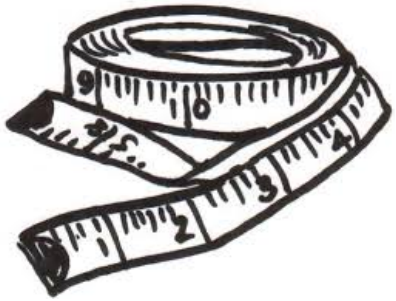

Weights and Measures: Fast and Active Prediction in the Era of Big Data

Lev Reyzin
UIC, Math Department

In a world of Big Data…

- Google/Yahoo! approaching 1B daily visitors.
- 3B yearly doctor's visits in the US (records digitized).
- CERN currently storing > 100 petabytes of data.
- The NSA collects…. everything.

In a world of Big Data…

- Google/Yahoo! approaching 1B daily visitors.
  What ads to show?
- 3B yearly doctor's visits in the US (records digitized).
  How to treat patients?
- CERN currently storing > 100 petabytes of data.
  Which of the events contain an interesting signal?
- The NSA collects…. everything.
  Who is a terrorist?

Some Opportunities and Challenges

Can handle rich / interesting classes of functions.

Should be really fast (linear time, faster?)

Have petabytes of data

Can label only some of it.

Can finally predict accurately.

Must predict quickly! (Or otherwise limited.)

Weights and Measures

- weighing – keep track of "performance" without needing to remember the history.

- measuring – given the weights, how to properly "measure" them to determine the correct outcome?

Optimal and Efficient Contextual Bandits

"the world is a bandit problem"

with

Beygelzimer, Dudik, Hsu, Kale, Karampatziakis, Langford, Li, Schapire, Zhang

What is the Bandit Setting?

- The name bandit refers to slot machines in a casino.
- You choose actions (e.g. which machine to play), one after another. These actions come with rewards.
- Goal is to minimize your regret – informally, how well you did compared to how well you could have done.
    - Bound expected regret or bound regret w.h.p.

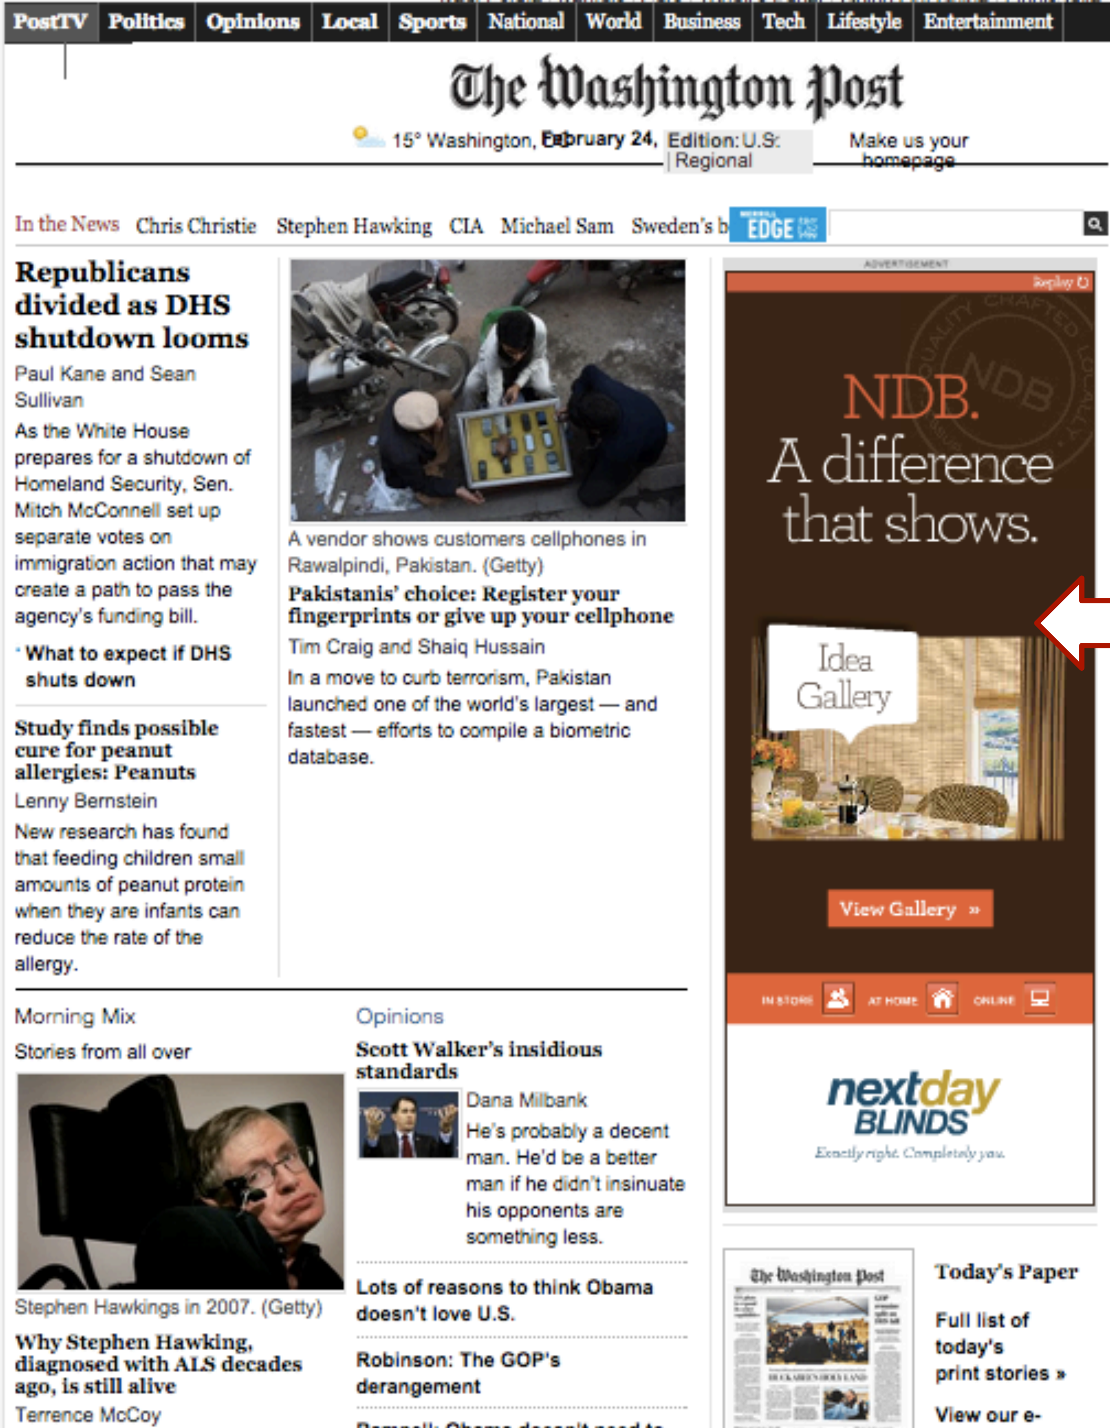

PostTV | Politics | Opinions | Local | Sports | National | World | Business | Tech | Lifestyle | Entertainment

The Washington Post

15° Washington, February 24, Edition: U.S. | Regional Make us your homepage

In the News Chris Christie Stephen Hawking CIA Michael Sam Sweden's b

Republicans divided as DHS shutdown looms

Paul Kane and Sean Sullivan

As the White House prepares for a shutdown of Homeland Security, Sen. Mitch McConnell set up separate votes on immigration action that may create a path to pass the agency's funding bill.

- What to expect if DHS shuts down

Study finds possible cure for peanut allergies: Peanuts

Lenny Bernstein

New research has found that feeding children small amounts of peanut protein when they are infants can reduce the rate of the allergy.

A vendor shows customers cellphones in Rawalpindi, Pakistan. (Getty)

Pakistanis' choice: Register your fingerprints or give up your cellphone

Tim Craig and Shaiq Hussain

In a move to curb terrorism, Pakistan launched one of the world's largest — and fastest — efforts to compile a biometric database.

Morning Mix

Stories from all over

Stephen Hawkings in 2007. (Getty)

Why Stephen Hawking, diagnosed with ALS decades ago, is still alive

Terrence McCoy

Opinions

Scott Walker's insidious standards

Dana Milbank

He's probably a decent man. He'd be a better man if he didn't insinuate his opponents are something less.

Lots of reasons to think Obama doesn't love U.S.

Robinson: The GOP's derangement

Rampell: Obama doesn't need to

Today's Paper

Full list of today's print stories »

View our e-

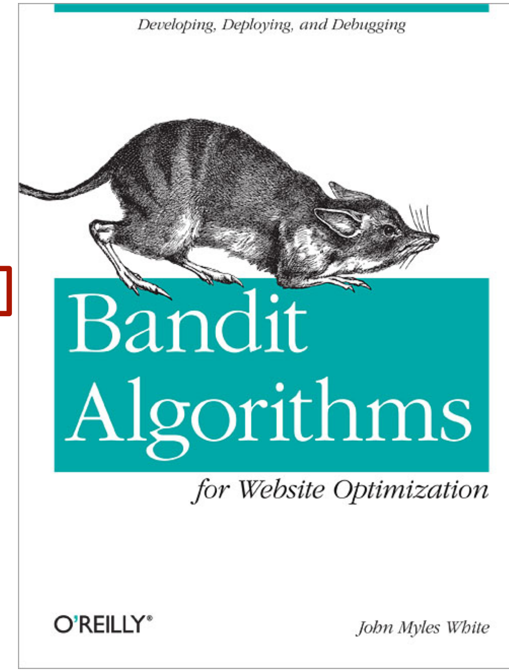

9/21/2010

Regretting the dead

Tags: Exploration,Machine Learning,Online — jl@ 9:31 pm

Nikos pointed out this new york times article about poor clinical design killing people. For those of us who study learning from exploration information this is a reminder that low regret algorithms are particularly important, as regret in clinical trials is measured by patient deaths.

The New York Times

Research

TARGET CANCER

New Drugs Stir Debate on Rules of Clinical Trials

Monica Almeida/The New York Times, left

Two Cousins, Two Paths Thomas McLaughlin, left, was given a promising experimental drug to treat his lethal skin cancer in a medical trial; Brandon Ryan had to go without it.

Contextual Bandits [Auer et al. '02]

context:

1 

2 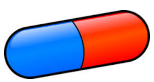

3 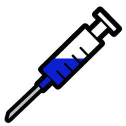

⋮

k 

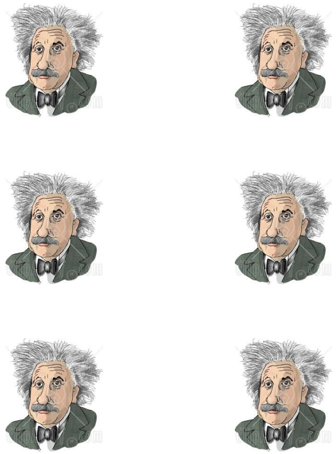

N experts/function
think of N >> K

Contextual Bandits [Auer et al. ’02]

context: x_1

1 

2 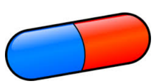

3 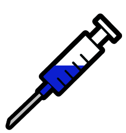

⋮

k 

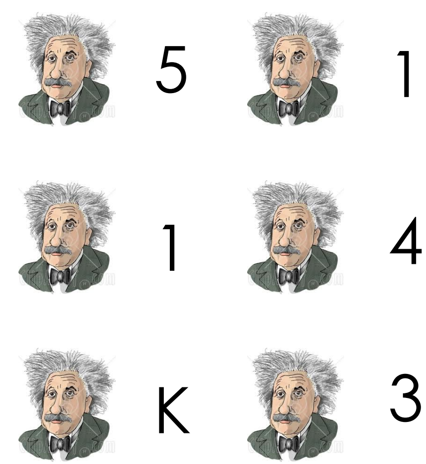

N experts/function
think of N >> K

Contextual Bandits [Auer et al. '02]

context: x_1

1  yes

2 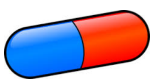

3 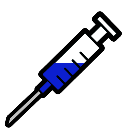

⋮

k 

 5  1

 1  4

 K  3

N experts/function
think of N >> K

Contextual Bandits [Auer et al. '02]

context: x_1 x_2


1  yes

2 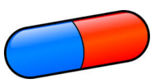

3  yes

⋮

k 

 2  1

 K  3

 2  1

N experts/function
think of N >> K

Contextual Bandits [Auer et al. '02]

context: x_1 x_2 x_3 x_4 … x_T

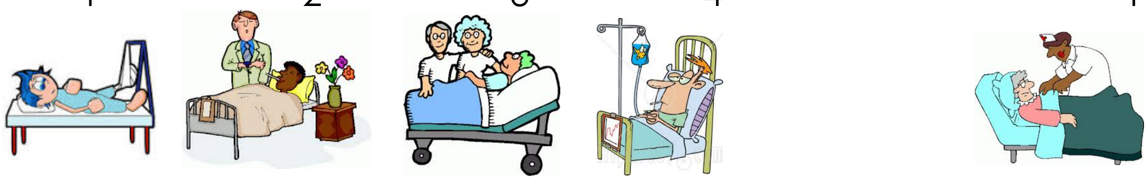

		x_1	x_2	x_3	x_4	…	x_T
1		yes					yes
2	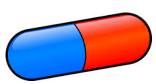						
3	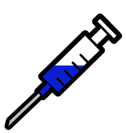		yes		no		
⋮							
k				no			

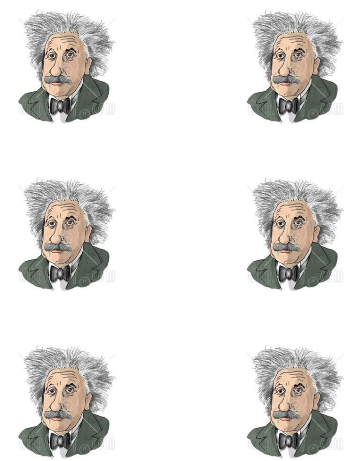

N experts/function
think of N >> K

Contextual Bandits [Auer et al. '02]

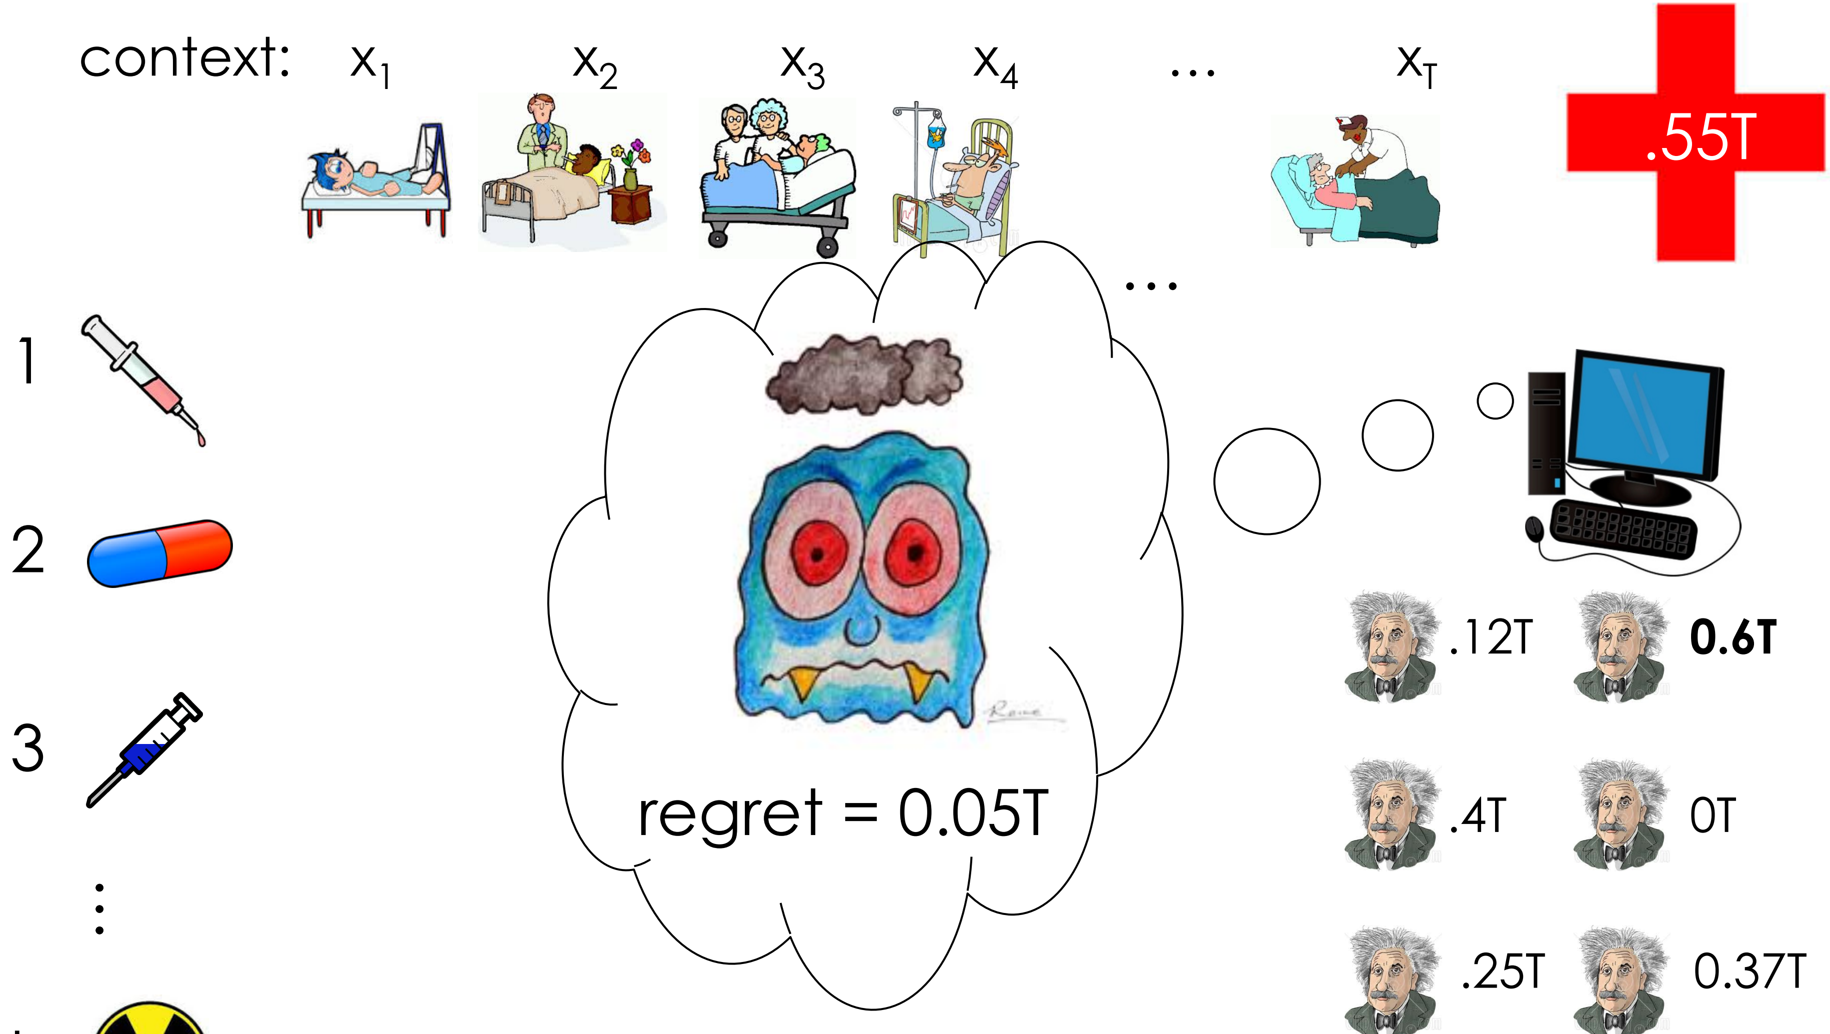

N experts/function
think of N >> K

Contextual Bandits [Auer et al. '02]

context: x_1 x_2 x_3 x_4 … x_T





…

1 

2 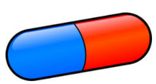

the contexts & rewards can come from a distribution (stochastic) or be adversarial

3 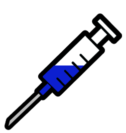

The experts can be present (contextual) or not.

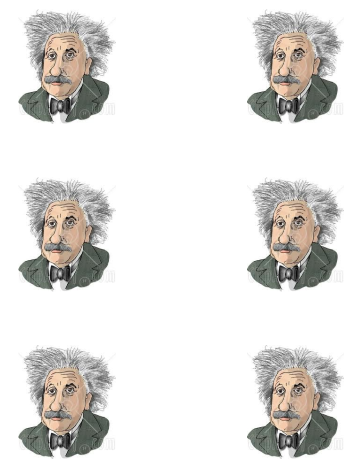

⋮

k 

- harder than supervised (usual) learning:

  In the bandit setting, we do not know the rewards of actions not taken.

- many applications:

  Medicine, ad auctions, finance, …

- exploration/exploitation dilemma:

  exploit policies you've learned to be good?

  explore policies you're not sure about?

- $\Omega(TK \ln N)^{1/2}$ is a known lower bound on regret [Auer et al. '02] even in the stochastic setting.
- Algorithms achieving $\tilde{O}(KT \text{ polylog } N)^{1/2}$ regret are said to be optimal.
- greedily first exploring (acting randomly) then exploiting (following best policy) cannot be optimal. Optimal algorithms must be adaptive.

Upper Confidence Bounds

[Auer '02]

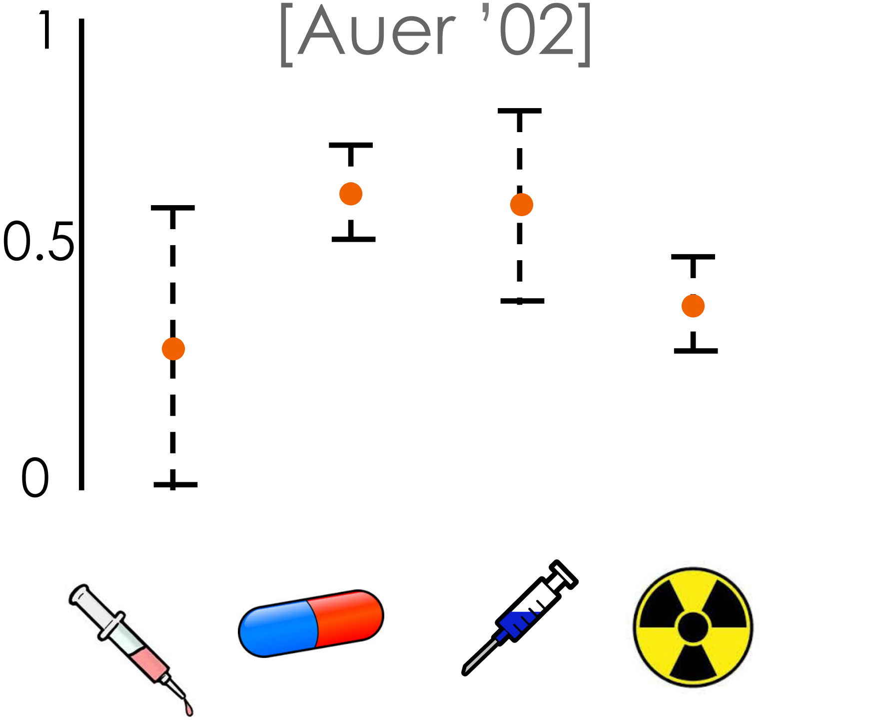

At every time step:

1) choose action with highest UCB
2) update confidence bound of the arm pulled.

EXP3 Exponential Weights

[Auer et al '02]

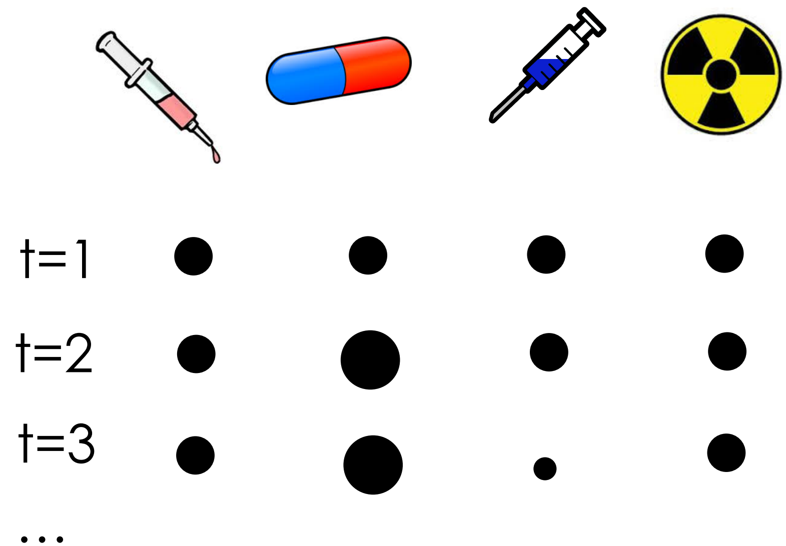

At every time step:

1) sample from distribution defined by weights (mixed w/ uniform)
2) "exponential" weight updates

UCB vs EXP3: A Comparison

UCB

[Auer '02]

◆ Pros

Optimal in stochastic setting.

Succeeds w.h.p.

◆ Cons

Fails in adversarial setting.

Not optimal in the contextual setting.

EXP3 & Friends

[Auer et al. '02]

◆ Pros

Optimal for adversarial and stochastic settings.

Adaptable to the contextual setting

◆ Cons

Succeeds only in expectation.

Algorithm	Optimal?	High Prob?	Context?
Exp4 [ACFS '02]	Yes	No	Yes
epoch-geedy [LZ '07]	No	Yes	Yes
Exp3.P [ACFS '02] UCB [Auer '00]	Yes	Yes	No

Algorithm	Optimal?	High Prob?	Context?
Exp4 [ACFS '02]	Yes	No	Yes
epoch-geedy [LZ '07]	No	Yes	Yes
Exp3.P [ACFS '02] UCB [Auer '00]	Yes	Yes	No
Exp4.P [BLLRS '11]	Yes	Yes	Yes

EXP4.P

[Beygelzimer-Langford-Li-R-Schapire '11]

Main Theorem [BLL**R**S '11]: For any $\delta > 0$, w.p. at least $1-\delta$, EXP4P has regret at most $O(KT \ln (N/\delta))^{1/2}$ in adversarial contextual bandit setting.

Combines advantages of Exponential Weights and UCB.

- Optimal for both the stochastic and adversarial settings
- Works for contextual case (also non-contextual case)
- A high probability result

Ideas Behind Exp4.P

(all appeared in previous algorithms)

- exponential weights

  keep a weight on each expert that drops exponentially in the expert's (estimated) performance

- upper confidence bounds

  adds an upper confidence bound on each expert's estimated reward

- ensuring exploration

  make sure each action is taken with some minimum probability

- importance sampling

  give rare events more importance to keep estimates unbiased

Key Insights

key insights
(on top of UCB/ EXP)

1) exponential weights and confidence bounds "stack"
2) generalized Bernstein's inequality for martingales

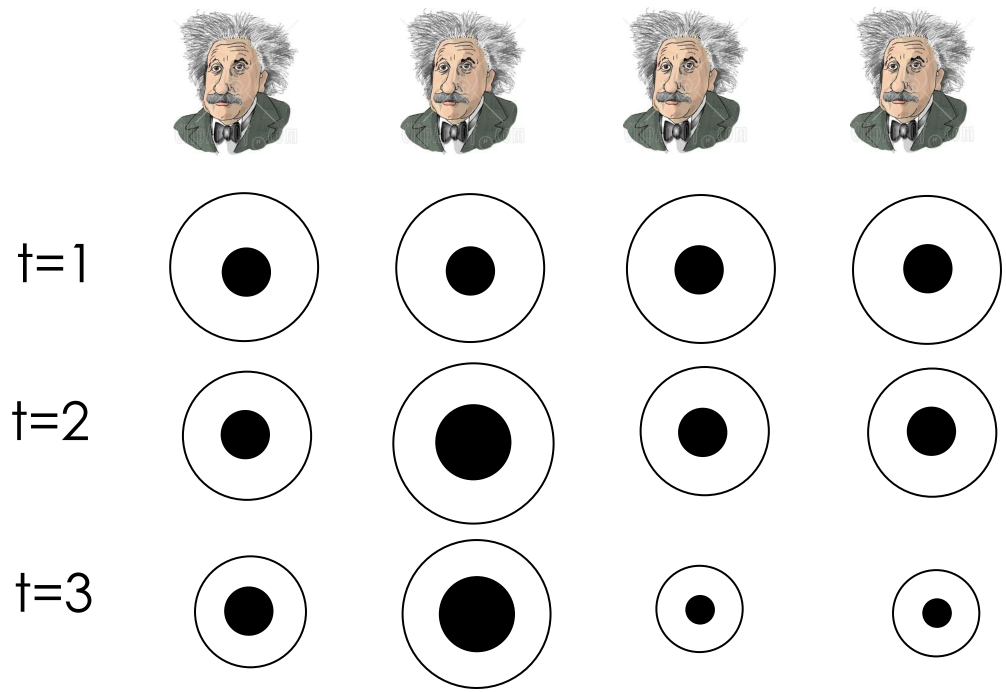

Algorithm	Optimal?	High Prob?	Context?	Efficient?
Exp4 [ACFS '02]	Yes	No	Yes	No
epoch-geedy [LZ '07]	No	Yes	Yes	Yes
Exp3.P [ACFS '02] UCB [A '00]	Yes	Yes	No	Yes
Exp4.P [BLL**R**S '10]	Yes	Yes	Yes	**No**

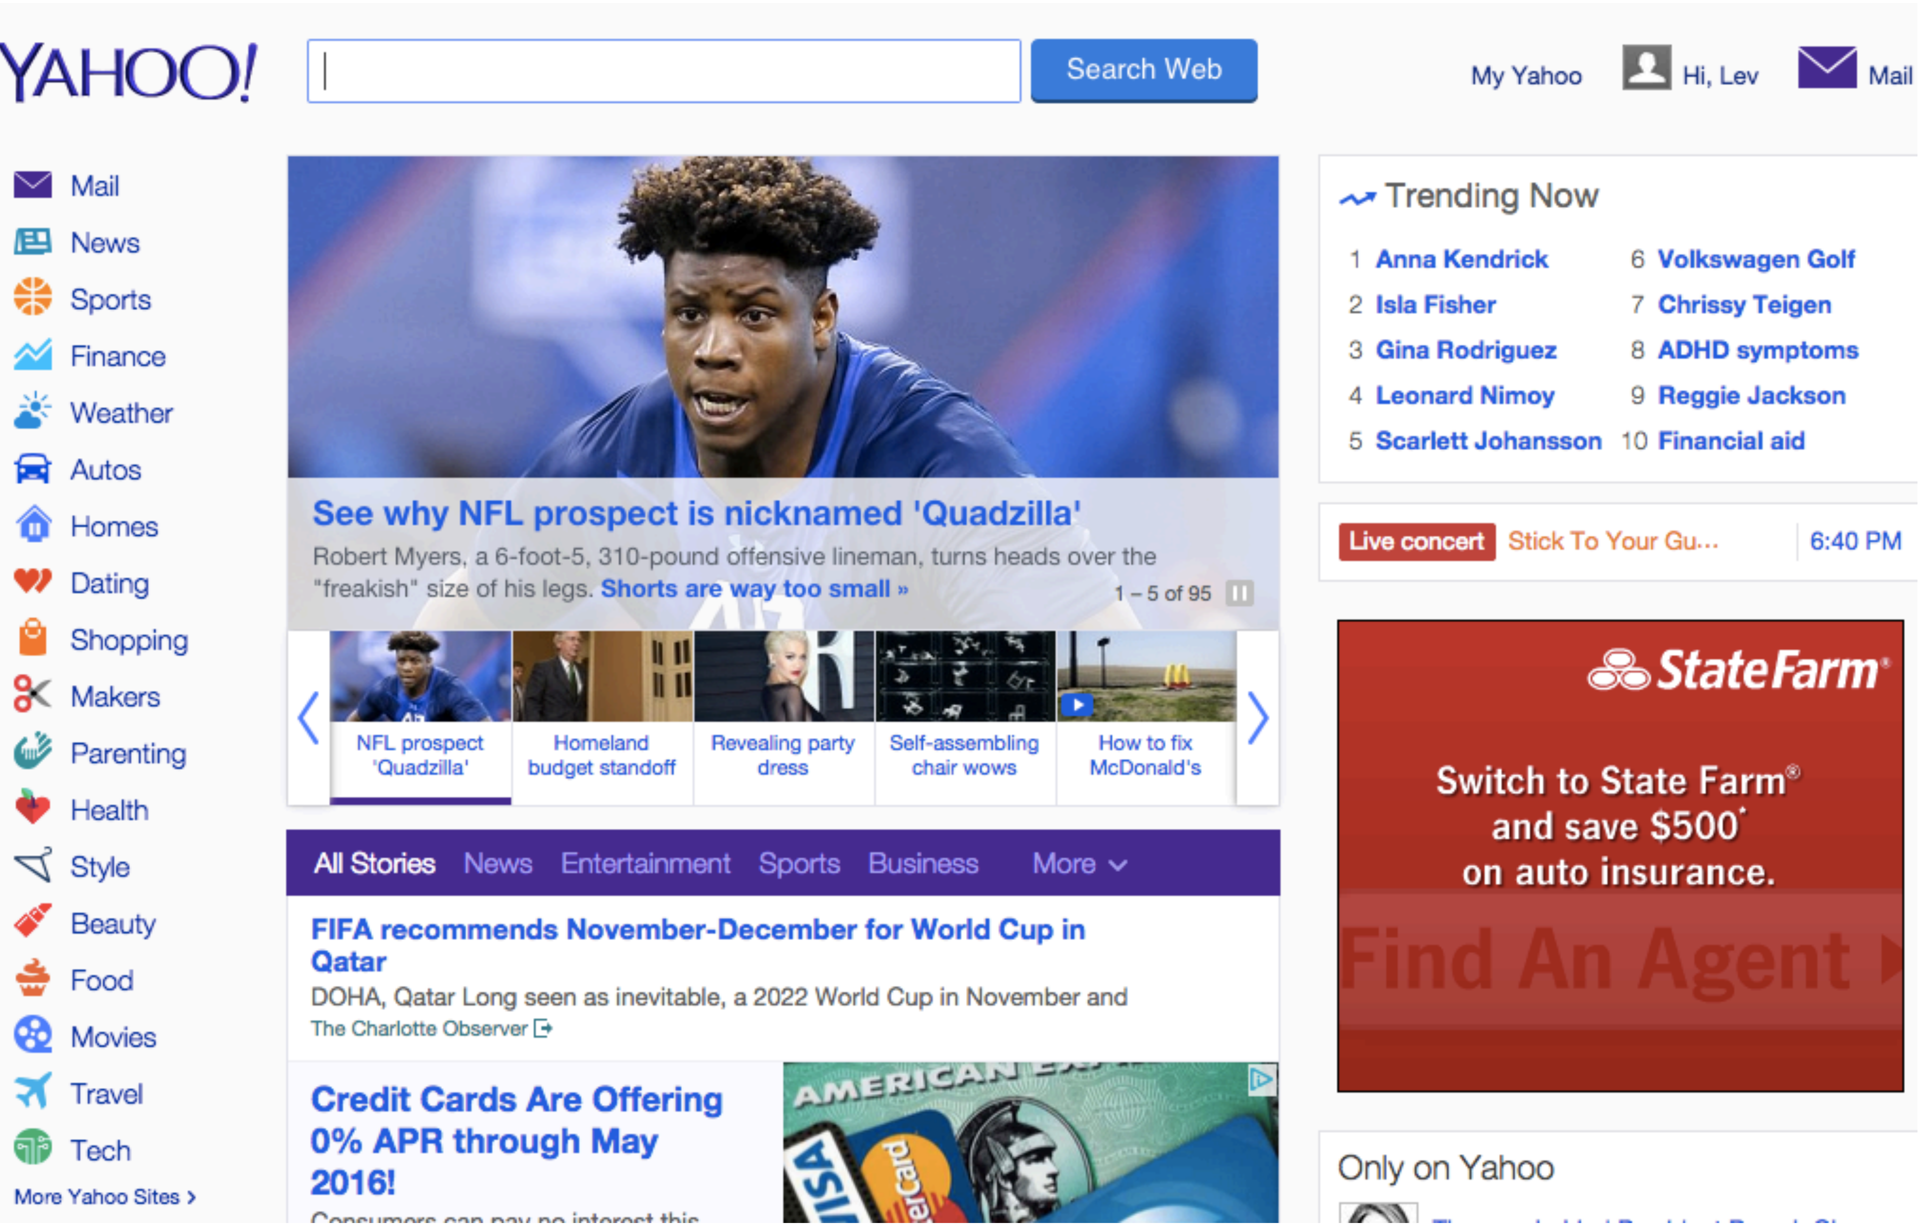
YAHOO!
Search Web
My Yahoo
Hi, Lev
Mail
Mail
News
Sports
Finance
Weather
Autos
Homes
Dating
Shopping
Makers
Parenting
Health
Style
Beauty
Food
Movies
Travel
Tech
More Yahoo Sites ›
See why NFL prospect is nicknamed 'Quadzilla'
Robert Myers, a 6-foot-5, 310-pound offensive lineman, turns heads over the "freakish" size of his legs. Shorts are way too small »
1 – 5 of 95
NFL prospect 'Quadzilla'
Homeland budget standoff
Revealing party dress
Self-assembling chair wows
How to fix McDonald's
All Stories
News
Entertainment
Sports
Business
More
FIFA recommends November-December for World Cup in Qatar
DOHA, Qatar Long seen as inevitable, a 2022 World Cup in November and
The Charlotte Observer
Credit Cards Are Offering 0% APR through May 2016!
Trending Now
1 Anna Kendrick
2 Isla Fisher
3 Gina Rodriguez
4 Leonard Nimoy
5 Scarlett Johansson
6 Volkswagen Golf
7 Chrissy Teigen
8 ADHD symptoms
9 Reggie Jackson
10 Financial aid
Live concert
Stick To Your Gu...
6:40 PM
StateFarm
Switch to State Farm® and save $500* on auto insurance.
Find An Agent
Only on Yahoo

Evaluation

- We chose a policy class for which we could efficiently keep track of the weights.
  - Created 5 clusters, using user features.
  - Policies mapped clusters to article choices.
  - Ran on personalized news article recommendations for Yahoo! front page.
- We used a learning bucket on which we ran the algorithms and a deployment bucket on which we ran the greedy (best) learned policy.

Experiments

Reported estimated (normalized) click-through rates on front page news. Over 41M user visits. 253 total articles. 21 candidate articles per visit.

	EXP4P	EXP4	ε-greedy
Learning eCTR	1.0525	1.0988	1.3829
Deployment eCTR	**1.6512**	1.5309	1.4290

Why does this work in practice?

[McMahan '11]

A General Efficient Algorithm? [DHKKL**R**Z '11]

EXP4.P's regret grows only logarithmically with N.

this suggests

We could compete with a huge set of policies!
(e.g. **N=K**^{100} becomes **10 log**$^{1/2}$ **K** in the regret)

however

Exp4.P explicitly "keeps track" of all policies. Reading in all recommendations, for large N, would take too long.

Reduce to Supervised Learning!

(Idea from [Langford-Zhang '07])

- "Competing" with an exponentially large set of policies is commonplace in supervised learning.
- Recommendations of the policies/functions don't need to be explicitly read when the policy class has structure!

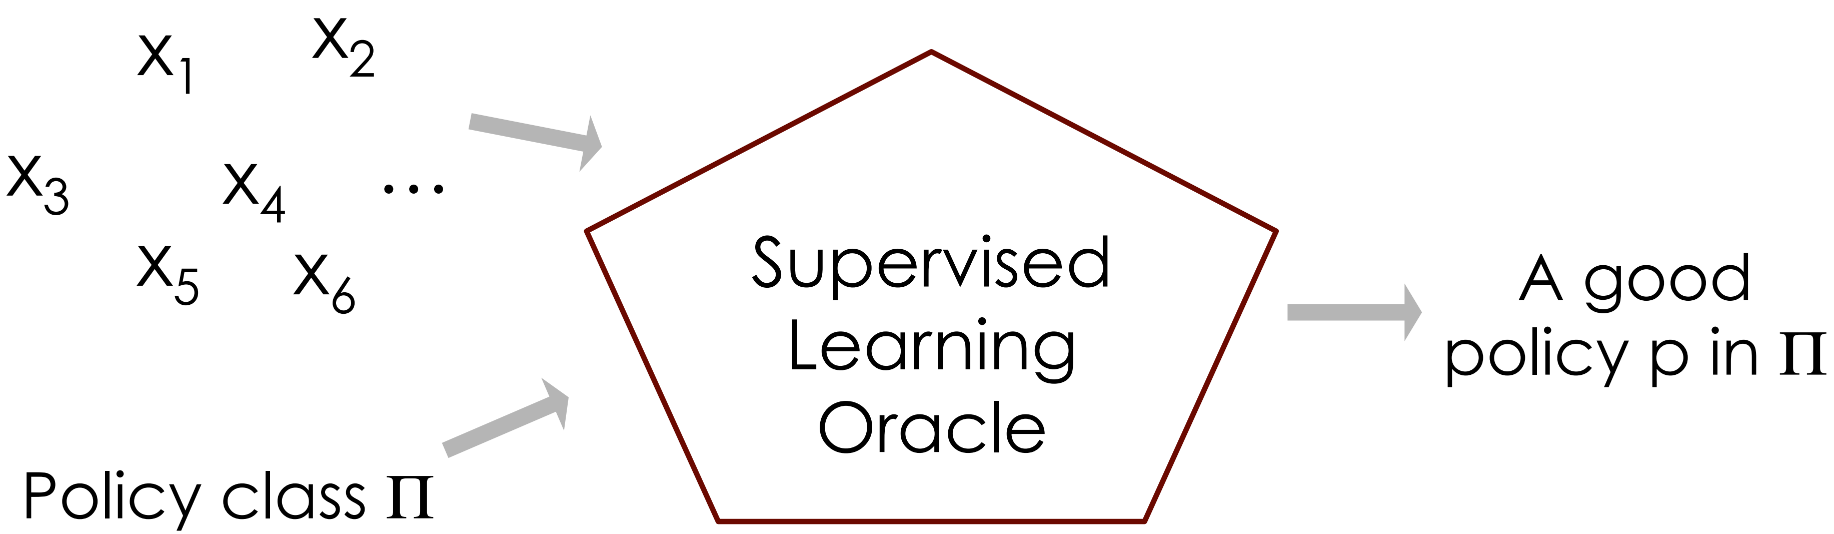

Reduce to Supervised Learning!

(Idea from [Langford-Zhang '07])

- "Competing" with an exponentially large set of policies is commonplace in supervised learning.
- Recommendations of the policies/functions don't need to be explicitly read when the policy class has structure!

context: x_1 x_2 x_3

1 2 3 4 … T

1 yes

2 no

3

⋮

k

5 1

1 4

k 3

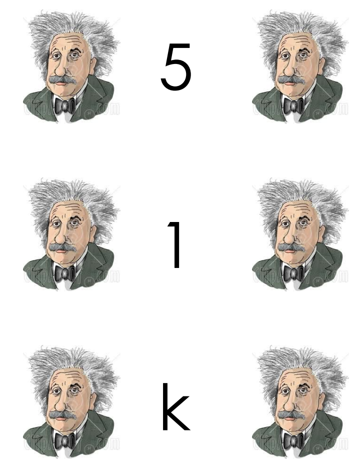

N experts/functions

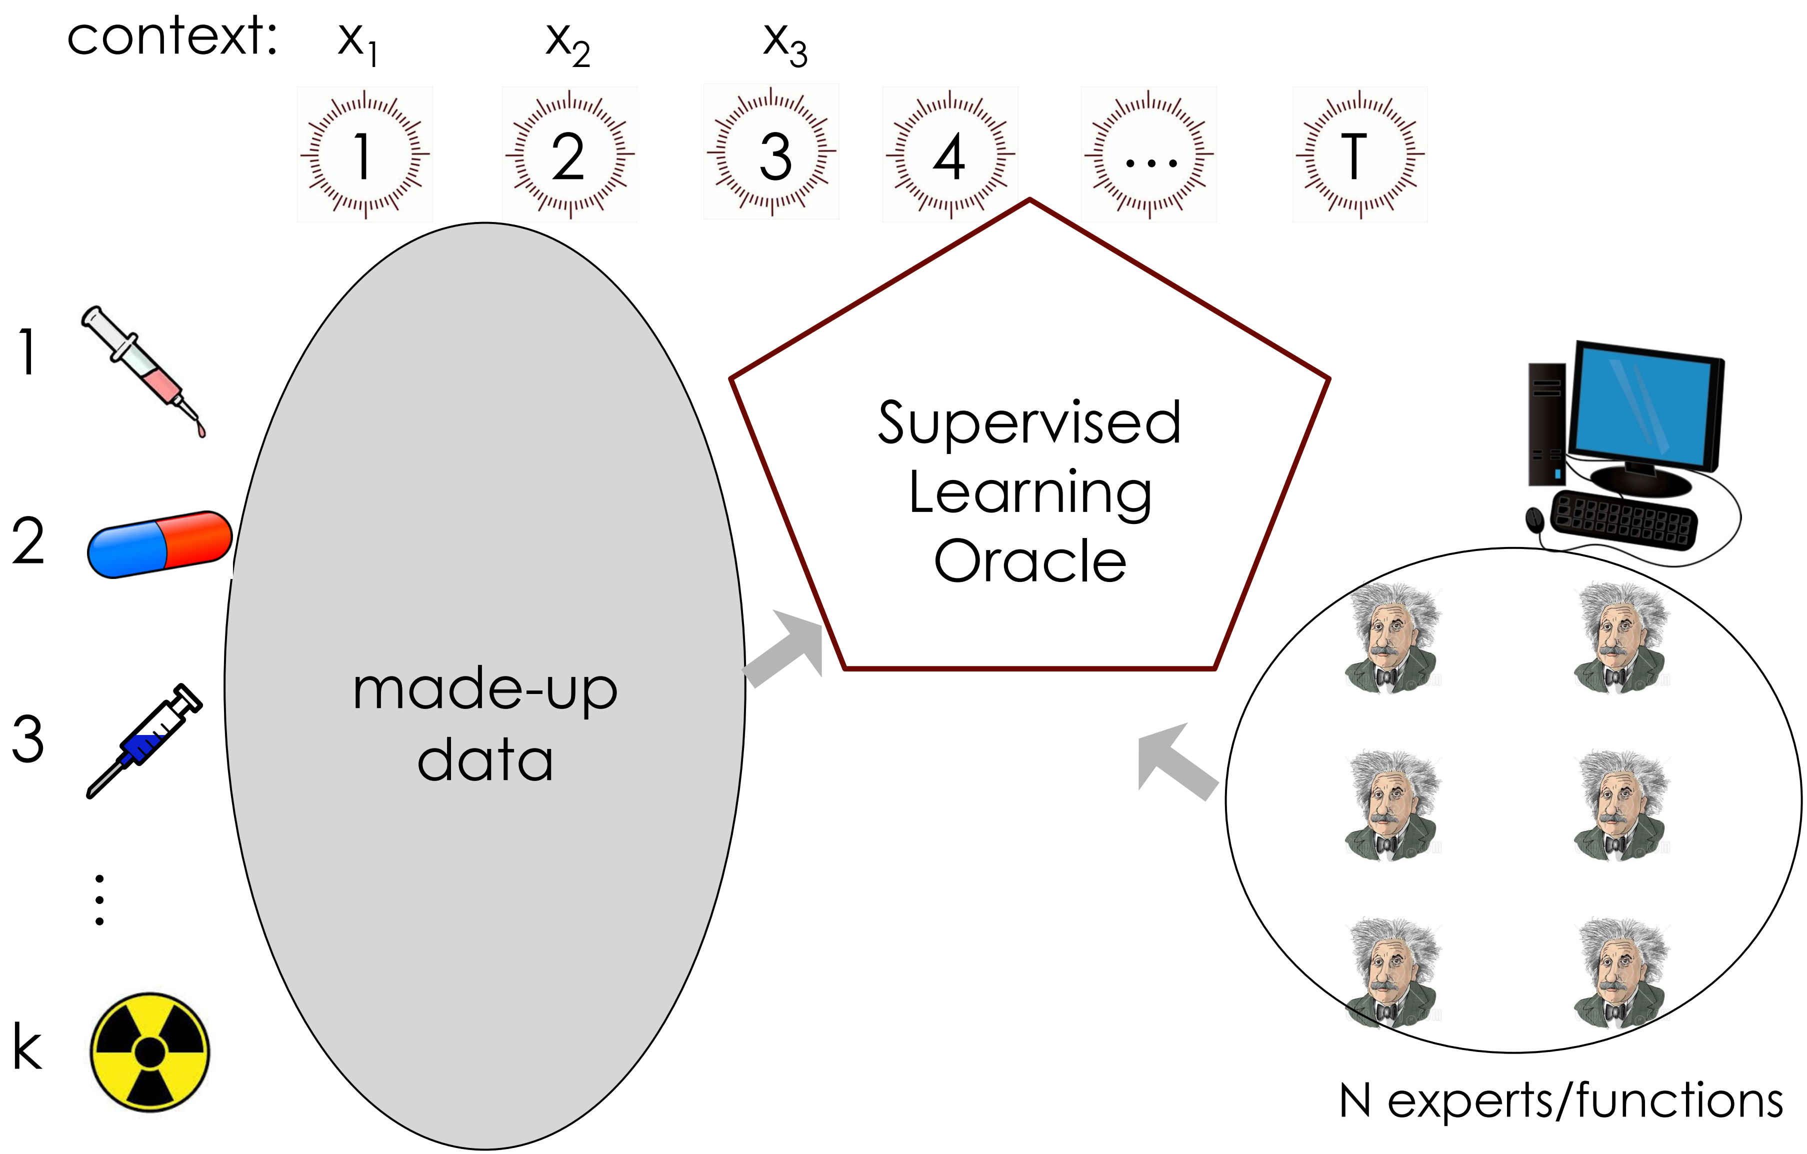
context: x_1 x_2 x_3
1
2
3
4
...
T
1
2
3
⋮
k
made-up data
Supervised Learning Oracle
N experts/functions

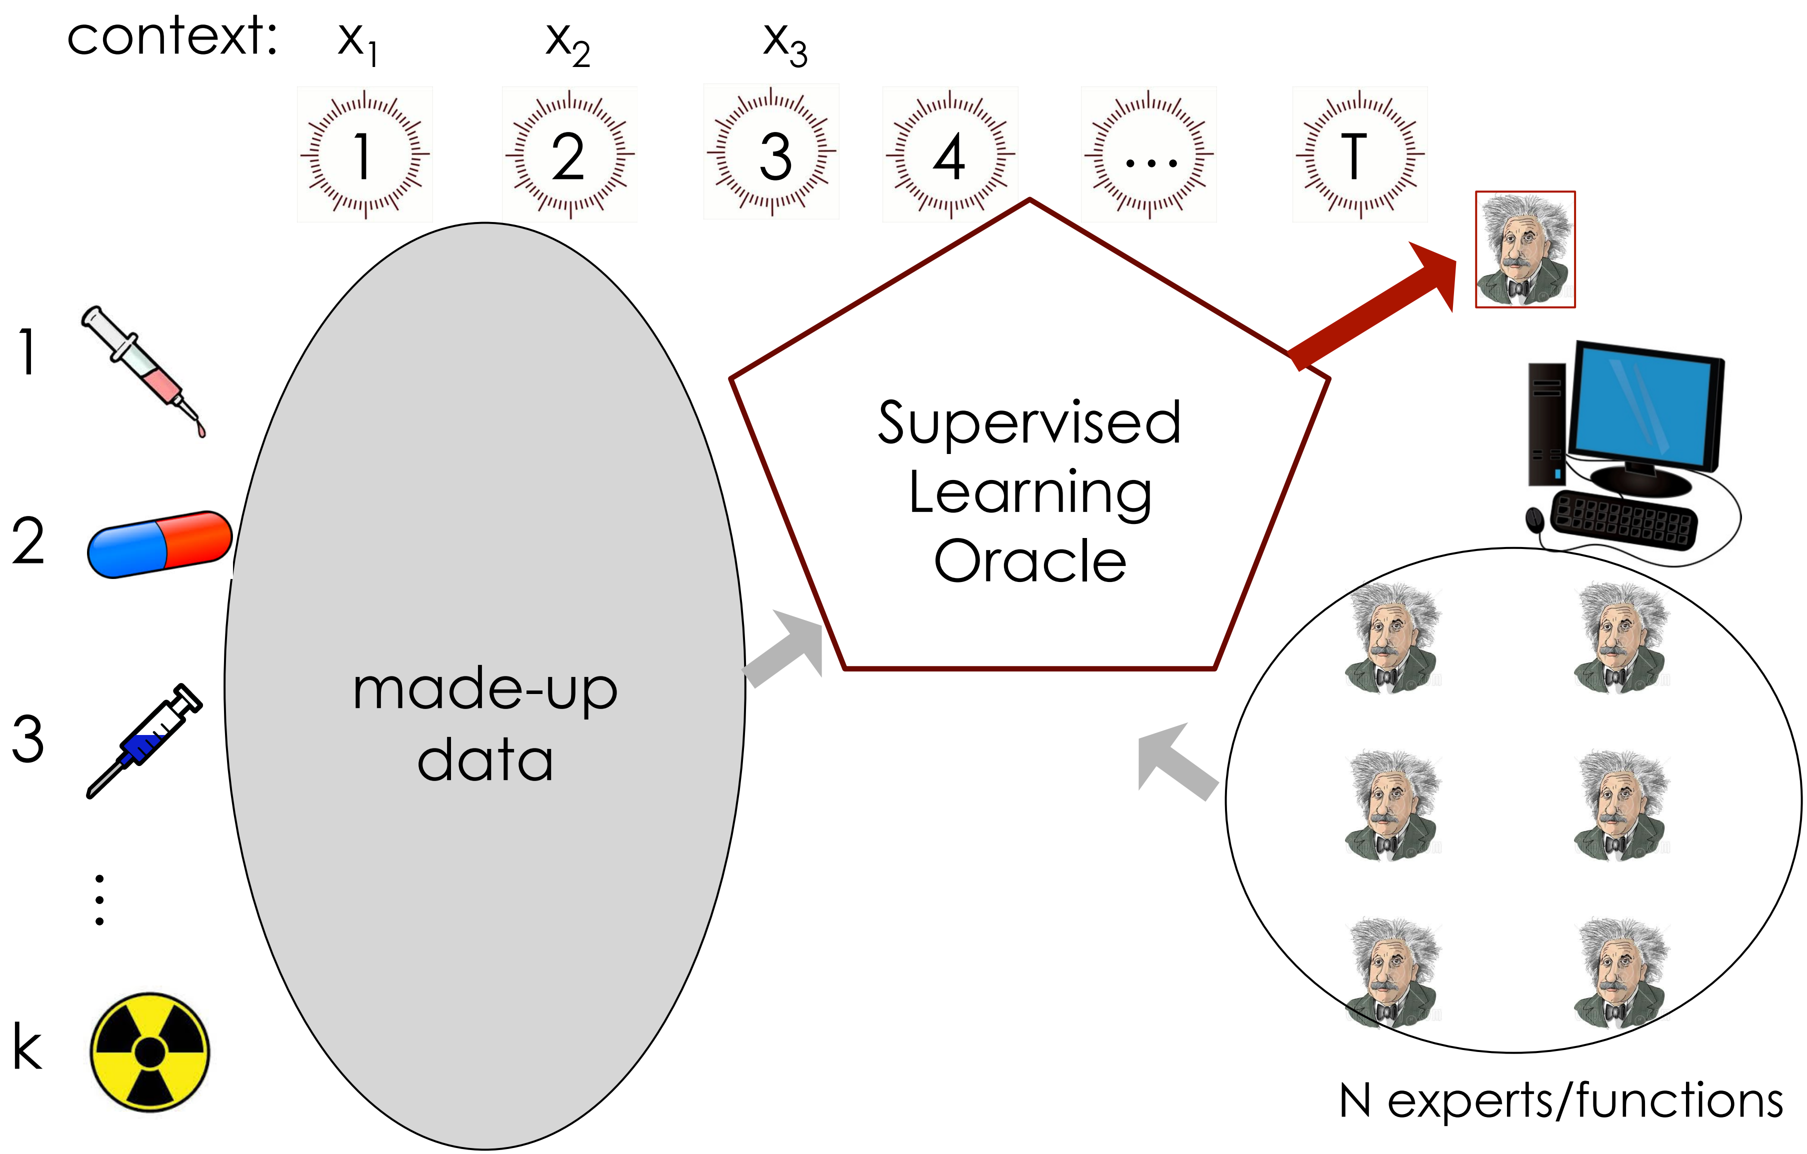
context: x_1 x_2 x_3
1
2
3
4
...
T
1
2
3
⋮
k
made-up data
Supervised Learning Oracle
N experts/functions

Thm: [Dudik-Hsu-Kale-Karampatziakis-Langford-**R**-Zhang '11]: For any $\delta > 0$, w.p. at least $1 - \delta$, given access to a learning oracle, R-UCB has regret $O((KT \ln (NT/\delta))^{1/2}$ in the stochastic contextual bandit setting and runs in time poly(K,T, ln N).

Thm: [Dudik-Hsu-Kale-Karampatziakis-Langford-**R**-Zhang '11]:
For any $\delta > 0$, w.p. at least $1-\delta$, given access to a learning oracle, **R-UCB** has regret **$O((KT \ln (NT/\delta))^{1/2}$** in the stochastic contextual bandit setting and runs in time **poly(K,T, ln N).**

if arms chosen among only "good" policies w/ variance < 2K, we win
can prove this exists via a minimax theorem

↓

this can be softened to occasionally allow choosing of bad policies
via "randomized" upper confidence bounds

↓

creates a problem of how to choose arms as to satisfy the constraints
expressed as convex optimization problem

↓

solvable by ellipsoid algorithm
can implement a separation oracle with the supervised learning oracle

Thm: [Dudik-Hsu-Ka ... -R-Zhang '11]:
For any δ >0 ... learning
oracle, R-... hastic
conte...).

if arm ... we win

thi ... ies

creat ... aints

can implement ... d learning oracle

Big theoretical
breakthrough!

But not practical.
(and needs stochastic assumption)

Taming the Monster: A Fast and Simple Algorithm for Contextual Bandits
[Agarwal-et al. '14]

A research goal of mine: make this work in adversarial model.

Applying to Public Health

UIC SCHOOL OF PUBLIC HEALTH
UNIVERSITY OF ILLINOIS AT CHICAGO

Search
GO

About SPH | Academics | Admissions | Departments | Research

Research
- Announcements
- Seed Funding Competition
- Annual Research and Practice Awards Day
- New Awards
- Getting To Know UIC SPH Researchers
- Faculty Research Interests
- SPH Proposal Development
- Community Connections
- Current Sponsored Activities
- Centers Institutes & Programs
- Traineeships & Fellowships
- Contact Office of Research Services

Research / Centers Institutes & Programs /

Community Outreach Intervention Projects

ABOUT COIP

The Community Outreach Intervention Projects ("COIP"), School of Public Health, University of Illinois at Chicago was founded in 1986 to address HIV/AIDS, particularly among people who use drugs. COIP operates from storefront sites in Austin, Humboldt Park, West Englewood, South Chicago, and Uptown. Other neighborhoods are served by COIP's motorhome and mobile van units. COIP's interventions are known for their use of the Indigenous Leader Outreach Model, which employs former drug users to deliver services and assist in conducting research.

COIP's services include street outreach, counseling and testing for HIV, syphilis and other infectious diseases associated with substance use, case management for people who are HIV positive, syringe exchange, drug abuse and risk reduction counseling, support groups, educational activities, and a program that enhances linkages to care for HIV positive women exiting jail. COIP also makes many referrals to other providers such as drug treatment programs. Through a collaboration with UIC's Community Clinic Network, all COIP's storefront sites provide free medical, mental health and pharmacy care for people living with HIV, and one site offers free dental care.

COIP also conducts research to better understand HIV/AIDS in Chicago communities. Recent studies have examined or

Consortium for Modeling and Analysis of Treatments and Interventions

c-mati.org

Feature-Efficient Prediction

with

Huang and Powers

Feature-Efficient Prediction Examples

- Medical testing

  Want to predict what patients are sick with, but tests might be expensive or dangerous.

- Displaying internet results

  Want to give users the best results you can, but if you don't give results within 300 milliseconds, users will leave.

Model

- Goal is to do supervised learning, using a limited number of features in test-time.
  - Given a budget on total cost: on each example, the learner must look at no more features than allowed by the budget.
  - Each feature has an associated cost.
  - Budget only limited in test data, not training.

- Predictors that do this are feature-efficient.

Lots of work on this problem

- Sequential analysis: when to stop sequential clinical trials.
  [Wald '47] and [Chernoff '72]
- PAC learning with incomplete features.
  [Ben-David-Dichterman '93] and [Greiner et al. '02]
- Robust prediction with missing features.
  [Globerson-Roweis '06]
- Learning linear functions by few features
  [Cesa-Bianchi et al. '10]
- Incorporating feature costs in CART impurity [Xu et al. '12]
- MDPs for feature selection [He et al. '13]

A “Weigh + Measure” Idea

[**R** ’11]

- An ensemble is usually a weighted vote of many simple rules.
- The simple rules are usually feature-efficient.
- Take a vote of only a few of the rules.

AdaBoost in Pictures (Slides from Schapire)

Toy Example

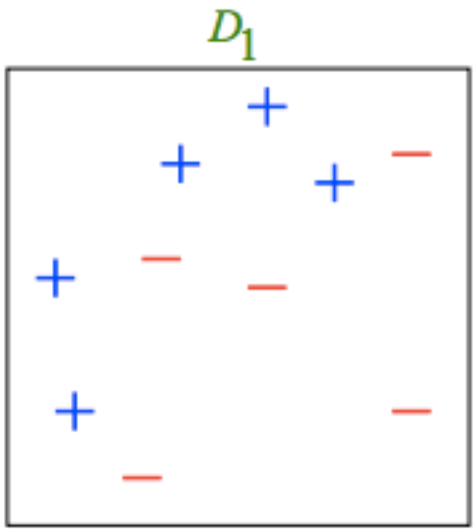

weak classifiers = vertical or horizontal half-planes

AdaBoost in Pictures (Slides from Schapire)

Round 1

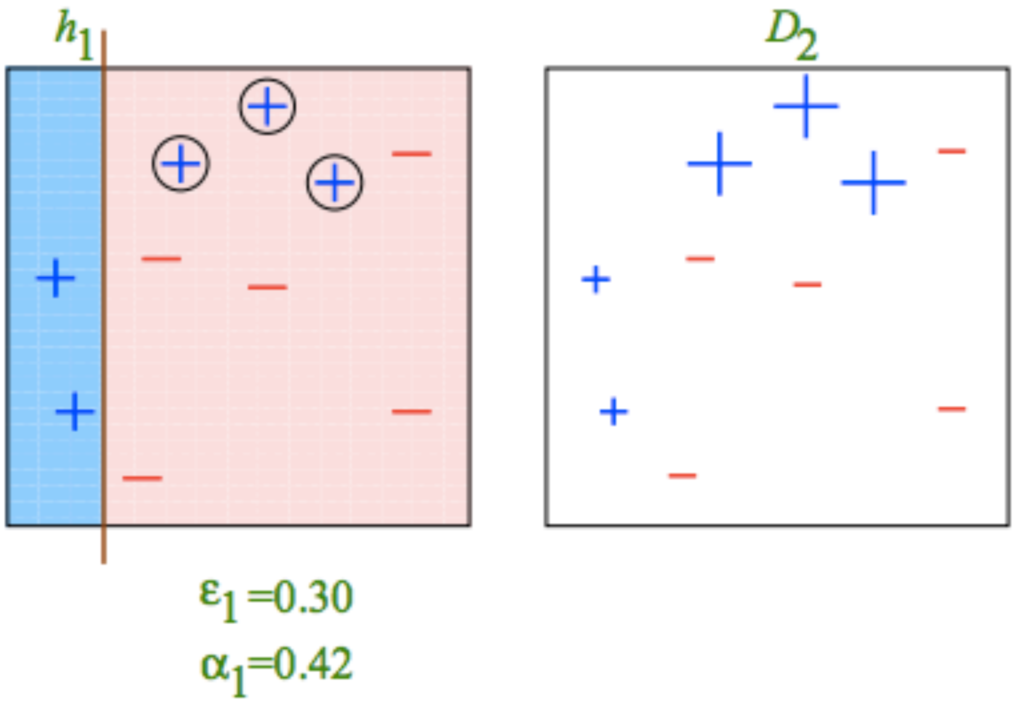

AdaBoost in Pictures (Slides from Schapire)

Round 2

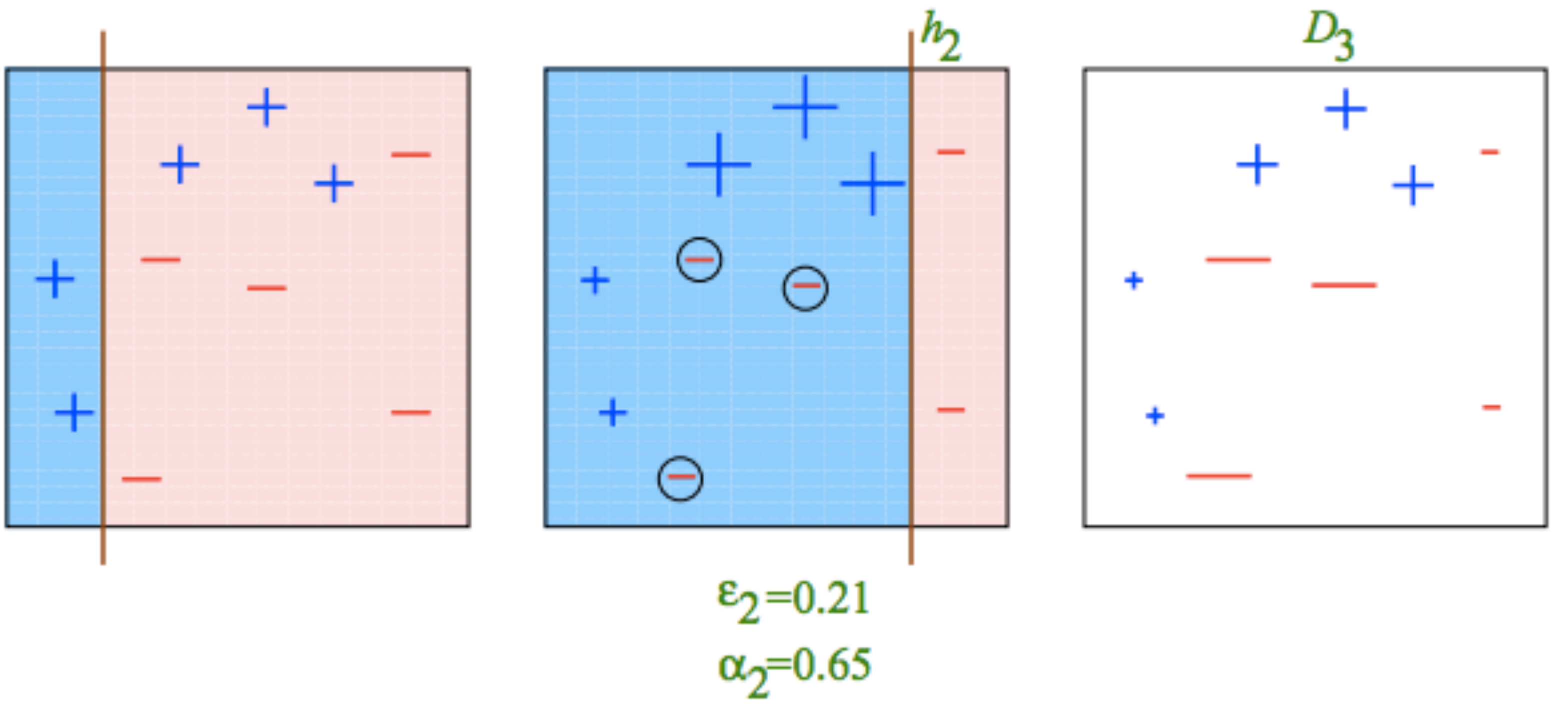

AdaBoost in Pictures (Slides from Schapire)

Round 3

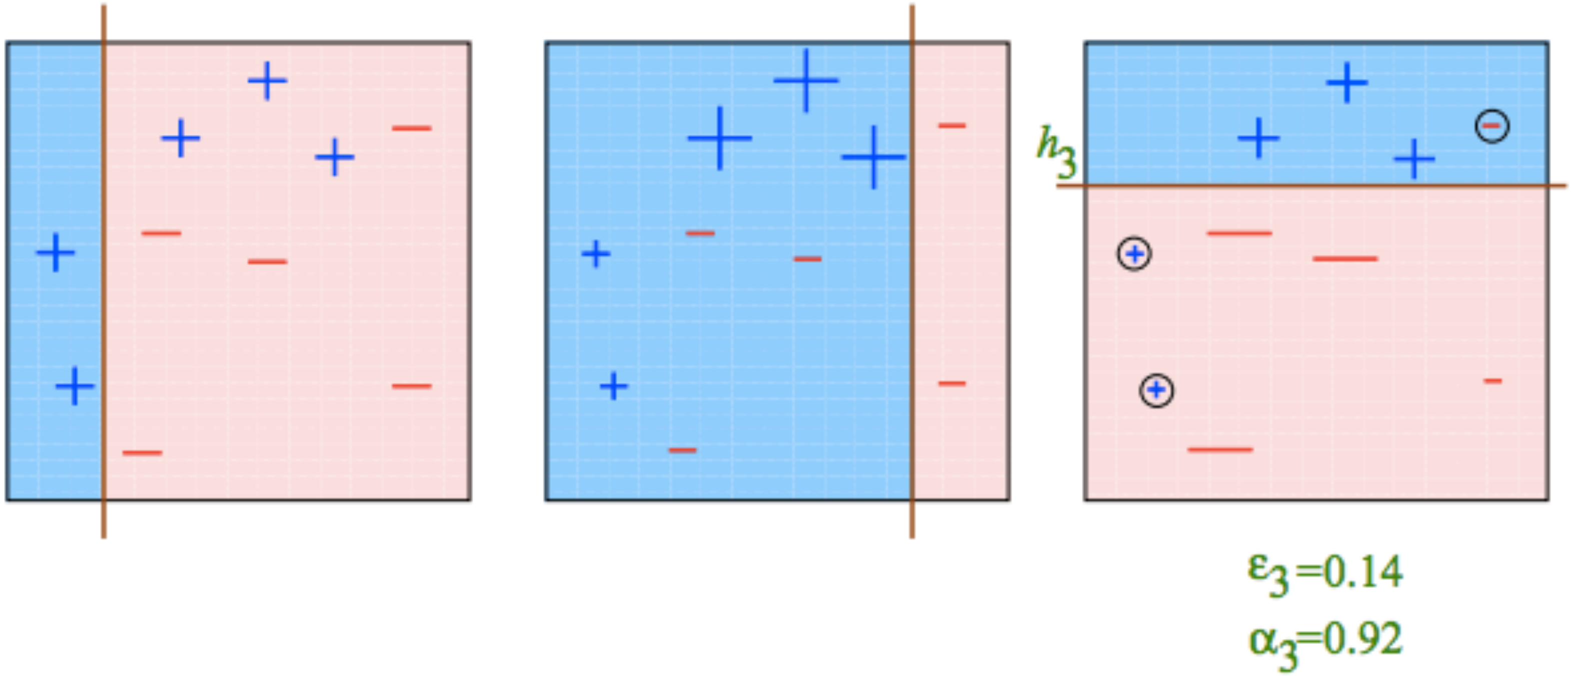

AdaBoost in Pictures (Slides from Schapire)

Final Classifier

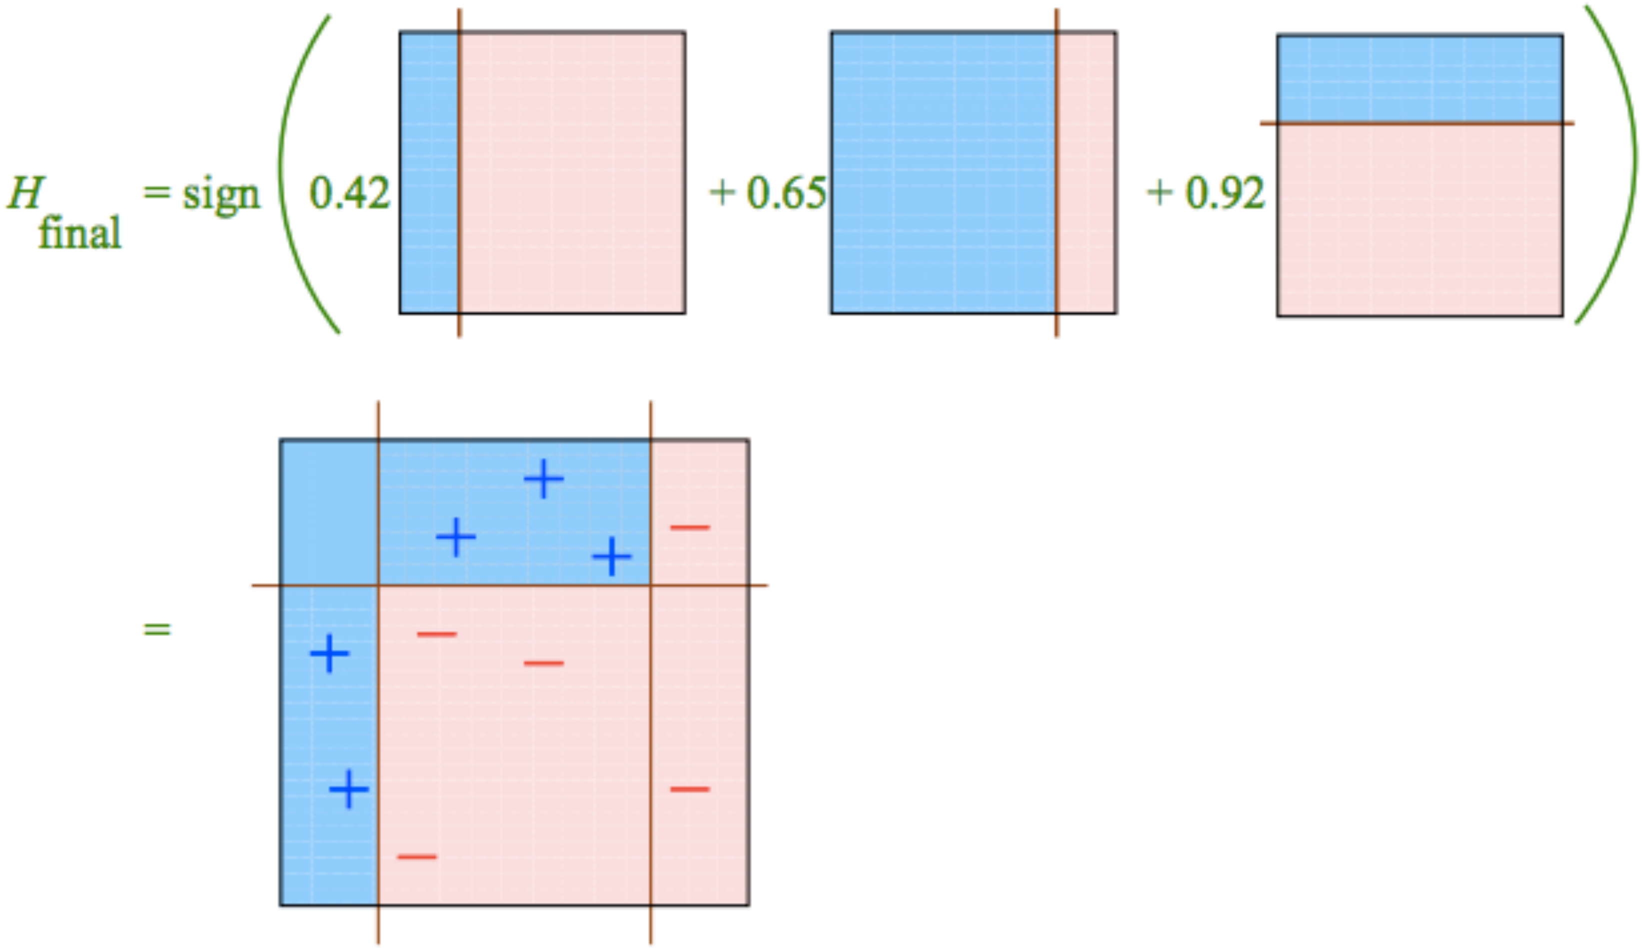

AdaBoostRS [R '11]

Training: train AdaBoost (or any ensemble).

Prediction:

1. Sample the weak learners depending on their voting weights and feature costs.
2. Take a importance-weighted vote of the sampled weak learners.

Intuition:

If ensemble has strong preference, sampling will converge fast. If ensemble is split, who cares?

(Thm resembles margin bound [Schapire et al. '98])

Experiments with AdaBoostRS

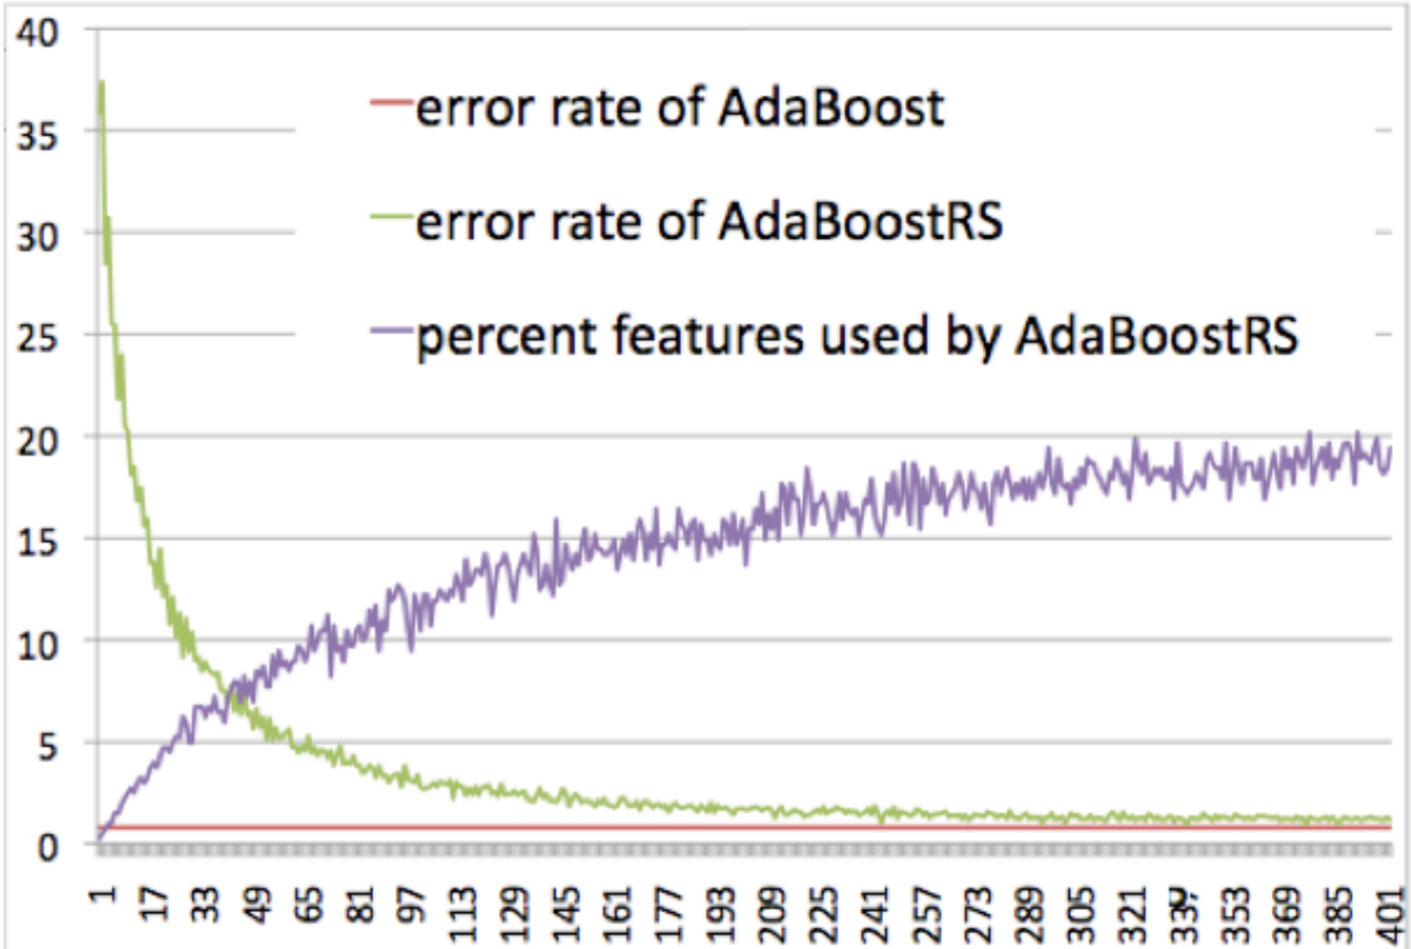

On ocr17 dataset. x-axis is number of samples taken.

Room for Improvement

Can we improve by moving the optimization into training?

Turns out: yes, by a lot! [Huang-Powers-**R** '14]

- Naïve idea: train AdaBoost until budget runs out
- Improvement: choose weak learners more wisely

`AdaBoost (S )` where: $S \subset X \times \{-1, +1\}$

given: $(x_1, y_1), \ldots, (x_m, y_m) \in S$
initialize $D_1(i) = \frac{1}{m}$
for $t = 1, \ldots, T$ **do**
    train base learner using distribution D_t.
    get $h_t \in \mathcal{H} : X \to \{-1, +1\}$.

    choose $\alpha_t = \frac{1}{2} \ln \frac{1+\gamma_t}{1-\gamma_t}$, where $\gamma_t = \sum_i D_t(i) y_i h_t(x_i)$.
    update $D_{t+1}(i) = D_t(i) \exp(\alpha_t y_i h_t(x_i)) / Z_t$,
end for
output the final classifier $H(x) = \mathrm{sign}\left(\sum_{t=1}^{T} \alpha_t h_t(x)\right)$

`AdaBoostBT(S,B,C)` where: $S \subset X \times \{-1,+1\}$, $B > 0$, $C : [n] \to \mathbb{R}^+$

given: $(x_1, y_1), \ldots, (x_m, y_m) \in S$
initialize $D_1(i) = \frac{1}{m}$, $B_1 = B$
for $t = 1, \ldots, T$ **do**
    train base learner using distribution D_t.
    get $\boxed{h_t \in \mathcal{H} : X \to \{-1,+1\}}$.
    if the total cost of the unpaid features of h_t exceeds B_t **then**
        set $T = t - 1$ and **end for**
    else set B_{t+1} as B_t minus the total cost of the unpaid features of h_t, marking them as paid
    choose $\alpha_t = \frac{1}{2}\ln\frac{1+\gamma_t}{1-\gamma_t}$, where $\gamma_t = \sum_i D_t(i) y_i h_t(x_i)$.
    update $D_{t+1}(i) = D_t(i) \exp(\alpha_t y_i h_t(x_i))/Z_t$,
end for
output the final classifier $H(x) = \mathrm{sign}\left(\sum_{t=1}^{T} \alpha_t h_t(x)\right)$

How to choose weak learner h_t?

Training error of AdaBoost is bounded by

[Freund-Schapire '97]

$$\hat{\Pr}[H(x) \neq y] \leq \prod_{t=1}^{T} \sqrt{1 - \gamma_t^2}$$

With budgets, we need to consider two effects:

- high edges make individual terms smaller
- low costs allow for more terms in the product

Two Optimizations

[Huang-Powers-**R** '14]

First idea: assume all future rounds will behave like current. Leads to optimization

1)
$$h_t = \text{argmin}_{h \in \mathcal{H}} \left(\left(1 - \gamma_t(h)^2\right)^{\frac{1}{c(h)}} \right)$$

Second idea: smoothed version of first.

2)
$$h_t = \text{argmin}_{h \in \mathcal{H}} \left(\left(1 - \gamma_t(h)^2\right)^{\frac{1}{(B - B_t) + c(h)}} \right)$$

Experiments with costs ~ U(0,2)

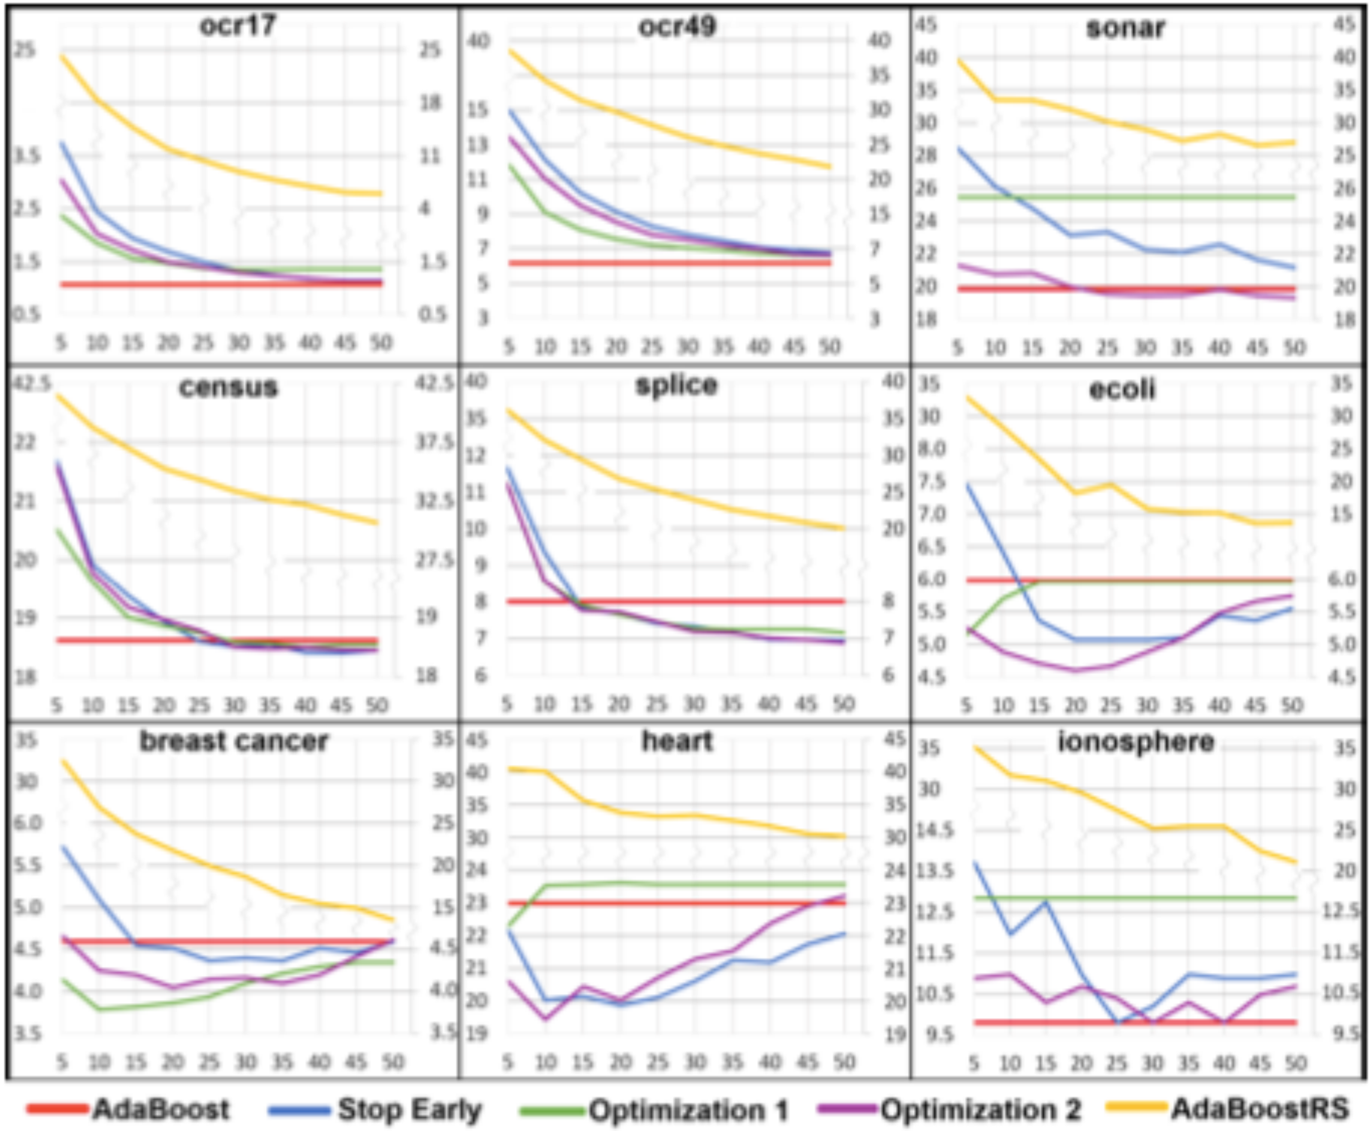

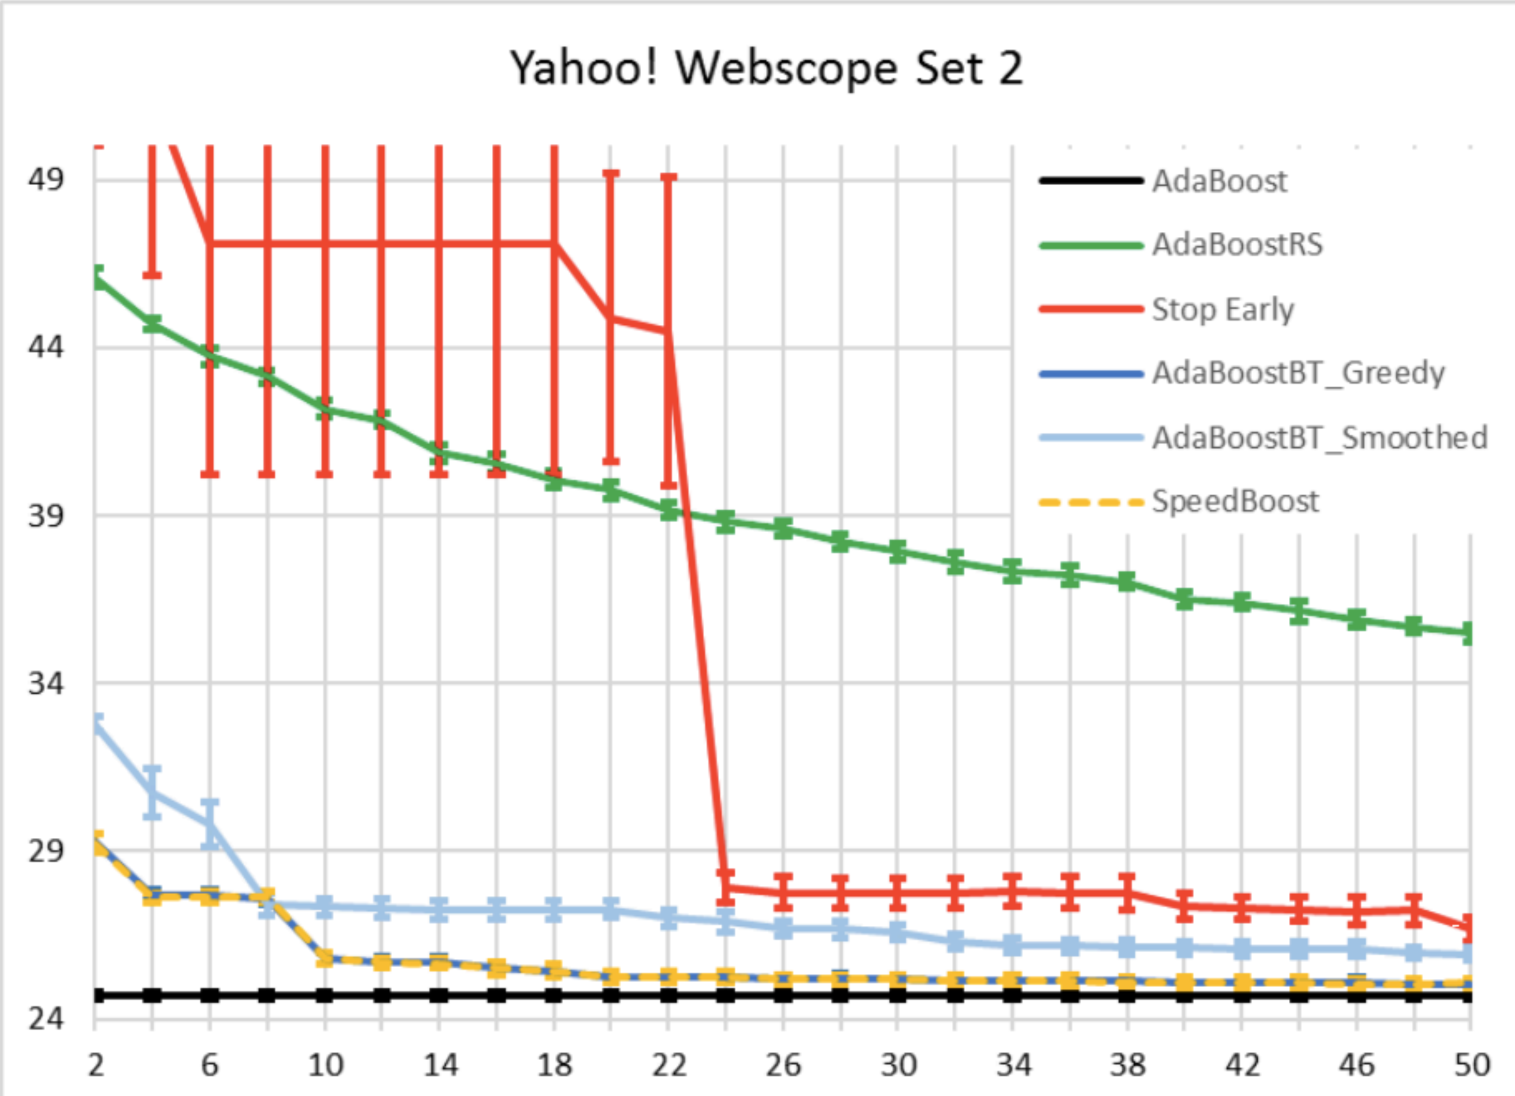

Yahoo! Webscope Set 2
49
44
39
34
29
24
2
6
10
14
18
22
26
30
34
38
42
46
50
AdaBoost
AdaBoostRS
Stop Early
AdaBoostBT_Greedy
AdaBoostBT_Smoothed
SpeedBoost

Active Learning

with

Liu and Ziebart

Pool-Based Active Learning

- A pool based active learning algorithm sequentially chooses which labels to solicit from a pool of examples. [Lewis-Gale '94]
  - Usually constructs estimate of conditional label distribution $P(y|x)$ from labeled dataset.
  - Uses own estimate to select next datapoint label.

(I'll focus on logloss, but ideas are more general)

Uncertainty Sampling

- Many active learning strategies employ uncertainty sampling – selecting examples about which the algorithm is least certain.
- Other strategies assess how a label:
  - is expected to change model [Settles-Craven '08]
  - reduces an upper bound on the generalization error in expectation [Mackay '92]
  - represents the input patterns of remaining unlabeled data [Settles '12]

A Problem

Current active learning algorithms often perform poorly in practice [Attenberg-Provost '11].

Why?

- In order to be take advantage of active learning, a biased label solicitation strategy should be used.
- Most current active learning strategies are overconfident, given this bias.

Typical Active Learner Behavior

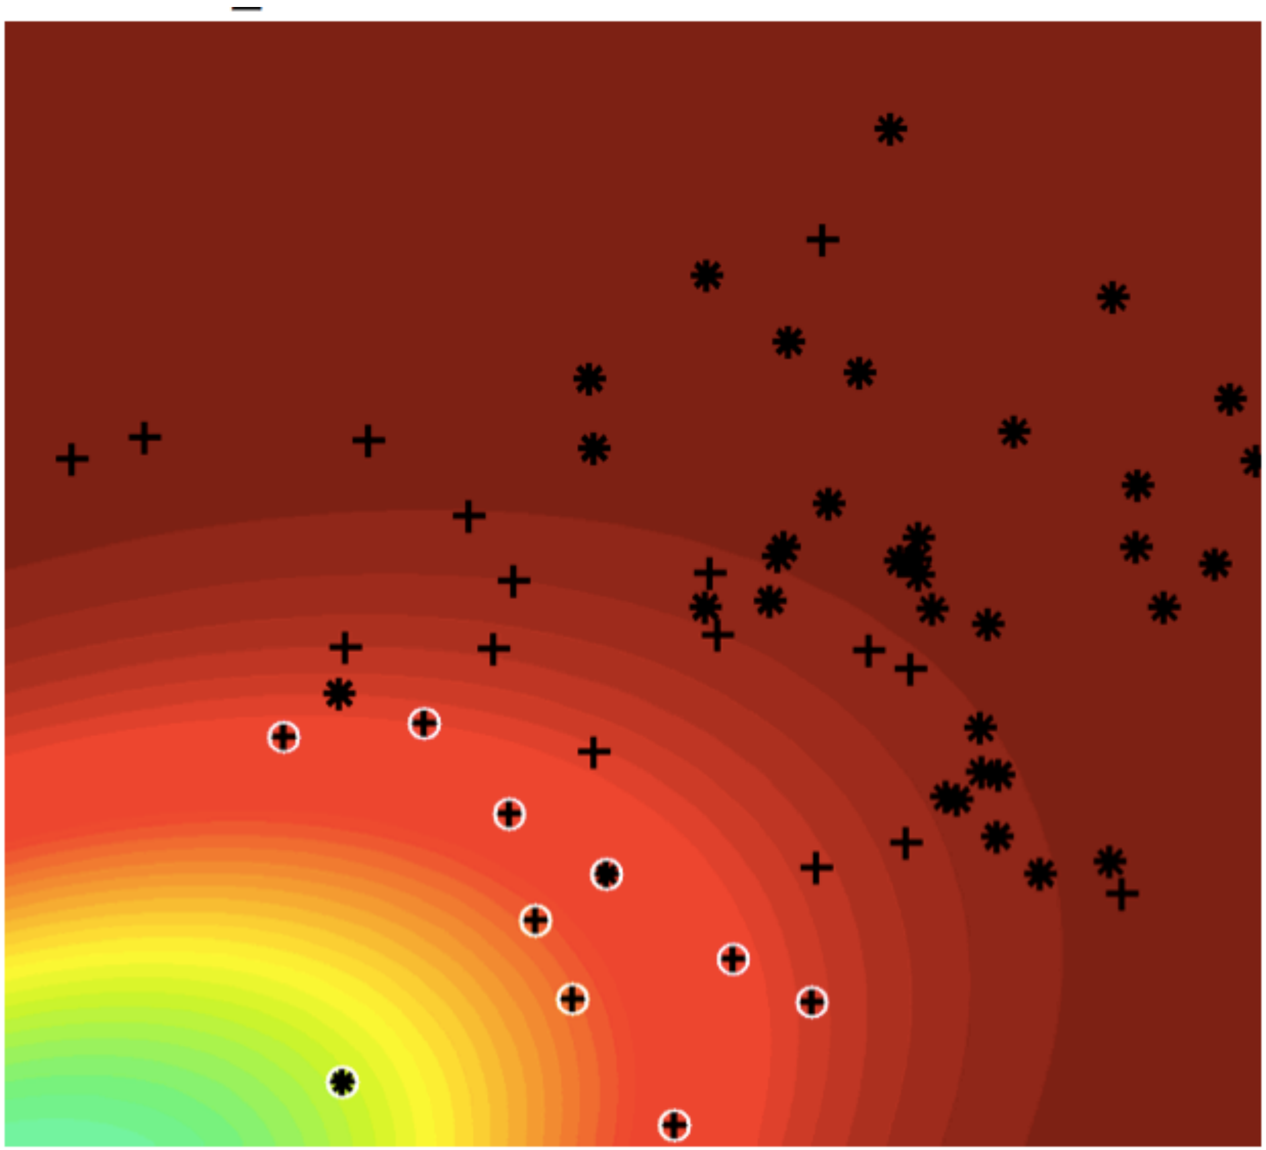

Desired Behavior

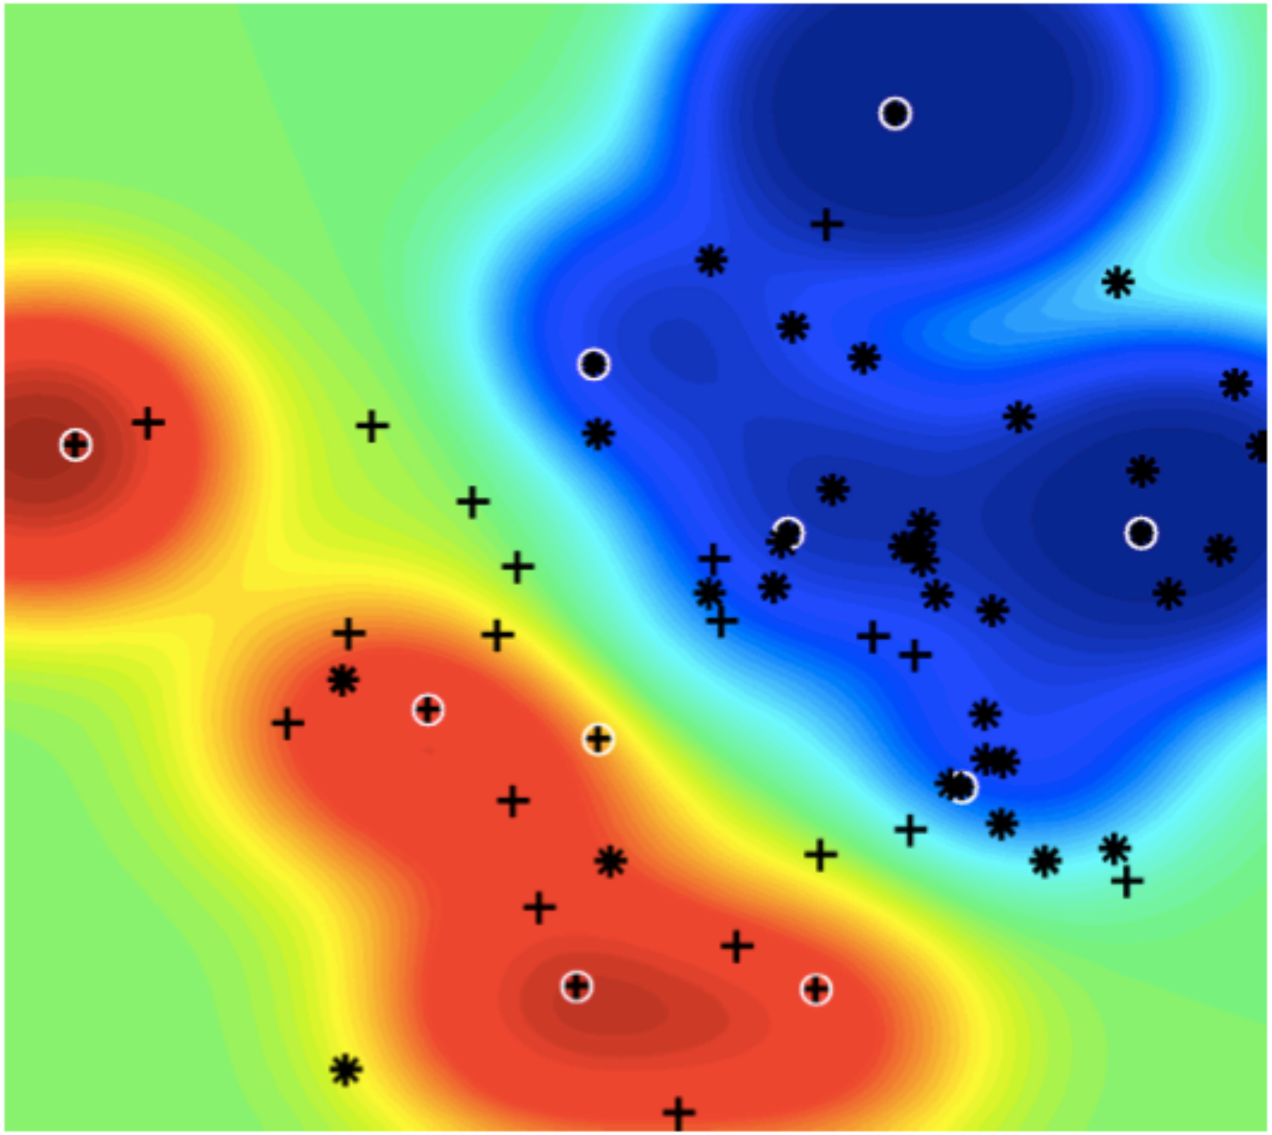

Some Attempts to Fix This

- Seeding the active learner with a small random set [Dligach-Palmer '11].
- Restricting the active learner to a small set of examples [Schein-Ungar '07].
- Etc.

However, these modifications treat the symptoms of optimistic modeling and biased sampling and restrict the active learner, undermining its benefit.

Biased Label Solicitation

Key idea: Active learning always leads to sample selection bias exists. Here, known as covariate shift -- $P(Y \mid X)$ is shared in source and target distributions.

Tackling covariate shift is difficult. A common approach is importance re-weighting of source samples x according to $P_{trg}(x)/P_{src}(x)$ and minimizing a reweighted version of the loss [Shimodaria '00].

- This converges slowly [Cortes-Mansour-Mohri '10] and the variance of estimates is too high to be useful.

Logistic Regression Models

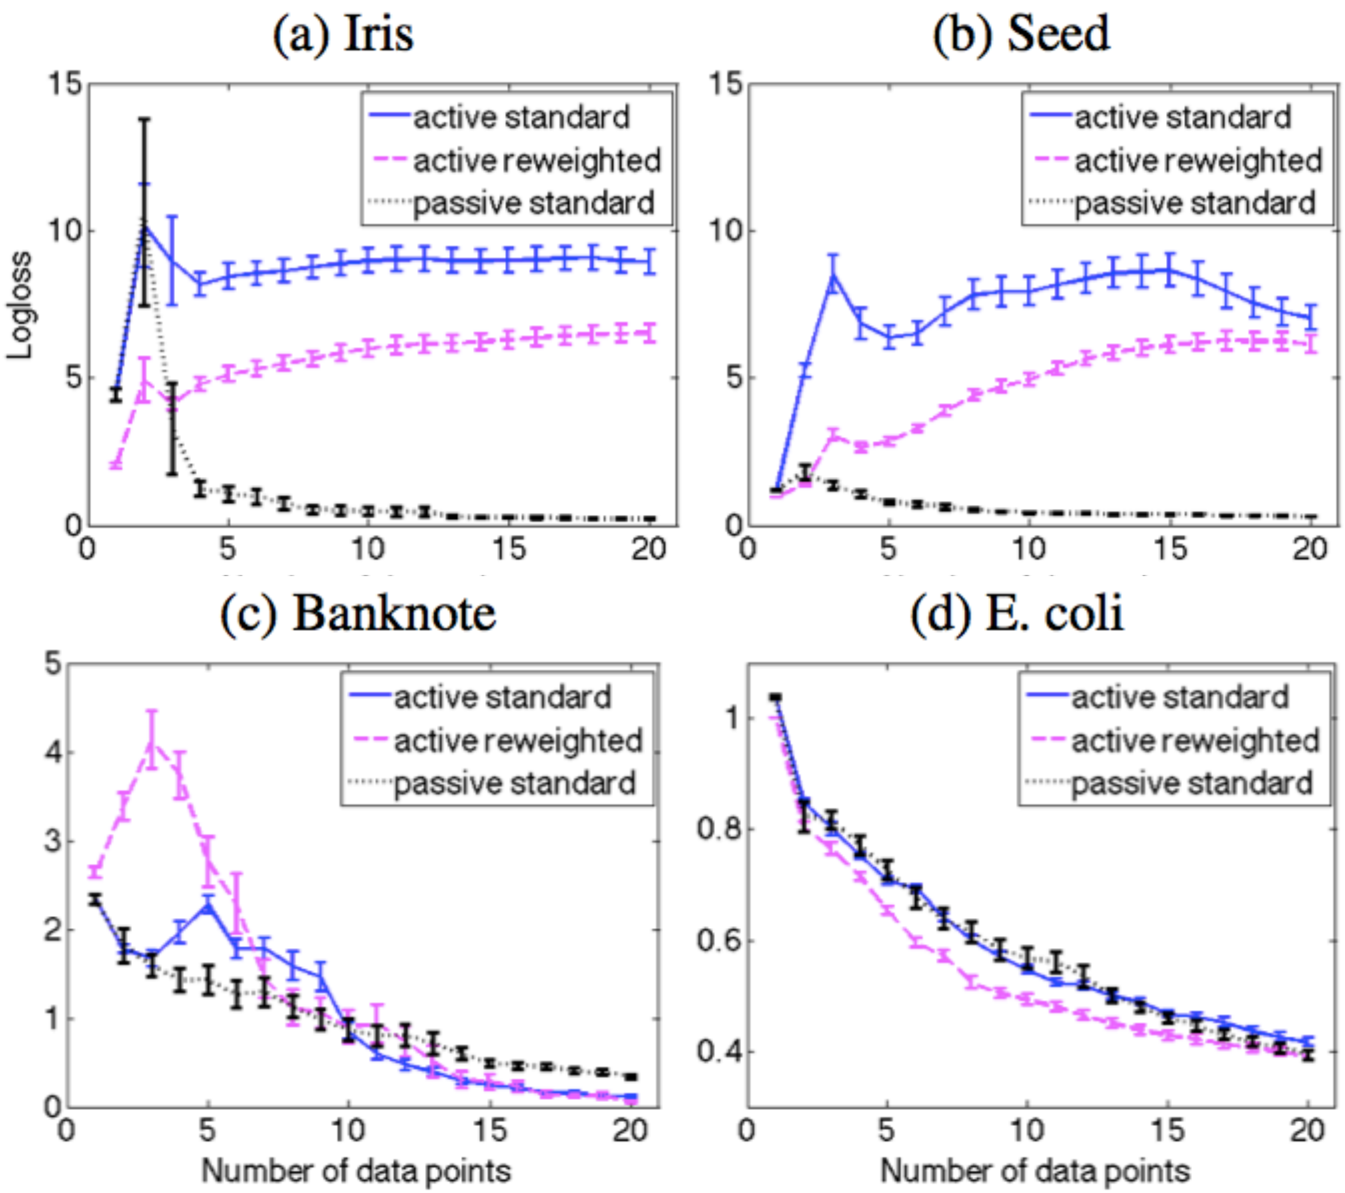

Approach

- We use the recently developed RBA (robust bias-aware prediction) framework for tackling covariate shift [Liu-Ziebart '14].
- RBA solves a game against a constrained adversary that chooses an evaluation distribution:

$$\min_{\hat{P}(y|x)} \max_{\check{P}(y|x) \in \tilde{\Xi}} \mathbb{E}_{P_{\mathcal{D}}(x)\check{P}(y|x)} \overbrace{[-\log \hat{P}(Y|X)]}^{\text{logarithmic loss}}$$

The set $\tilde{\Xi}$ constrains the adversary

Robust Prediction Strategy

- The RBA predictor can be obtained by solving the dual of a conditional max entropy estimation problem. [Liu-Ziebart '14]
- Can be shown to upper bound the the generalization loss, under some assumptions. [Grunwald-Dawid '04]
- $P_{src}(x)$ needs to be estimated – we use kernel density estimation with Gaussian kernels for $P_{src}(x)$.
- The RBA predictor turns out to less certain where the labeled data underrepresents the full data distribution.

Sampling Strategies

- active robust – select point with largest value conditioned entropy
- active random – select point at random
- active density – select point with highest density ratio of $P_D(x)/P_L(x)$

Standard Logistic Regression Models

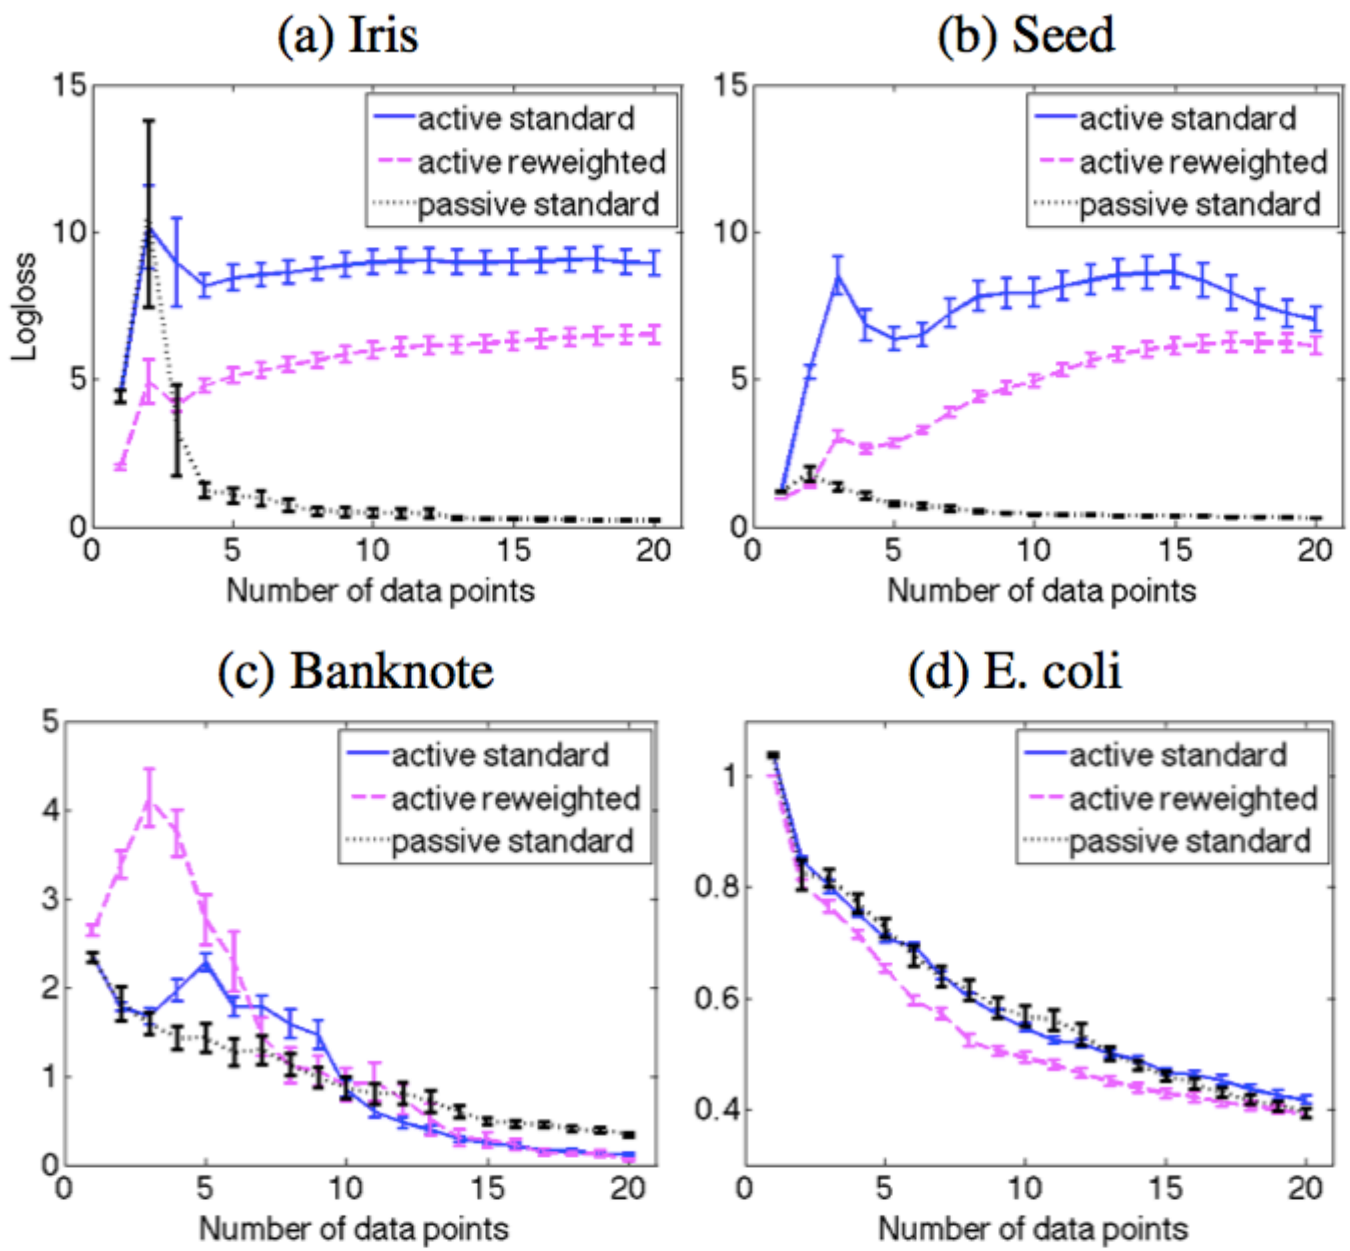

Our Results (logloss) [Liu-**R**-Ziebart '15]

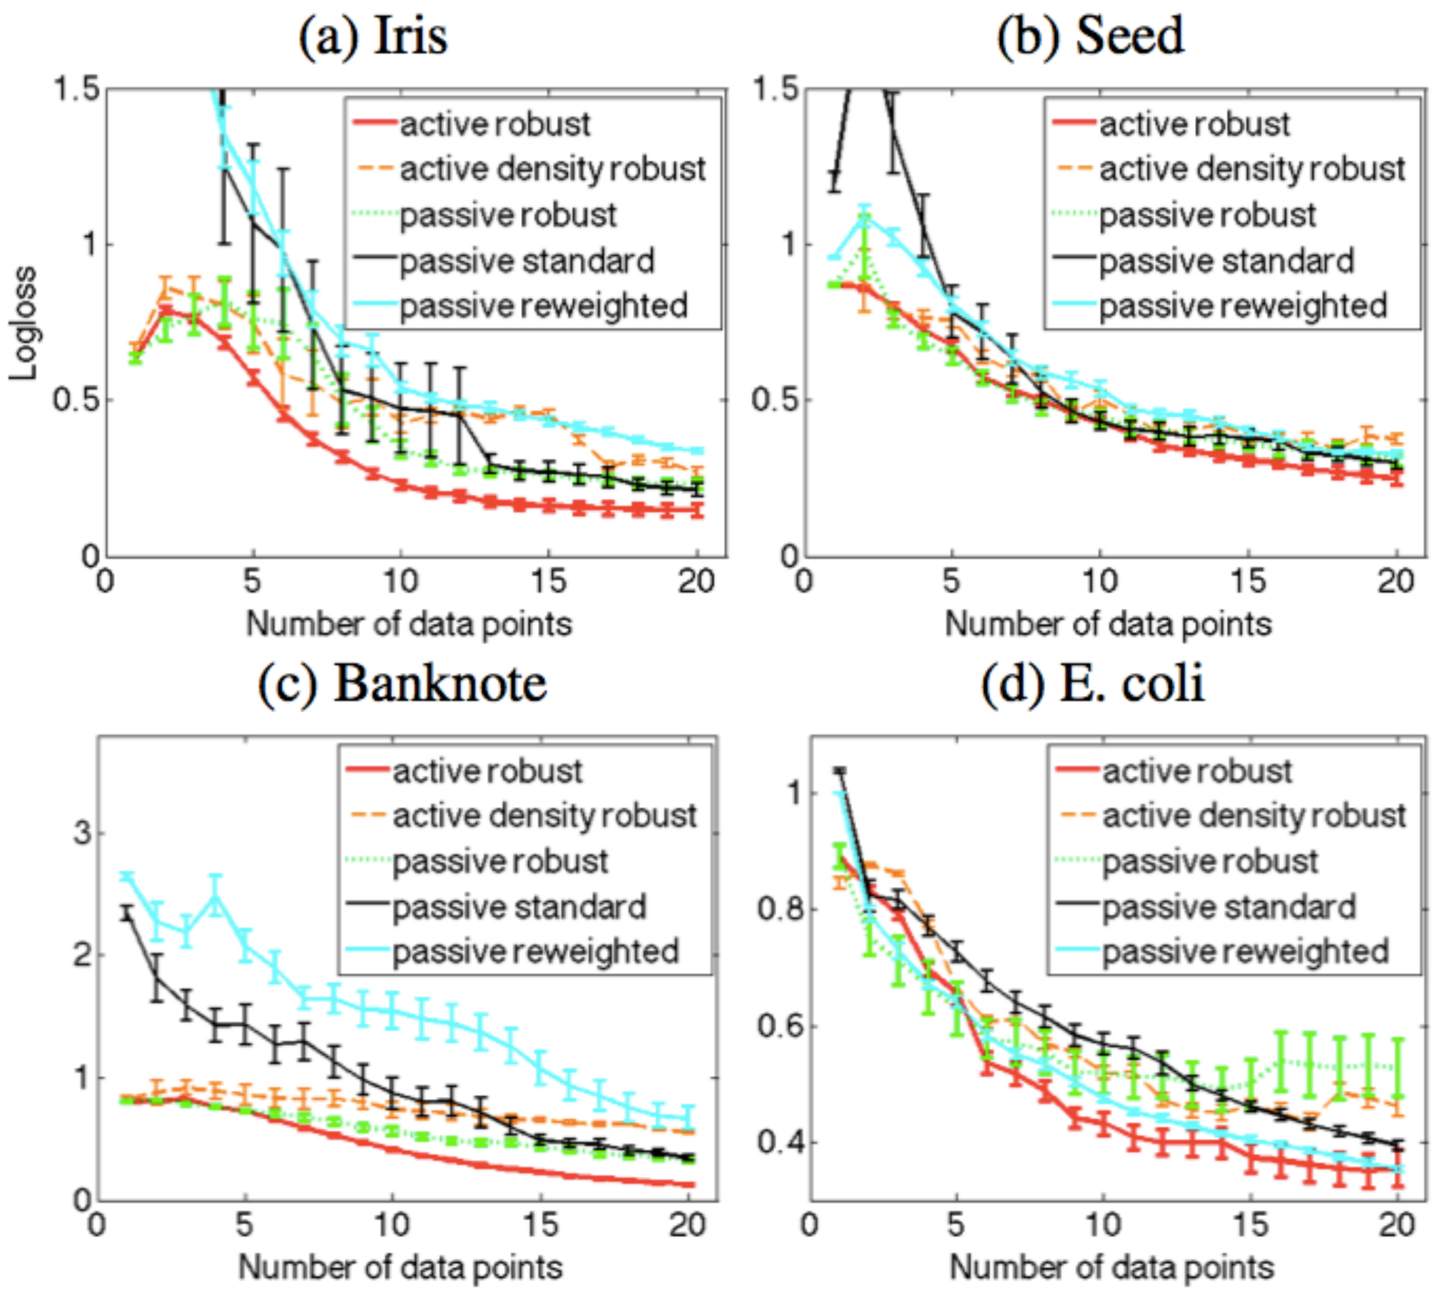

Our Results (error) [Liu-**R**-Ziebart '15]

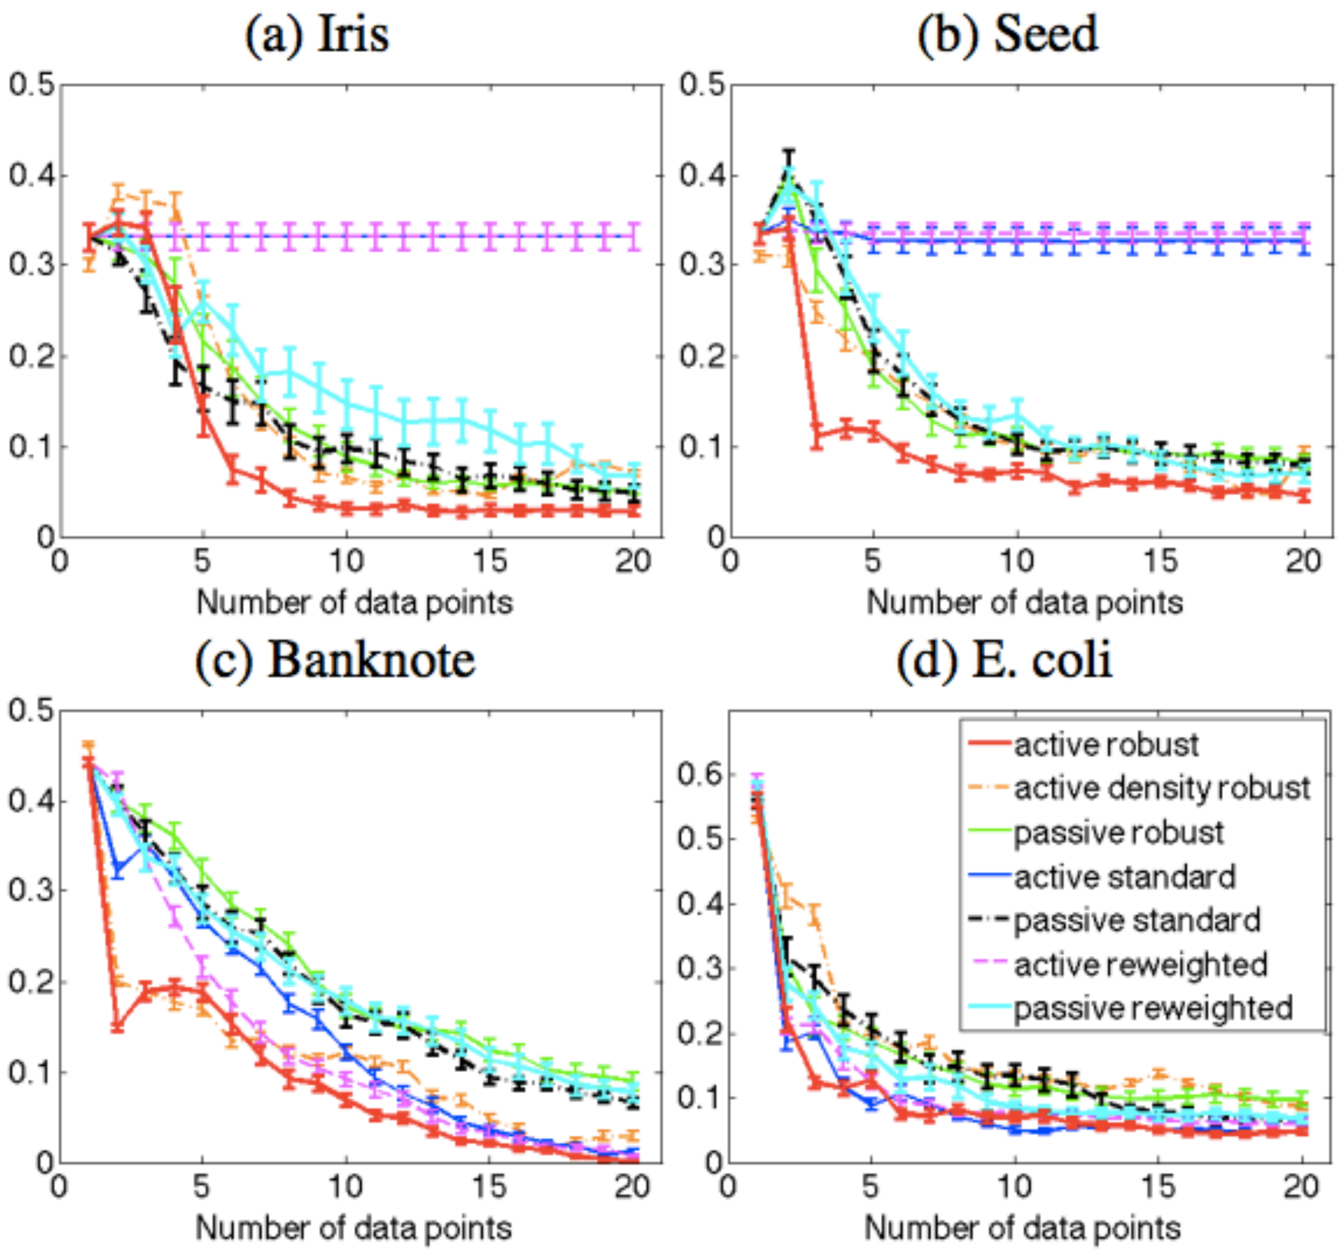

Predictions

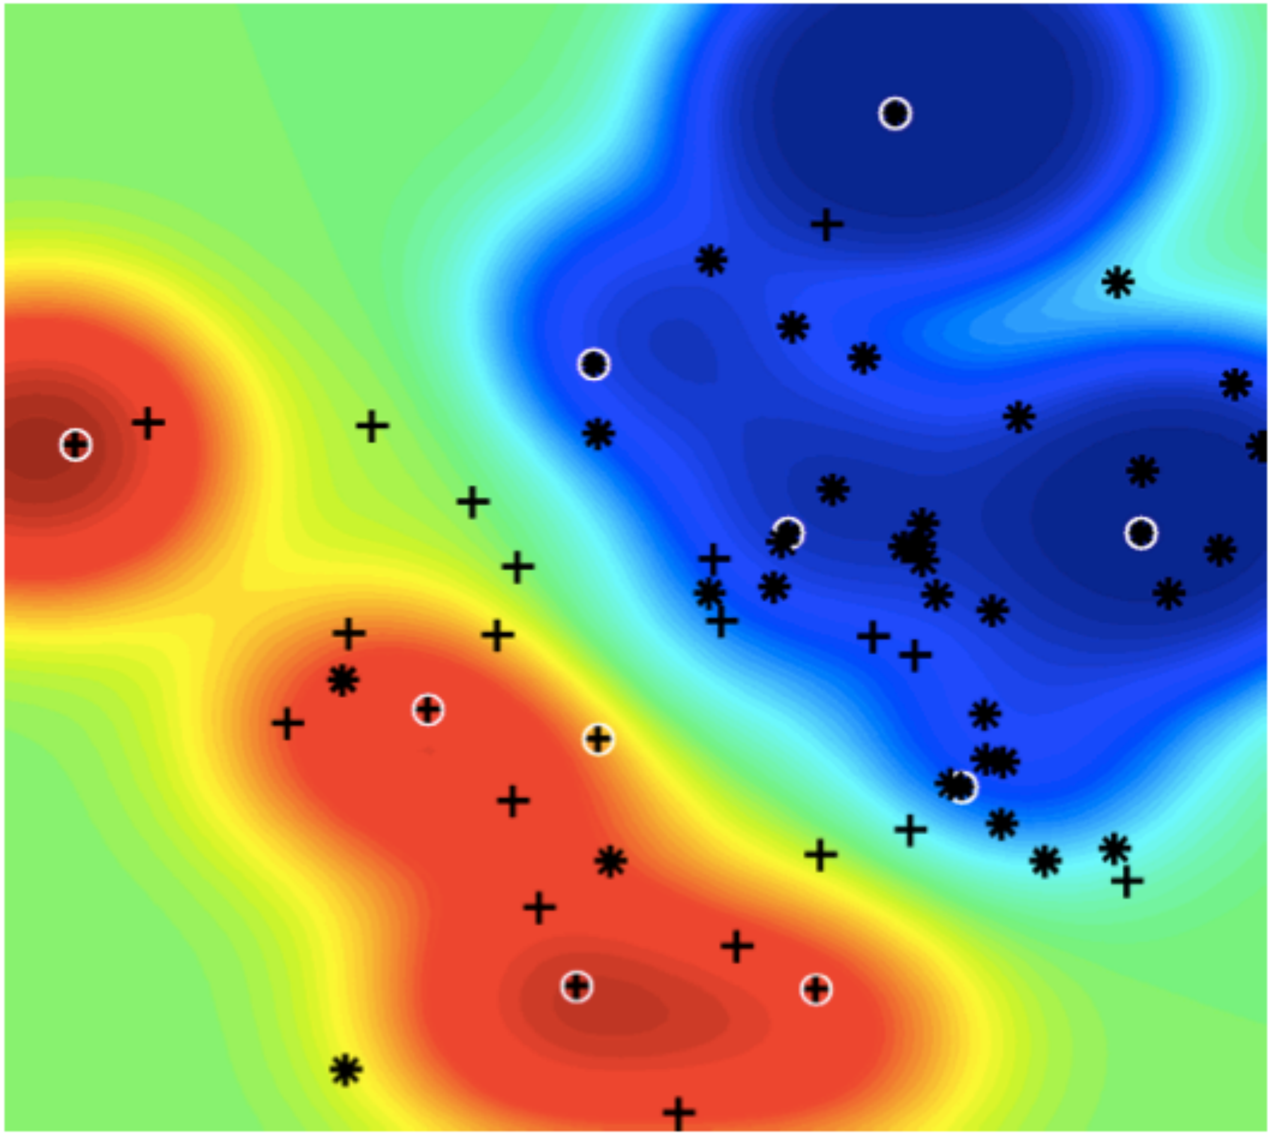

Discussion

Summary

- Gave first optimal and (theoretically) efficient algorithm for contextual bandits by reducing to supervised learning.
- Showed how to make any ensemble algorithm "feature efficient".
- Gave an a principled active learning algorithm with impressive empirical performance.

Observations

- Progress occurred in stages – eg. contextual algorithm was first made optimal but inefficient, then theoretically efficient, then practical (by others).
- All these problems originate from practical concerns.
- Each of these areas heavily used theory, but included practical elements, verified on large data. E.g. active learning's density estimation.

Open Problems

- A practical, optimal contextual bandit algorithm that works in adversarial setting.
- Better models for trading off error / prediction time.
- Pessimistic active learning applied directly to classification error.
- **In general**: *machine learning problems where principled new algorithms can tackle important application areas!*

The End

Thank you! Any Questions?