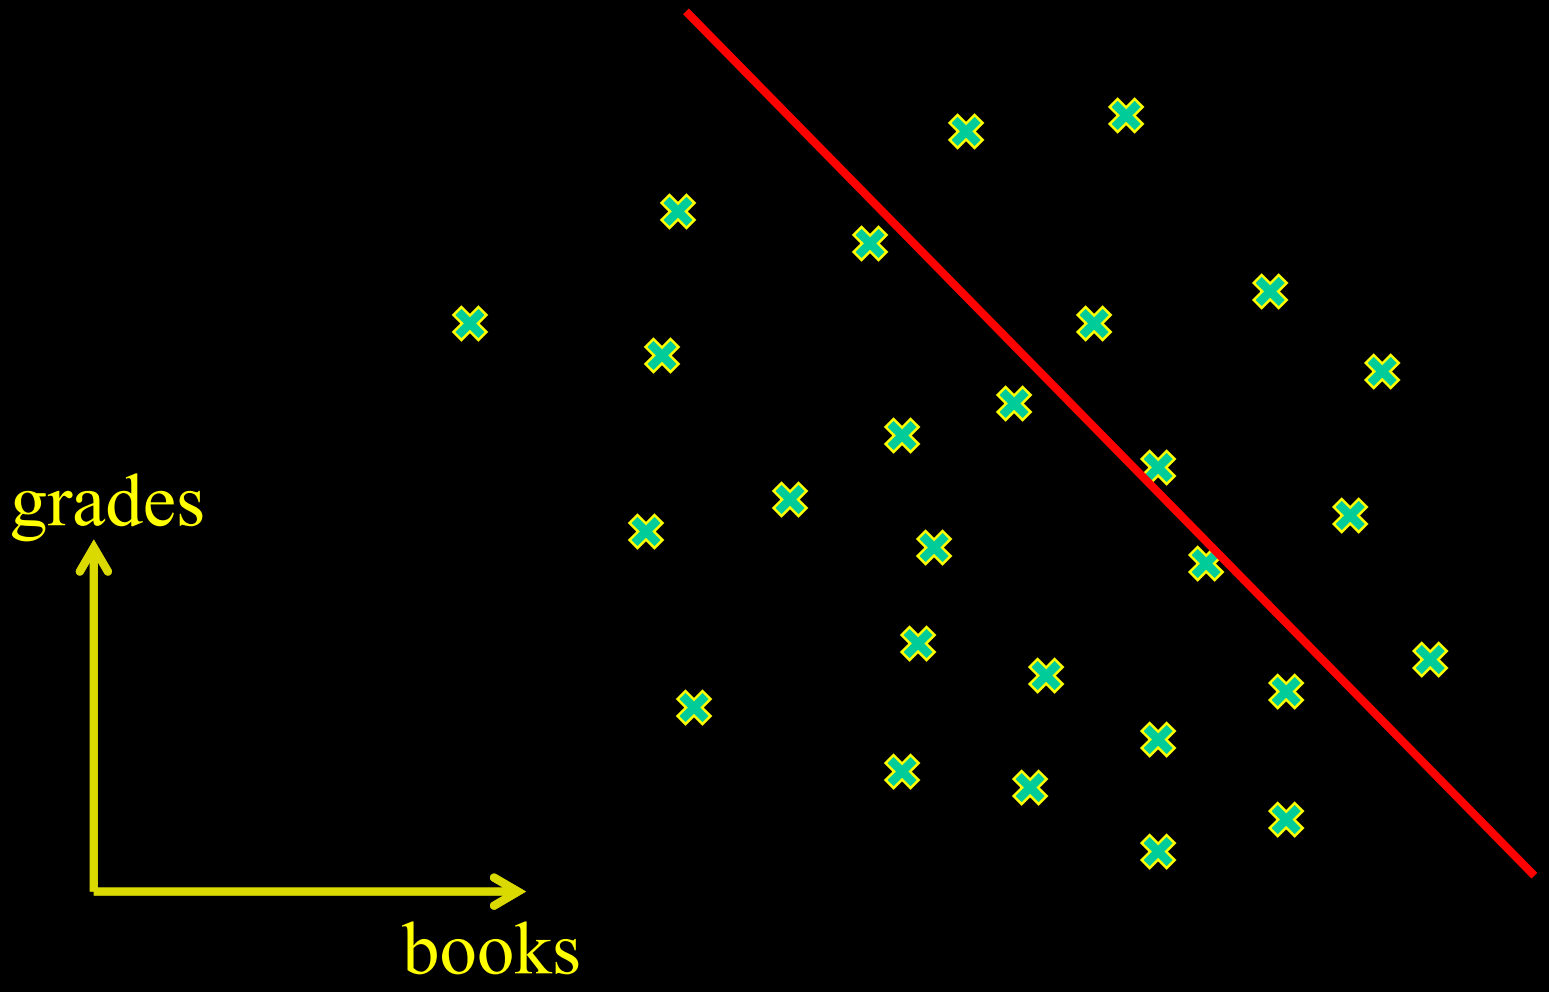

But of course…

- …this way you may lose some good students who are poor…
- Suppose you want to maximize the ***score***
  = #correctly classified - # misclassified
- Applicants: unit utility for being admitted
- Each has a cost $c_i(x, y)$ for moving from x to y (not necessarily a metric)

Stackelberg!

- Given: Data distribution, cost function, ideal classifier H…
- …come up with modified classifier H^*…
- …such that, when the data respond by moving to the closest admitted point (if cost of moving < 1)…
- …the expected score is maximized

Good news, bad news

Theorem 1: Can be done arbitrarily close to the optimum in polynomial sample complexity ***and*** time as long as the original classification problem has low sample complexity and the cost function has "finite s-dimension"

Theorem 2: NP-hard to approximate within any ratio even if you know the data points

Finally:
III Strategic Data *Sources*

- (Joint work with Cai and Daskalakis, in Arxiv)
- Context: Statistical estimation
- There is a ground truth function F(x)
- You estimate it from data points (x_i,y_i)
- You get the data points from workers
- Each worker i has a convex function σ_i
- With effort e >0 from x she comes up with y = F(x) in expectation, and variance $\sigma_i^2(e)$

The agents

- You (the Statistician) have a way, from data $Z=\{x_i,y_i\}$, to find an approximation G_Z of F
- (Think linear regression, much more general)
- Your loss is $L = E_{x\sim D}\,[(F(x) - G_Z(x))^2]$
- Assumption: L depends only on $\{x_i\}$, D, $\{\sigma_i\}$
- (Not on F; estimation is unbiased)
- Plus any payments to workers, $\Sigma_i\, p_i$
- Each worker i wants to minimize $e_i - p_i$

Social Optimum

$$OPT = \min_{w', x, e} \; [\, L(\{x_i\}, D, \{e_i\}) + \Sigma_i \, e_i \,]$$

- Notice that OPT is your (the Statistician's) wildest dream
- It means that you have persuaded the workers to exert optimum effort at no surplus (can't do better because of IR)

Surprise: It can be achieved!

Theorem: There is a mechanism which achieves "your wildest dream" utility OPT as dominant strategy equilibrium

- Trick: Promise payments

$$p_i = A_i - B_i\,(y_i - G_{-i}\,(x_i))^2$$

zero surplus

optimum effort

Caveat

- This means that the incentive problem can be solved, essentially by a surprisingsly powerful contract
- What is less clear is how to solve the computational problem – how to compute OPT
- Btw, if the x_i's are given, even ridge regression can be managed the same way

Soooo….

- There seems to be a secret affinity between AGT and ML
- Key complementary ways of dealing with the information environment
- Their interface is ancient and growing fast


谢谢