

BANDIT ALGORITHMS FOR INTERNET-SCALE APPLICATIONS

LEV REYZIN

UIC

CASSC



Lev Reyzin

0

+ Share



Google



Google Search

I'm Feeling Lucky


Google



Google Search

I'm Feeling Lucky

GOOGLE NEWS, AN EXAMPLE



Grouping Articles

News

U.S. edition ▾

Modern ▾



Top Stories

Top Stories



- Gwyneth Paltrow
- Beyoncé
- Justin Bieber
- Akamai Technologies
- United Parcel Service
- Verizon Wireless
- Safeway Inc.
- National Football League
- Harley-Davidson Motor Company
- Google



Philly.com

See realtime coverage



Live video: Bush library opening to draw four former presidents

Los Angeles Ti...
29 minutes ago



Written by Molly Hennessy...



DALLAS -- Members of George W. Bush's administration gathered at Southern Methodist University in anticipation of the dedication of the former president's library, a ceremony which President Obama and all four living former presidents are expected to ...

Featured: [So, are George W. Bush's paintings any good?](#)
Pegasus News

In Depth: [Presidential Library Funding Lacks Donor Disclosure Requirements](#)
Huffington Post

Related: [George W. Bush Presidential Center »](#)
[George Bush »](#)




CNN



CNN



CNN



NBCNe



NDTV

Body found in water that of missing Brown student

Boston.com - 29 minutes ago

PROVIDENCE, R.I. (AP) - The Rhode Island medical examiner's office says a body found in Providence is that of a Brown University student missing since last month.



New York ...

Explosions on River in Alabama Injure 3

New York Times - 32 minutes ago

Fires ignited by a series of at least seven explosions aboard two fuel barges on the Mobile River in Alabama burned themselves out Thursday morning, the authorities said.

Recent

Syria Claims Disruption of Rebel Supply Lines

New York Times - 18 minutes ago

Convergence of Presidents at Bush Library Dedication

New York Times - 24 minutes ago

'America took my kids away': Mother of Boston suspects insists sons not ...

NBCNews.com - 20 minutes ago

Weather for Hanover Park, IL

Today



54° 34°

Fri



66° 45°

Sat



64° 46°

Sun



68° 46°

The Weather Channel - Weather Underground - AccuWeather

Illinois » - Change location

Unity Weekend draws more than 530 volunteers for interfaith community service ...

Chicago Tribune - 3 minutes ago

What's new to eat at the ballparks?

Chicago Tribune - 3 minutes ago

Shea Soucie And Martin Homer, Chicago-Based Interior Designers, Tell Us ...

Huffington Post - 2 hours ago

Editors' Picks

Illinois

small plane crash

Recommended

U.S.

World

Sci/Tech

Health

Business

More Top Stories

Spotlight


Entertainment

Sports

5 presidents on hand for Bush Center dedication

CNN - 13 minutes ago   

Dallas (CNN) – Some may disagree whether George W. Bush was a "uniter, not a divider," as he liked to say, but he is expected to get all five living presidents together for the dedication of his presidential library. Five presidents – three Democrats and two ...



CNN

CNN

6:38a The following heads of state will be in attendance at the George W. Bush ...

KHOU - 3 minutes ago

George W. Bush library dedication ceremony begins in Dallas; 5 presidents ...

Minneapolis Star Tribune - 3 minutes ago

5 Presidents Appear at Bush Library...

ABC News - 2 minutes ago

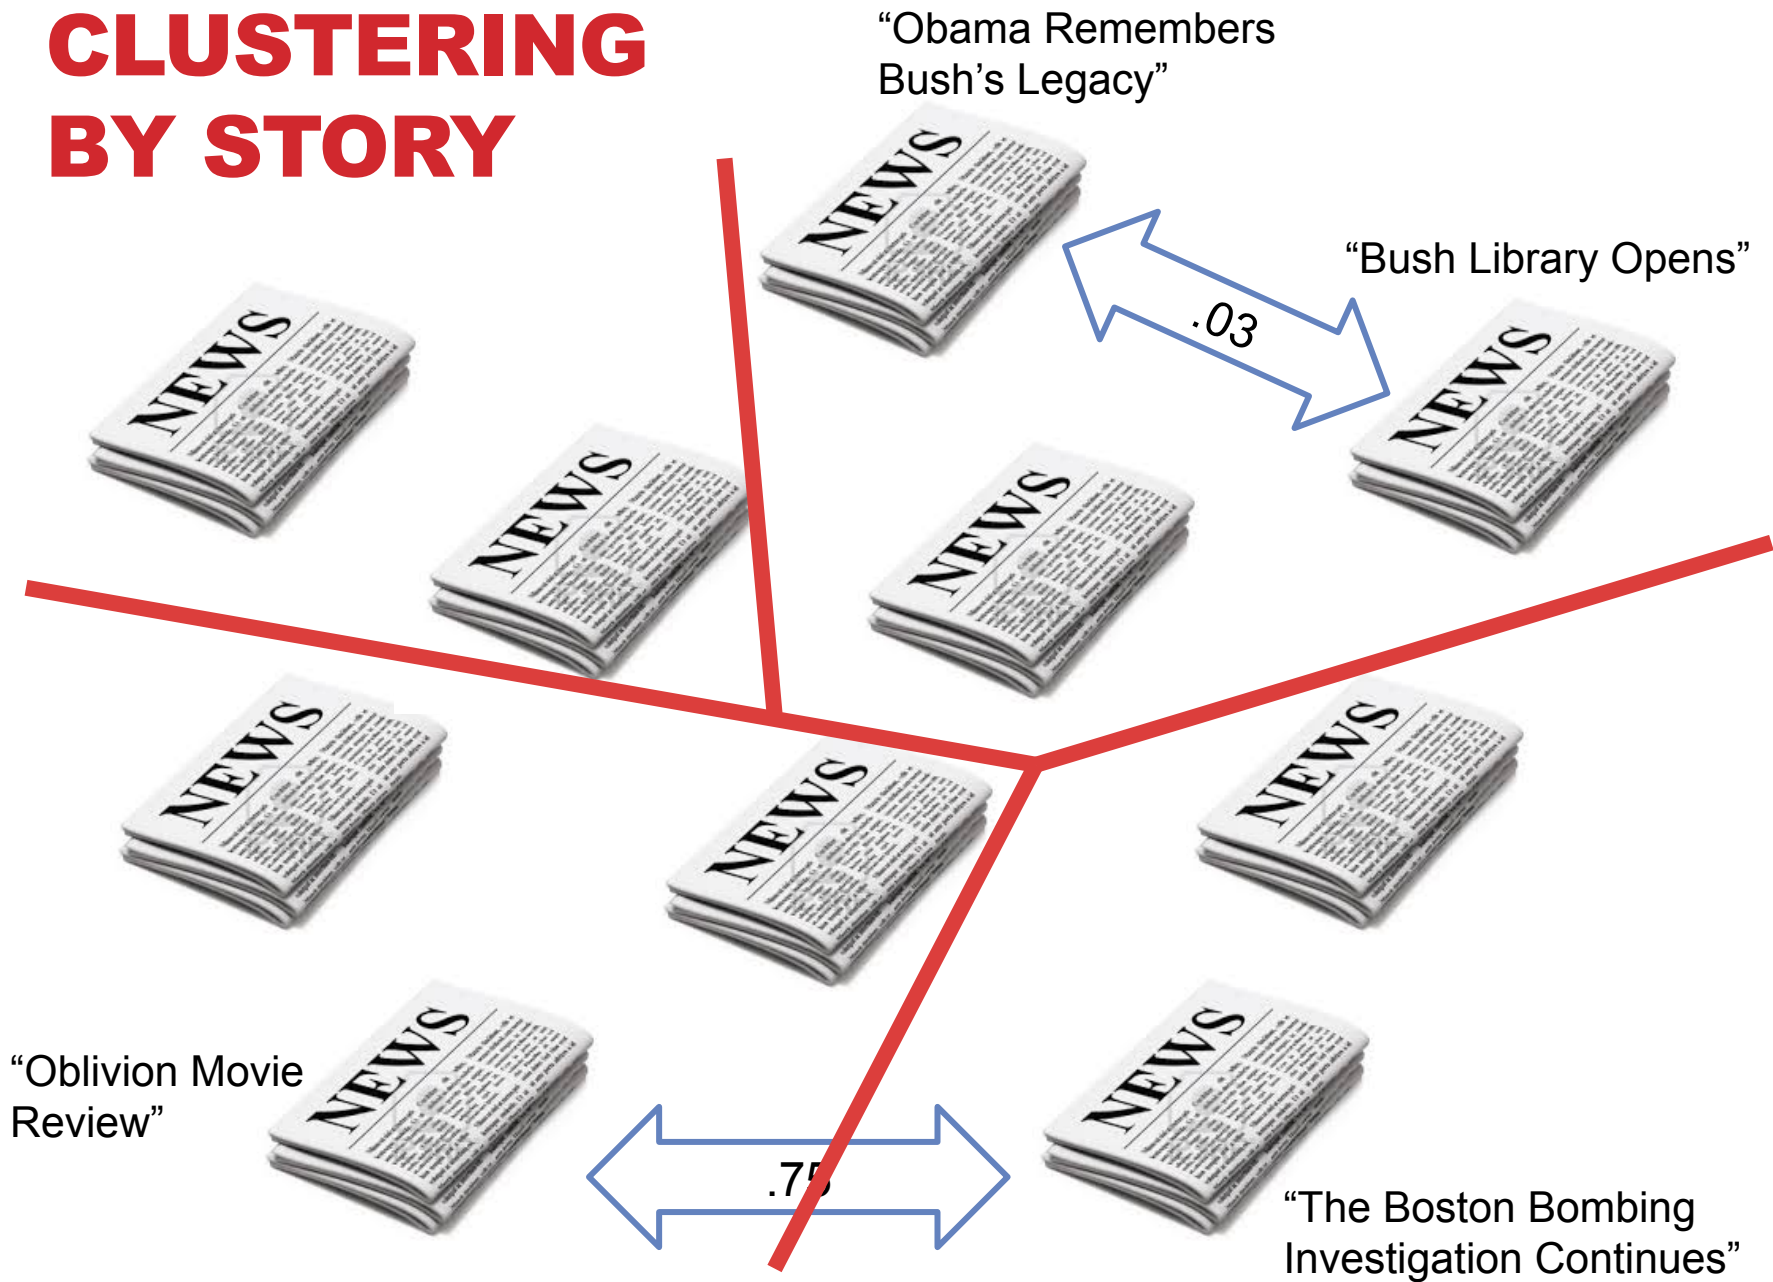
Presidents past and present converge in Texas, saluting Bush at library dedication

Montreal Gazette - 2 minutes ago

Blog: Five living presidents announced at Bush Center ceremony

WFAA - 4 minutes ago

CLUSTERING BY STORY



CLUSTERING

1. How to choose a distance function?

one popular one is cosine similarity = $a \cdot b / (\|a\| \|b\|)$

2. What is the right number of clusters?

3. What should the clustering objective function be?

k-means?, k-median?, min-sum?, max-cut?,...

4. How do you optimize the objective?

all of the reasonable ones are NP-Hard!

5. How to “fade out” old clusters?

6. How to evaluate how you’re doing?

Putting the Articles into Categories

News

U.S. edition ▾

Modern ▾



Top Stories

- Gwyneth Paltrow
- Beyoncé
- Justin Bieber
- Akamai Technologies
- United Parcel Service
- Verizon Wireless
- Safeway Inc.
- National Football League
- Harley-Davidson Motor Company
- Google

Illinois

small plane crash

Recommended

U.S.

World

Sci/Tech

Health

Business

More Top Stories

Spotlight

Entertainment

Sports

Top Stories



Philly.com

See
realtime
coverage

Live video: Bush library opening to draw four former presidents

Los Angeles Ti...
29 minutes ago



Written by
Molly Hennessy...



DALLAS -- Members of George W. Bush's administration gathered at Southern Methodist University in anticipation of the dedication of the former president's library, a ceremony which President Obama and all four living former presidents are expected to ...


Featured: [So, are George W. Bush's paintings any good?](#)
Pegasus News

In Depth: [Presidential Library Funding Lacks Donor Disclosure Requirements](#) Huffington Post

Related: [George W. Bush Presidential Center >](#)
[George Bush >](#)




CNN



CNN



CNN



NBCNe



NDTV

Body found in water that of missing Brown student

Boston.com - 29 minutes ago

PROVIDENCE, R.I. (AP) - The Rhode Island medical examiner's office says a body found in Providence is that of a Brown University student missing since last month.



New York ...

Explosions on River in Alabama Injure 3

New York Times - 32 minutes ago

Fires ignited by a series of at least seven explosions aboard two fuel barges on the Mobile River in Alabama burned themselves out Thursday morning, the authorities said.

Recent

Syria Claims Disruption of Rebel Supply Lines

New York Times - 18 minutes ago

Convergence of Presidents at Bush Library Dedication

New York Times - 24 minutes ago

'America took my kids away': Mother of Boston suspects insists sons not ...

NBCNews.com - 20 minutes ago

Weather for Hanover Park, IL

Today



54° 34°

Fri



66° 45°

Sat



64° 46°

Sun



68° 46°

The Weather Channel - Weather Underground - AccuWeather

Illinois > - [Change location](#)

Unity Weekend draws more than 530 volunteers for interfaith community service ...

Chicago Tribune - 3 minutes ago

[What's new to eat at the ballparks?](#)

Chicago Tribune - 3 minutes ago

[Shea Soucie And Martin Horner, Chicago-Based Interior Designers, Tell Us ...](#)

Huffington Post - 2 hours ago

Editors' Picks

Putting the Articles into Categories

News

U.S. edition ▾

Modern ▾



Top Stories

News near you

small plane crash

Recommended

U.S.

World

Sci/Tech

Health

Business

More Top Stories

Spotlight

Entertainment

Sports

Tyrann Mathieu
Dale Earnhardt, Jr.
New England Patriots
Petra Kvitová
New York Yankees
Brooklyn Nets
Lance Armstrong
Barry Sanders
San Francisco 49ers
Erik Karlsson

Sports



MassLive....

QB Ryan Mallett in play for trade

Boston.com
1 hour ago



Written by
Greg Bedard



Multiple teams have called the Patriots this week about a potential trade for backup quarterback Ryan Mallett, according to three league sources.

See
realtime
coverage

Highly Cited: Source: Browns have inquired about Ryan Mallett again, Pats want second ... Boston Herald (blog)

Related: [New England Patriots](#) » [Ryan Mallett](#) »




Newsday

No post-NFL draft party tonight for LSU's Tyrann Mathieu after all

Los Angeles Times - 15 minutes ago

Tyrann Mathieu is not having a post-draft party for himself Thursday night. I don't care what you think you heard, it's not true.



Northwest ...

Last-minute Lions prognostications as NFL draft looms

The Detroit News - 1 hour ago

With just hours before the 2013 NFL draft, pundits are still scratching their heads about where most prospects will land. The consensus is the top two players in the draft are offensive tackles Luke Joeckel of Texas A&M and Eric Fisher of Central Michigan, but ...



The Hock...

Jamie Benn: Dallas Stars is must-win mode

Rotoworld.com - 2 hours ago

The Dallas Stars need to beat Columbus on Thursday to stay in the hunt for a playoff spot. "Desperation, do or die," said Jamie Benn of the Stars' mindset.



Dwight Howard can't hold it all together for Lakers

Los Angeles Times - 10 hours ago

Popular video



This Is SportsCenter - Clayton Kershaw
ESPN

Spotlight

[Mark Sanchez questions Jets' trade for Tim Tebow - USA Today](#)

USA TODAY - Apr 19, 2013

[Board play - Cincinnati Bengals](#)

Bengals.com - Apr 21, 2013

[Justice Dept. lays out suit against Lance Armstrong, says he was 'unjustly enriched'](#)

Fox News - Apr 23, 2013


[David Ortiz Speech: Pre-Game Ceremony At Fenway Punctuated By F-Bomb ...](#)

Huffington Post - Apr 20, 2013

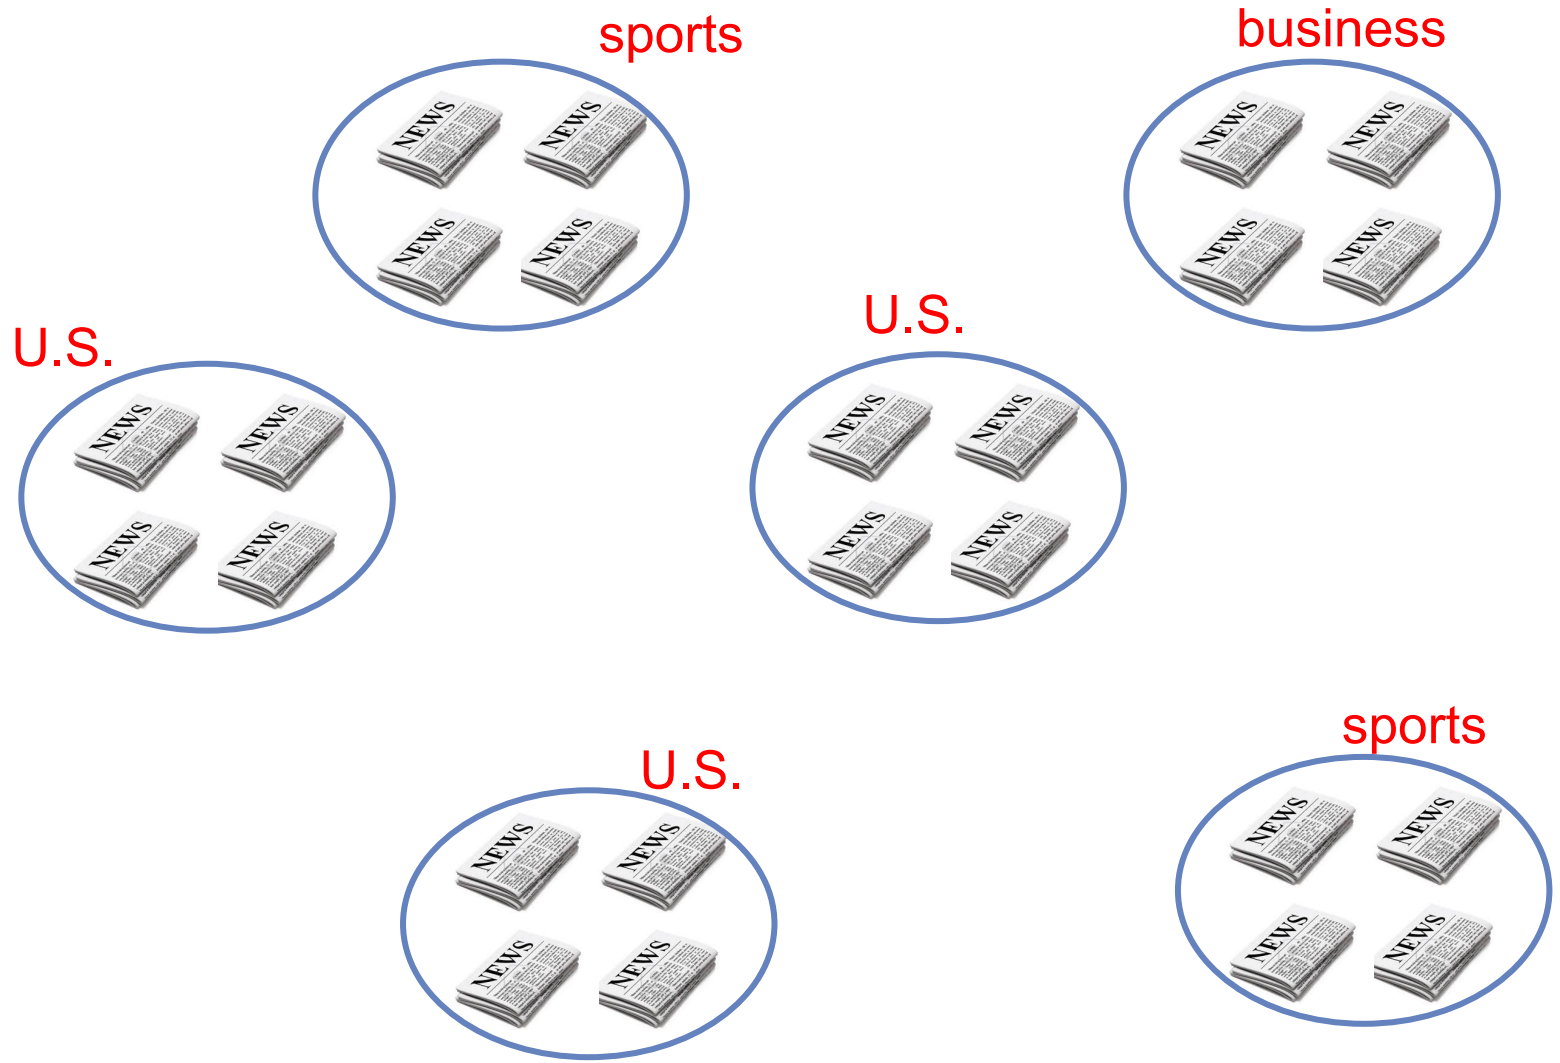
[Manti Te'o on the Hoax and Life After: 'Honestly, I'm Never Going to Be ...](#)

Vanity Fair - 1 hour ago

MULTICLASS CLASSIFICATION



MULTICLASS CLASSIFICATION



MULTICLASS CLASSIFICATION

How do you “classify” news articles into proper categories?

Thinking hard about what makes an article sports vs business and then programming it in won't work!

Main idea: use machine learning!

What does that mean? Give a computer many correctly classified articles and use a *generic algorithm* to make predictions from this data.

MACHINE LEARNING EXAMPLE

sports



business



sports



business



politics



politics



MACHINE LEARNING EXAMPLE

sports



business



sports



business



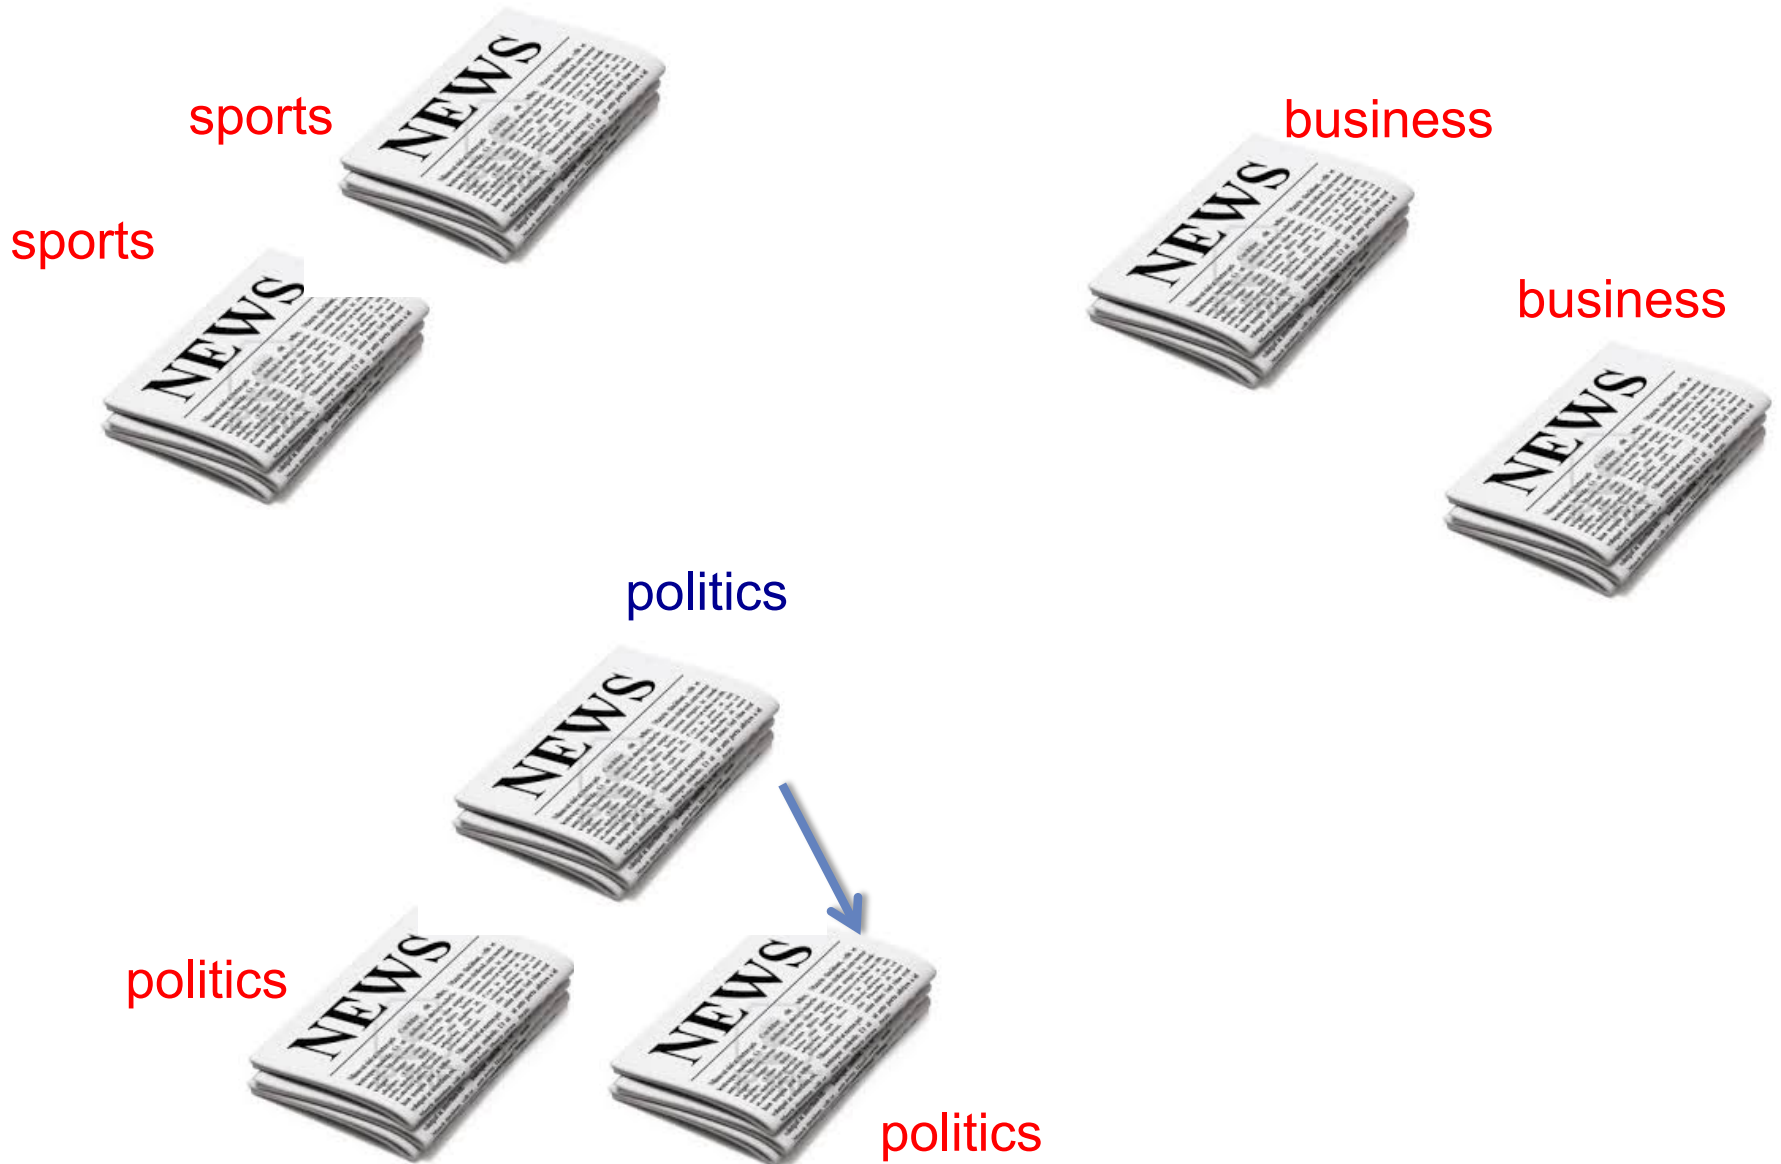
politics



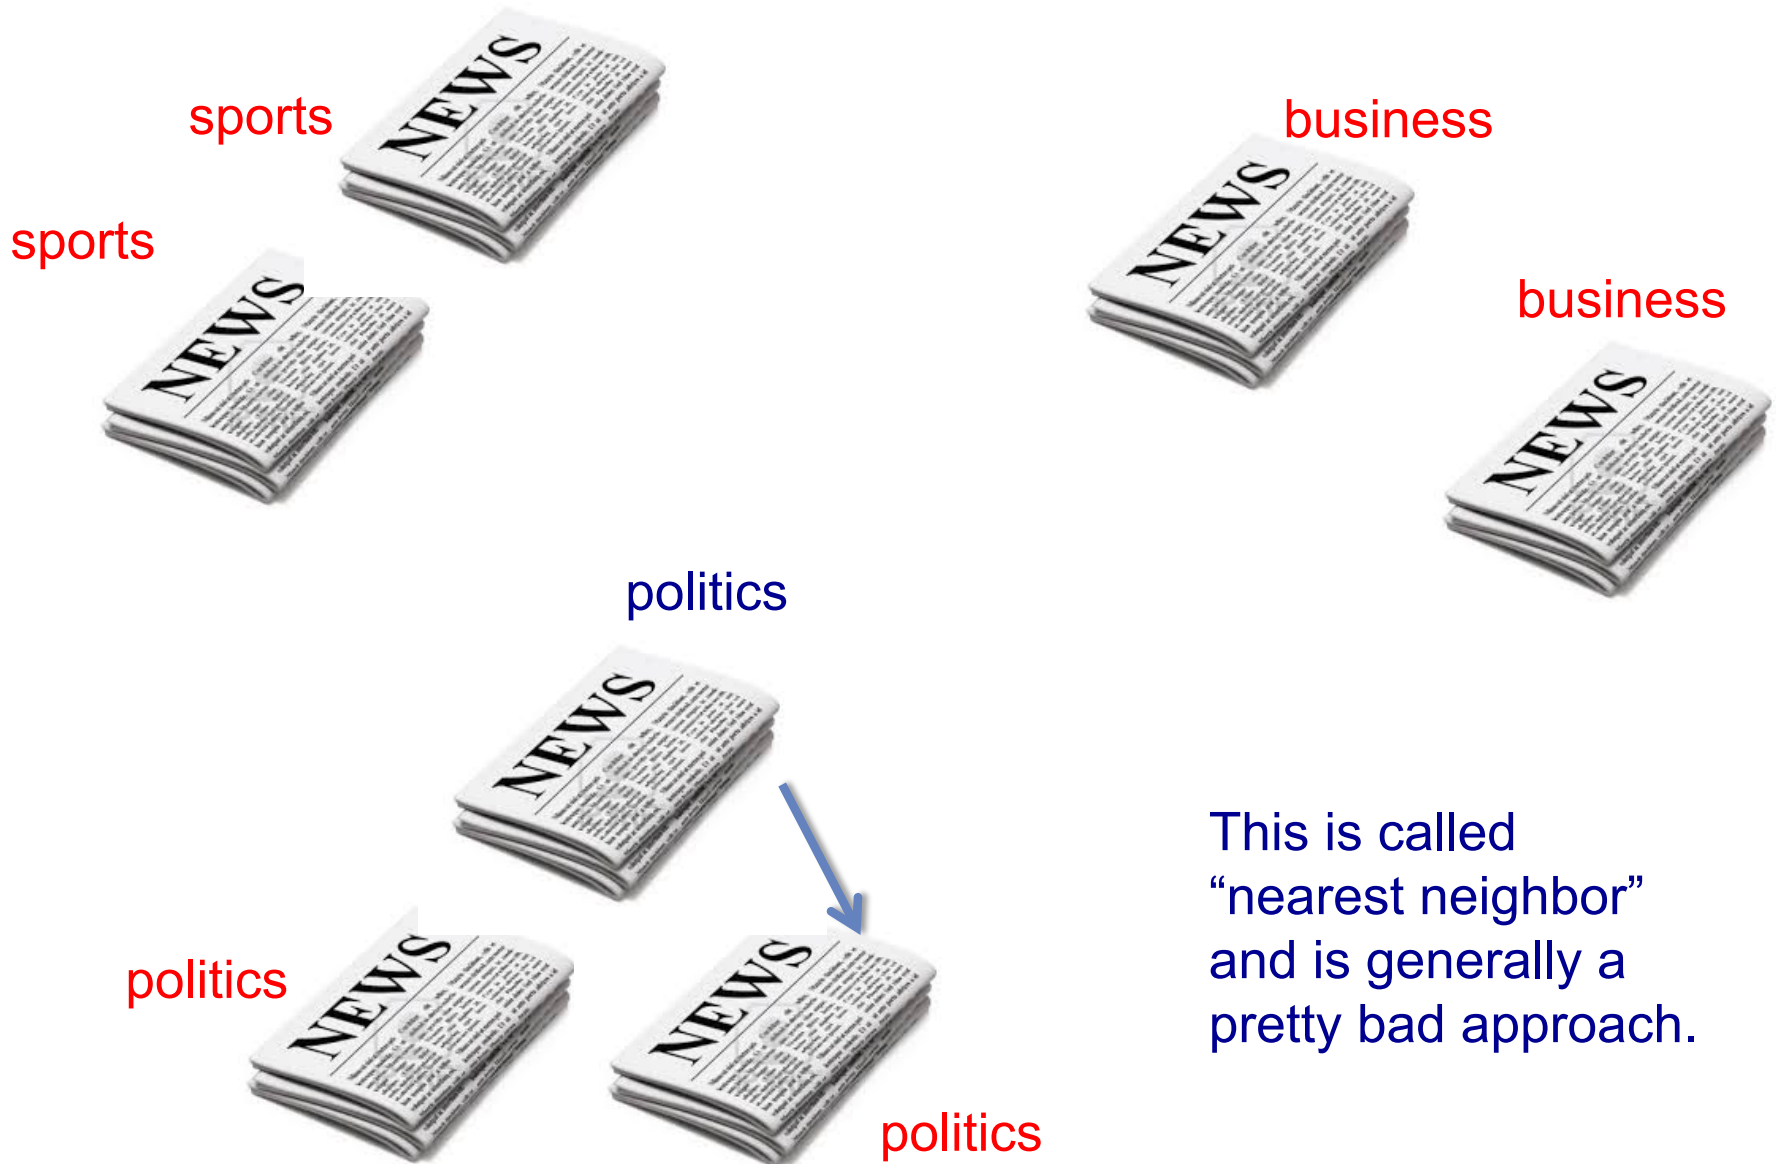
politics



MACHINE LEARNING EXAMPLE



MACHINE LEARNING EXAMPLE



This is called
“nearest neighbor”
and is generally a
pretty bad approach.

MULTICLASS CLASSIFICATION

PAC learning: *Given training data (articles) and their labels (sports/world/US/...) find a function that labels new data (articles) arbitrarily accurately without using too much labeled data (polynomial # of examples).*

Lots of work in this area. Unfortunately, the real world isn't as simple.

usual assumption: new data comes from same distribution as training data

False here. E.g. Chrome was a color before the browser came out.

What's the right assumption?

questions: How does one get labeled data? (pay people!)

Should the data that is labeled be labeled randomly? (no!) How?

How do you know when you're making mistakes?

Also this process is "online."



Top Stories

- Gwyneth Paltrow
- Beyoncé
- Justin Bieber
- Akamai Technologies
- United Parcel Service
- Verizon Wireless
- Safeway Inc.
- National Football League
- Harley-Davidson Motor Company
- Google

- Illinois
- small plane crash

Recommended

- U.S.
- World
- Sci/Tech
- Health
- Business
- More Top Stories
- Spotlight
- Entertainment
- Sports

Top Stories




See realtime coverage

Live video: Bush library opening to draw four former presidents

Los Angeles Ti... 29 minutes ago
 Written by Molly Henness...
 DALLAS — Members of George W. Bush's administration gathered at Southern Methodist University in anticipation of the dedication of the former president's library, a ceremony which President Obama and all four living former presidents are expected to ...

Featured: [So, are George W. Bush's paintings any good?](#)
 Pegasus News
 In Depth: [Presidential Library Funding Lacks Donor Disclosure Requirements](#) Huffington Post

Related: [George W. Bush Presidential Center](#) »
[George Bush](#) »



Body found in water that of missing Brown student

Boston.com - 29 minutes ago
 PROVIDENCE, R.I. (AP) - The Rhode Island medical examiner's office says a body found in Providence is that of a Brown University student missing since last month.



Explosions on River in Alabama Injure 3

New York Times - 32 minutes ago
 Fires ignited by a series of at least seven explosions aboard two fuel barges on the Mobile River in Alabama burned themselves out Thursday morning, the authorities said.

Recent

- Syria Claims Disruption of Rebel Supply Lines**
New York Times - 18 minutes ago
- Convergence of Presidents at Bush Library Dedication**
New York Times - 24 minutes ago
- 'America took my kids away': Mother of Boston suspects insists sons not ...**
NBCNews.com - 20 minutes ago

Weather for Hanover Park, IL

Today	Fri	Sat	Sun
54° 34°	66° 45°	64° 46°	68° 46°

The Weather Channel - Weather Underground - AccuWeather

Illinois » - Change location

- Unity Weekend draws more than 530 volunteers for interfaith community service ...**
Chicago Tribune - 3 minutes ago
- What's new to eat at the ballparks?**
Chicago Tribune - 3 minutes ago
- Shea Soucie And Martin Homer, Chicago-Based Interior Designers, Tell Us ...**
Huffington Post - 2 hours ago

Editors' Picks

Library dedication offers glimpse of tenuous Obama-Bush relationship

Washington Post - 21 minutes ago

After the killing of Osama bin Laden in May 2011, President Obama invited former president George W. Bush to join him at Ground Zero for a ceremony marking the event. Bush declined, saying through a spokesman that he wanted to stay out of the spotlight.



ABC News

CNN

Politics on hold at the dedication of Bush library

Exponent-telegram - 6 minutes ago

Clinton, Carter note Bush's humanitarian efforts - NECN.com

NECN - 7 minutes ago

Obama, former presidents rally around George W. Bush as library opens

Reuters India - 20 minutes ago

George W. Bush Tears: See Former President Give Emotional Speech At Bush ...

iDigitalTimes.com - 26 minutes ago

Bush ends library dedication with emotional speech - NECN.com

NECN - 35 minutes ago

Presidents praise George W. Bush at new library - USA Today

USA TODAY - 46 minutes ago

Obama Honors Bush at 'Texas-Sized' Library Dedication

ABC News - 57 minutes ago

Bush library dedication both a look back and a prologue

Washington Post - 1 hour ago

WHAT TO DISPLAY?

This looks like a “ranking problem,” but a more important issue lies underneath.

What should be displayed?

nothing offensive

articles from diverse sources

articles from trusted sources

articles people are likely to read

- Mail
- News
- Finance
- Sports
- Movies
- omg!
- Shine
- Autos
- Shopping
- Travel
- Dating
- Jobs

More Y! Sites >

Make **YAHOO!** your homepage



Netflix-Try for free. Instantly watch TV shows and movies on your TV or computer. Get 1 month free.



Rare Yosemite spectacle caught on film

A photographer snaps a once-in-a-lifetime image of a bolt of lightning cracking through a rainbow. [See this picture »](#)

--	--	--	--	--

All Stories News Local Entertainment Sports More v



Rod Stewart Plays Intimate L.A. Club Gig for Diehard Fans, Harry Styles

On May 7, rock icon Rod Stewart, who experienced a career rebirth in the 1990s and 2000s with his multi-volume Great American Songbook Maximum Performance (NEW)



Beyonce and Blue Ivy Step Out in Paris, Gwyneth Paltrow Wears a Sheer Dress and No Underwear: Today's Top Stories

Take a look at Us Weekly's most-read stories from Thursday, April 25 Us Weekly

Cycling-UCI go on the attack after latest accusations by USADA

Trending Now

Watch the show »

- | | |
|---------------------------|--------------------------|
| 1 Jessica Alba double... | 6 Psy knocked off |
| 2 Ace of Base Nazi past | 7 Danica Patrick divo... |
| 3 38 die in mental hos... | 8 T-Mobile 'deceptive' |
| 4 Jennifer Love Hewitt | 9 HIV vaccine fails |
| 5 Edward Norton fight | 10 Hyundai suicide ad |

EVONY
FREE FOREVER
PLAY NOW

Ad Feedback

AdChoices [D]

Mount Laurel

48°F Fair



Today 65° 46°



Tomorrow 60° 44°



Sunday 60° 44°

- Mail
- News
- Finance
- Sports
- Movies
- omg!
- Shine
- Autos
- Shopping
- Travel
- Dating
- Jobs

More Y! Sites >

Make YAHOO! your homepage



Netflix-Try for free. Instantly watch TV shows and movies on your TV or computer. Get 1 month free.



Rare Yosemite spectacle caught on film

A photographer snaps a once-in-a-lifetime image of a bolt of lightning cracking through a rainbow. [See this picture »](#)

6 - 10 of 80

--	--	--	--	--

All Stories News Local Entertainment Sports More v



Rod Stewart Plays Intimate L.A. Club Gig for Diehard Fans, Harry Styles

On May 7, rock icon Rod Stewart, who experienced a career rebirth in the 1990s and 2000s with his multi-volume Great American Songbook Maximum Performance (NEW)



Beyonce and Blue Ivy Step Out in Paris, Gwyneth Paltrow Wears a Sheer Dress and No Underwear: Today's Top Stories

Take a look at Us Weekly's most-read stories from Thursday, April 25 Us Weekly

Cycling-UCI go on the attack after latest accusations by USADA

Trending Now

Watch the show »

- | | |
|---------------------------|--------------------------|
| 1 Jessica Alba double... | 6 Psy knocked off |
| 2 Ace of Base Nazi past | 7 Danica Patrick divo... |
| 3 38 die in mental hos... | 8 T-Mobile 'deceptive' |
| 4 Jennifer Love Hewitt | 9 HIV vaccine fails |
| 5 Edward Norton fight | 10 Hyundai suicide ad |

FREE FOREVER

PLAY NOW

Ad Feedback

AdChoices

Mount Laurel

48°F Fair



Today 65° 40°



Tomorrow 60° 44°



Sunday 60° 54°

Watch the Internet of Everything in action. 



●●● LATEST HEADLINES

[MORE HEADLINES »](#)


7:48 AM  [Shiseido Sees Profit Next Year](#)

First-Quarter Growth, at 2.5%, Misses Expectations

The nation's gross domestic product, a measure of all goods and services produced in the economy, advanced at a 2.5% annual rate between January and March, the Commerce Department said. [17 min ago](#)

U.S. GDP

Quarterly change at an annualized rate, adjusted for inflation




Markets >

Sponsored by: **Vanguard**

[Overview](#) | [U.S.](#) | [Europe](#) | [Asia](#) | [FX](#) | [Rates](#) | [Futures](#)

	LAST	CHG	%CHG	RANGE: 1 DAY
DJIA	14700.80	+24.50	0.17%	14673 - 14731
Nasdaq	3289.99	+20.33	0.62%	3269.66 - 3290.32
FTSE 100	6401.63	-40.96	0.64%	6442.59 - 6360.67
Nikkei 225	13884.13	-41.95	0.30%	13926.08 - 13842.18
Crude Oil	93.50	-0.14	0.15%	93.64 - 93.36
Gold	1470.50	+8.50	0.58%	1462.00 - 1479.00

4/26/13 8:44 AM EDT 


[Market Data](#) | [MoneyBeat](#) | [Watchlist](#) | [Portfolio](#) | [Customize](#)

Disappointing GDP Data Weigh on Futures

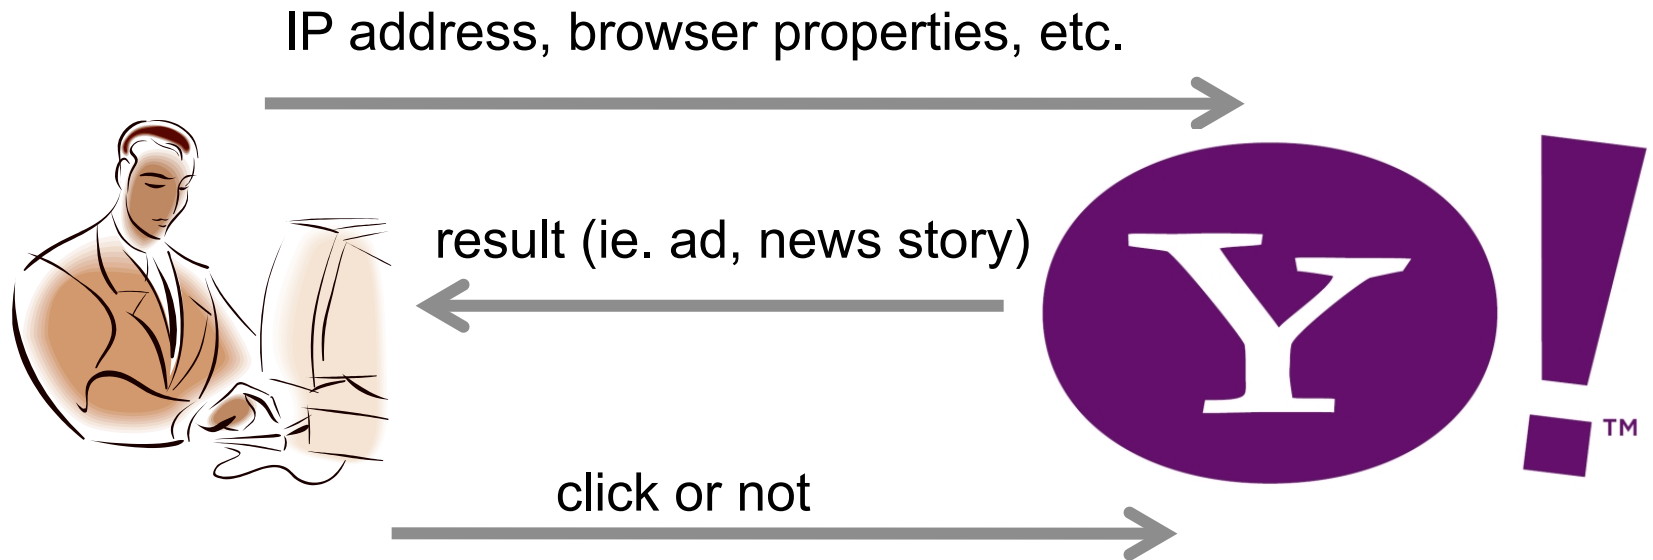
[Mansion: From Oil, a Land Rush](#)

[Einstein Proved Right--Again](#)

[NFL Draft](#)



SERVING CONTENT TO USERS



SERVING CONTENT TO USERS



IP address, browser properties, etc.



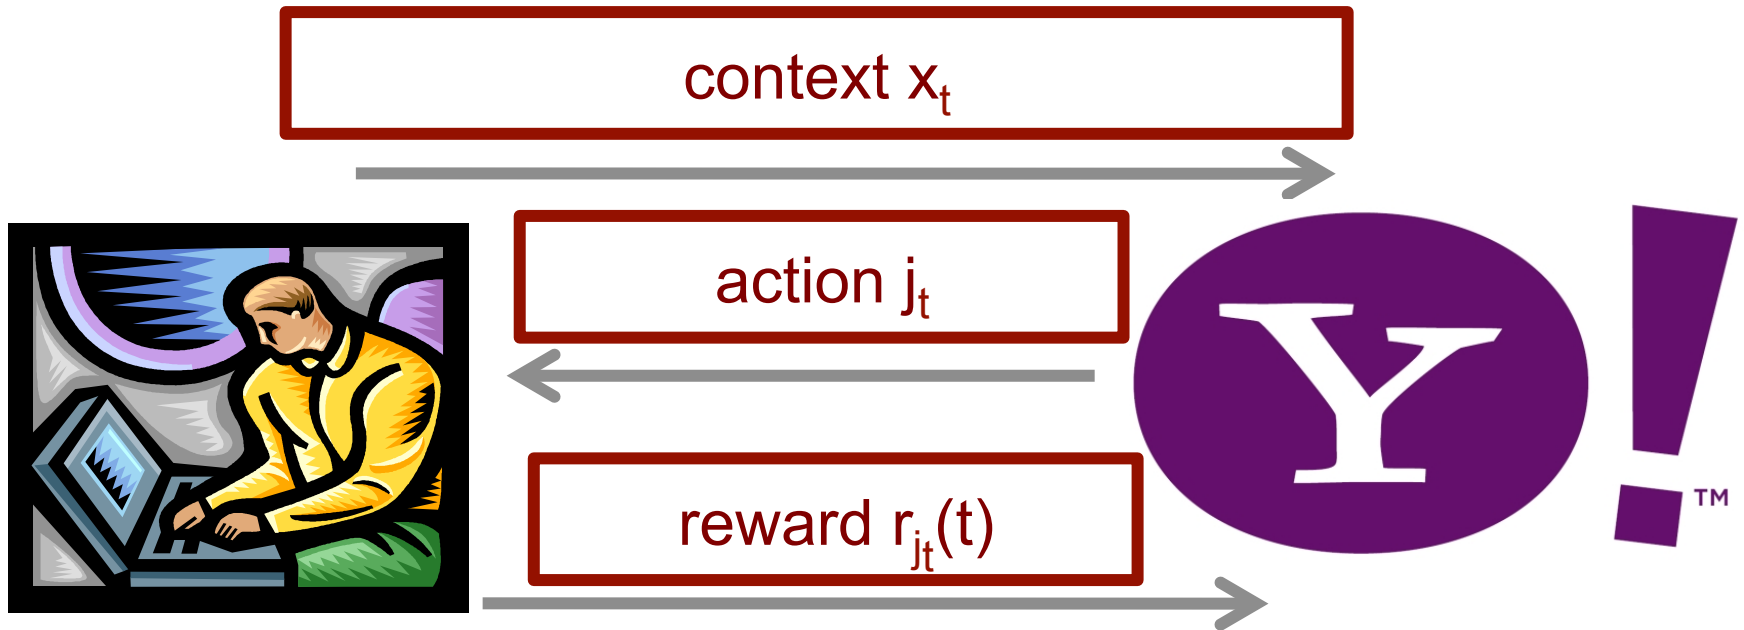
result (ie. ad, news story)



click or not




SERVING CONTENT TO USERS



MULTIARMED BANDITS

[ROBBINS '52]



1



2



3




⋮

k



MULTIARMED BANDITS

[ROBBINS '52]



click





⋮



MULTIARMED BANDITS

[ROBBINS '52]



click




click

⋮



MULTIARMED BANDITS

[ROBBINS '52]



click

no





click

⋮



MULTIARMED BANDITS

[ROBBINS '52]



click



click

no



click

⋮




click



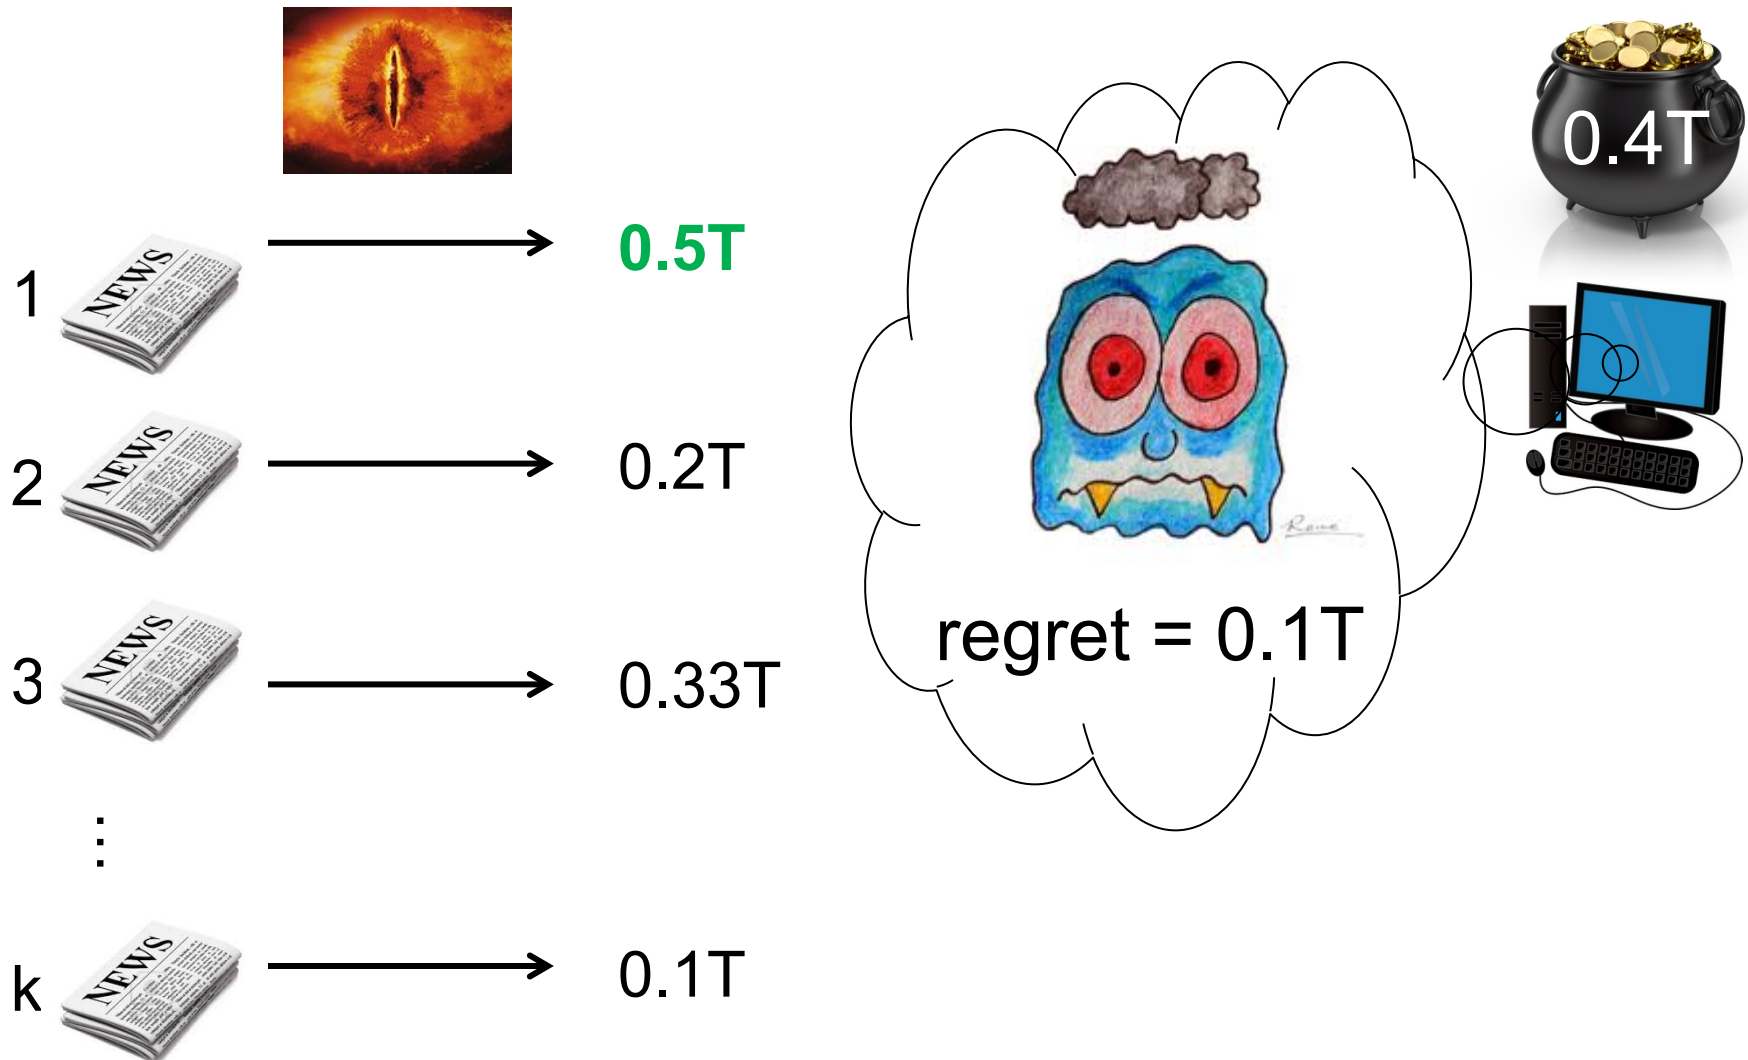
MULTIARMED BANDITS

[ROBBINS '52]



MULTIARMED BANDITS



[ROBBINS '52]



CONTEXTUAL BANDITS

[AUER-CESABIANCHI-FREUND-SCHAPIRE '02]

context:



1



2



3



⋮

k





N experts/policies/functions
think of $N \gg K$

CONTEXTUAL BANDITS

[AUER-CESABIANCHI-FREUND-SCHAPIRE '02]

context: x_1



⋮



5



1



1



4



K





3

N experts/policies/functions
think of $N \gg K$

CONTEXTUAL BANDITS

[AUER-CESABIANCHI-FREUND-SCHAPIRE '02]

context: x_1



⋮



5



1



1



4



K

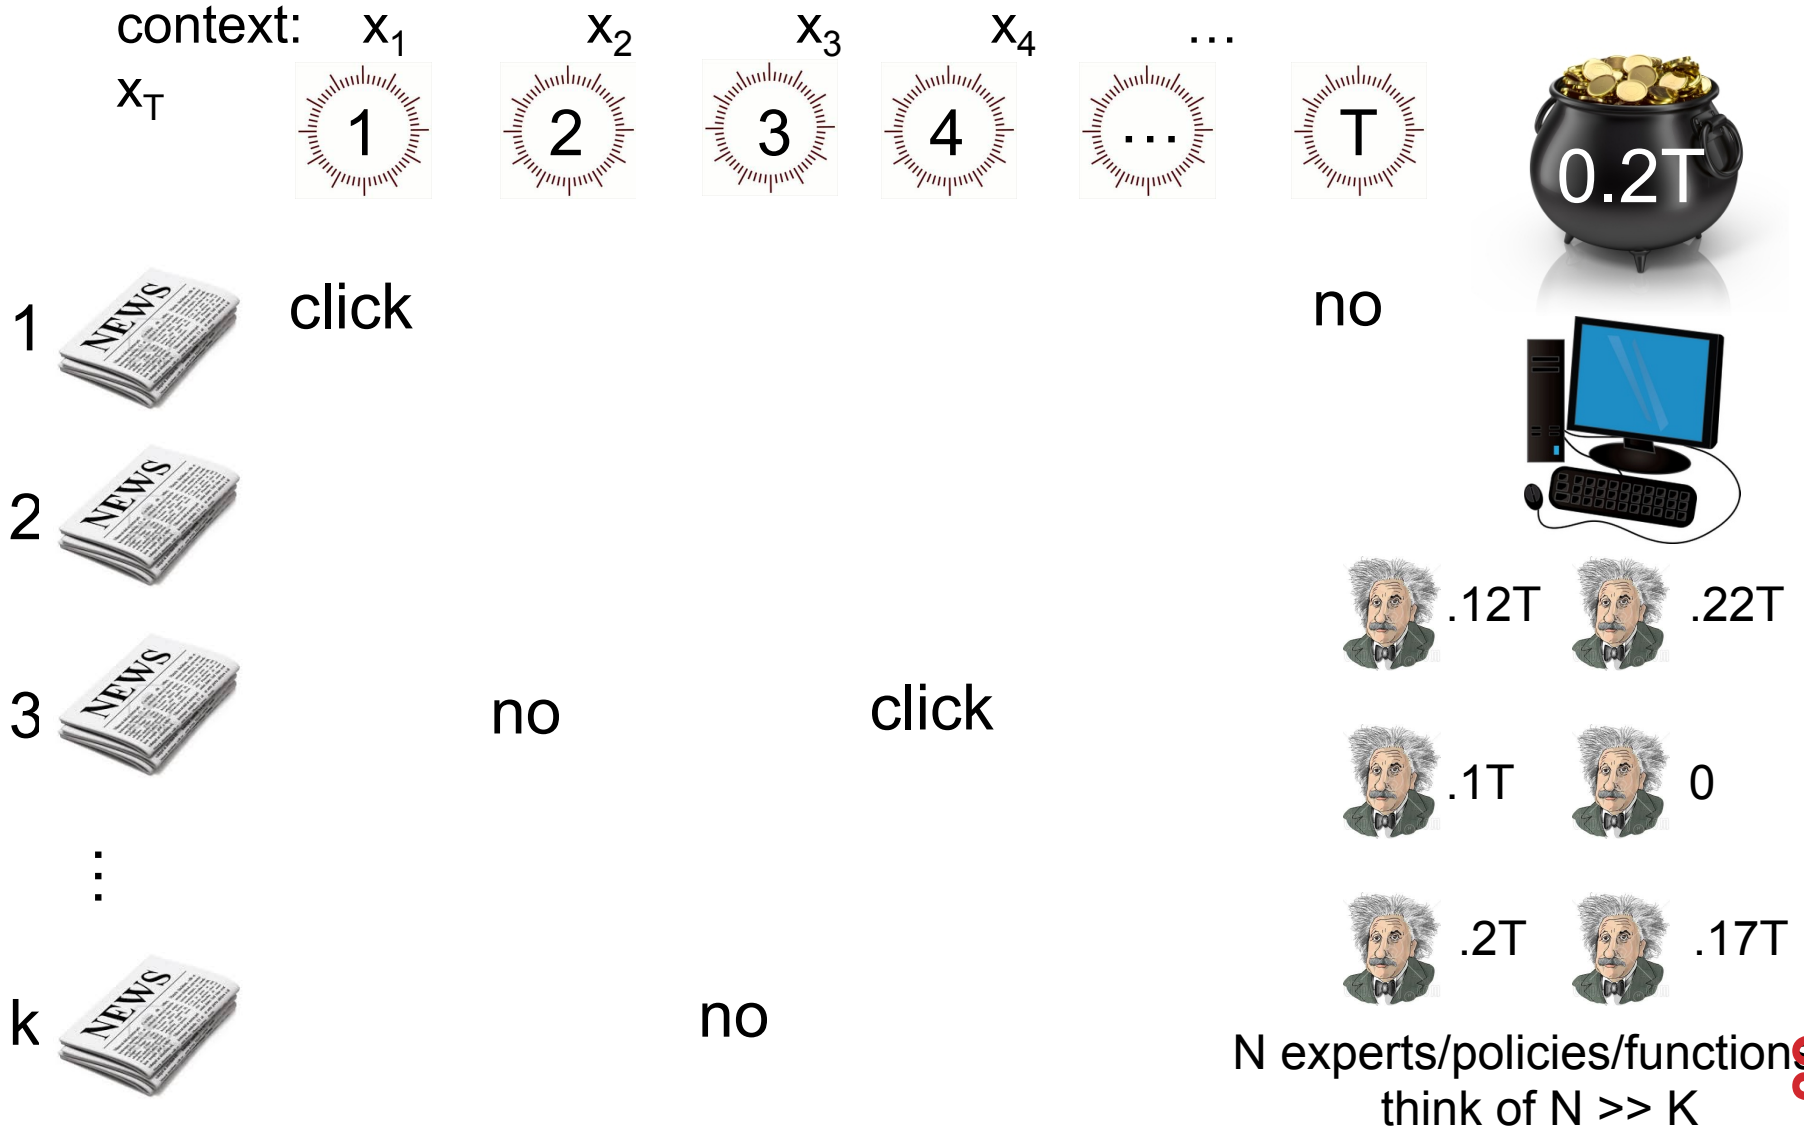


3

N experts/policies/functions
think of $N \gg K$


CONTEXTUAL BANDITS

[AUER-CESABIANCHI-FREUND-SCHAPIRE '02]




CONTEXTUAL BANDITS

[AUER-CESABIANCHI-FREUND-SCHAPIRE '02]




⋮



N experts/policies/functions
think of $N \gg K$

CONTEXTUAL BANDITS

[AUER-CESABIANCHI-FREUND-SCHAPIRE '02]



⋮

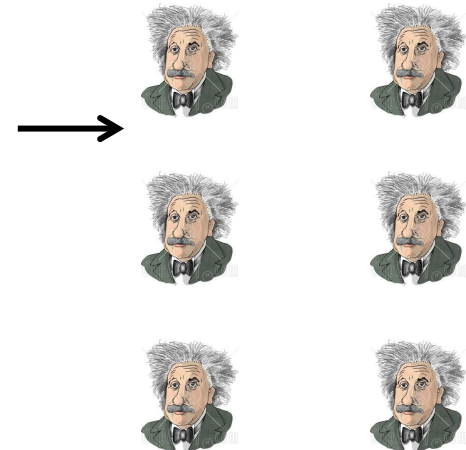


the clicks can come i.i.d. from a distribution or be arbitrary

stochastic / adversarial

The experts can be present or not.

contextual / non-contextual



BANDITS

Harder than supervised learning:

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

Many applications

Ad auctions, medicine, finance, ...

Exploration/Exploitation

Can **exploit** expert/article you've learned to be good.

Can **explore** expert/article you're not sure about.

EPSILON-FIRST

Rough idea of **ϵ -first** (or **ϵ -greedy**): act randomly for ϵ rounds, then go with best (arm or expert).

Rough analysis: even for 2 arms, we suffer regret $\epsilon + (T-\epsilon)/(\epsilon^{1/2})$.

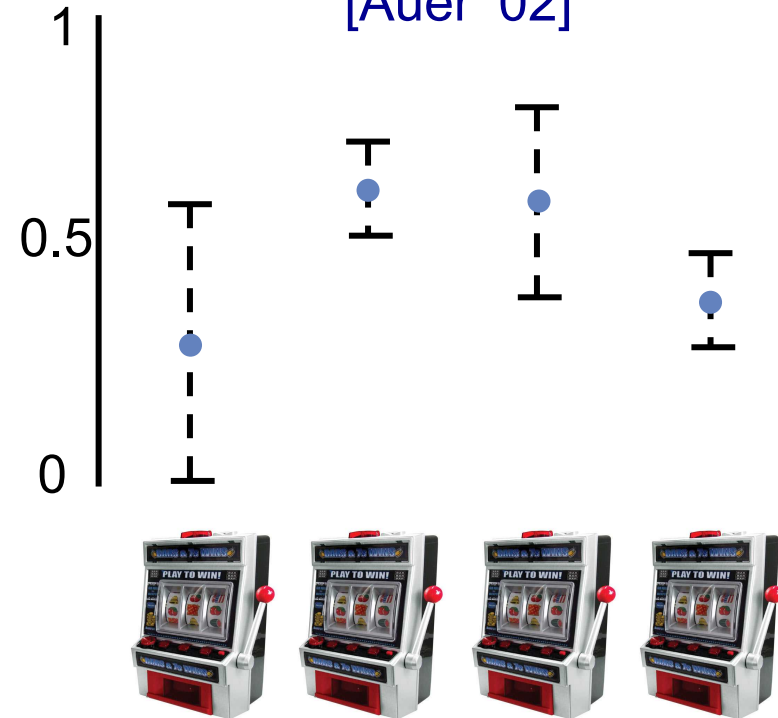
- $\epsilon \approx T^{2/3}$ is optimal tradeoff, gives regret $\approx T^{2/3}$

But actually $O(T)^{1/2}$ regret is possible!

TRADITIONAL BANDIT ALGORITHMS

UCB

[Auer '02]




Algorithm: at every time step

- 1) pull arm with highest UCB
- 2) update confidence bound of the arm pulled.

EXP3 / EW

[Littlestone-Warmuth '94]

[Auer et al. '02]



Algorithm: at every time step

- 1) sample from distribution defined by weights (mixed w/ uniform)
- 2) update weights “exponentially”

UCB VS EXP3: A COMPARISON

UCB

[AUER '02]

◆ Pros

- ◆ Optimal for the stochastic setting.
- ◆ Succeeds with high probability.

◆ Cons

- ◆ Does not work in the adversarial setting.
- ◆ Is not optimal in the contextual setting.

EXP3 & FRIENDS

[ACFS '02]

◆ Pros

- ◆ Optimal for both the adversarial and stochastic settings.
- ◆ Can be made to work in the contextual setting

◆ Cons

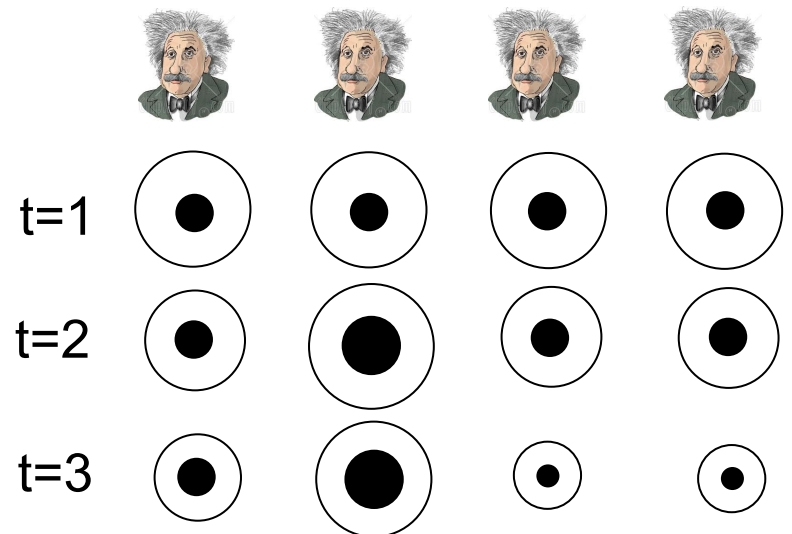
- ◆ Does not succeed with high probability in the contextual setting (only in expectation).

EXP4P

Main Theorem [Beygelzimer-Langford-Li-R-Schapire '11]:
For any $\delta > 0$, with probability at $> 1 - \delta$, EXP4P has “optimal”
regret in the adversarial contextual bandit setting.

key insights
on top of UCB/ EXP

- 1) exponential weights and upper confidence bounds “stack”
- 2) generalized Bernstein’s inequality for martingales



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

- use an upper confidence bound on each expert's estimated reward

ensuring exploration

- make sure each action is taken with some minimum probability

importance weighting

- give rare events more importance to keep estimates unbiased

Exponential Weight Algorithm for Exploration and Exploitation with Experts

Exp4.P [Beygelzimer, Langford, Li, R, Schapire '10]

Initialization: $\forall \pi \in \Pi : w_t(\pi) = 1$

For each $t = 1, 2, \dots$:

1. Observe x_t and let for $a = 1, \dots, K$

$$p_t(a) = (1 - K\rho_{\min}) \frac{\sum_{\pi} \mathbf{1}[\pi(x_t) = a] w_t(\pi)}{\sum_{\pi} w_t(\pi)} + \rho_{\min},$$

where $\rho_{\min} = \sqrt{\frac{\ln |\Pi|}{KT}}$.

2. Draw a_t from p_t , and observe reward $r_t(a_t)$.
3. Update for each $\pi \in \Pi$

$$w_{t+1}(\pi) = w_t(\pi) \exp \left(\frac{\rho_{\min}}{2} \left(\mathbf{1}[\pi(x_t) = a_t] \frac{r_t(a_t)}{p_t(a_t)} + \frac{1}{p_t(\pi(x_t))} \sqrt{\frac{\ln N/\delta}{KT}} \right) \right)$$

EXP4.P IN PRACTICE

- ◆ Application – Yahoo! front page
- ◆ We chose a special policy class for which we could efficiently keep track of the weights.
 - ◆ Created 5 clusters, with users (at each time step) getting features based on their distances to clusters.
 - ◆ Policies mapped clusters to article (action) 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.

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

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?

HOPE FOR AN EFFICIENT ALGORITHM?

[DUDIK-HSU-KALE-KARAMPATZIAKIS-LANGFORD-R-ZHANG '11]

For EXP4P, the dependence on N in the regret is logarithmic.

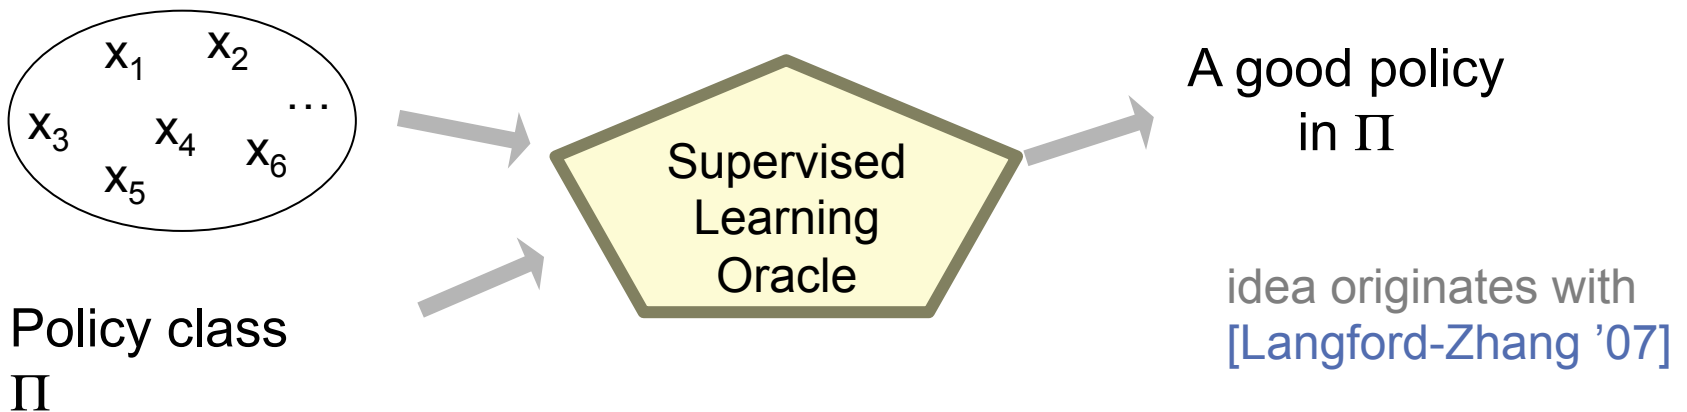
this suggests

We could compete with a large, even super-polynomial number of policies! (e.g. $N=K^{100}$ becomes $10 \log^{1/2} K$ in the regret)

however

All known contextual bandit algorithms explicitly “keep track” of the N policies. Even worse, just reading in the N would take too long for large N .

- ◆ “Competing” with a large (even exponentially large) set of policies is commonplace in supervised learning.
 - ◆ **Targets:** e.g. linear thresholds, CNF, decision trees (in practice)
 - ◆ **Methods:** e.g. boosting, SVM, neural networks, gradient descent
- ◆ The recommendations of the policies don’t need to be explicitly read in when the policy class has structure!




context:

x_1

x_2

x_3



1



2


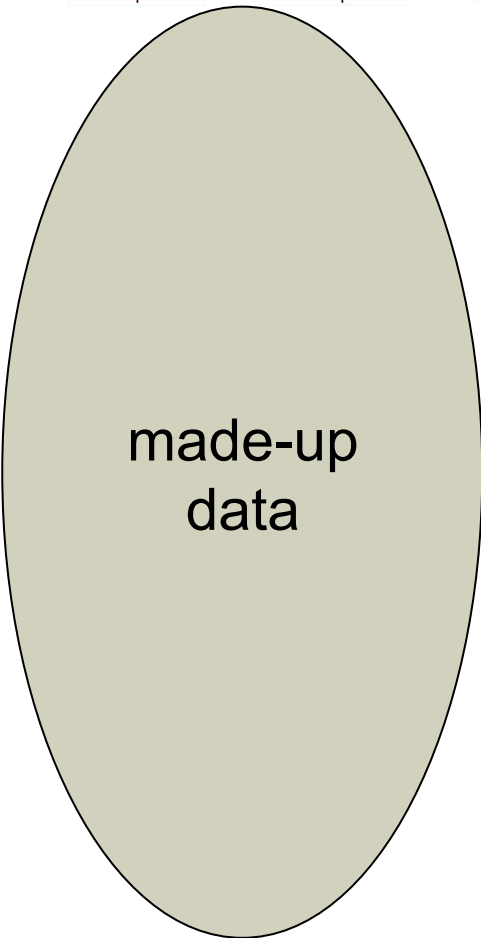


3



⋮

k







N experts/policies/functions
think of $N \gg K$



Bandit Slate Problems



[Kale-R-Schapire '11]

Sci/Tech


 **LivingSocial hacked; 50 million affected**
CNET | 1 hour ago | Written by  Seth Rosenblatt
Hackers target LivingSocial, stealing the personal data of more than 50 million people in an enormous security breach. Seth Rosenblatt.

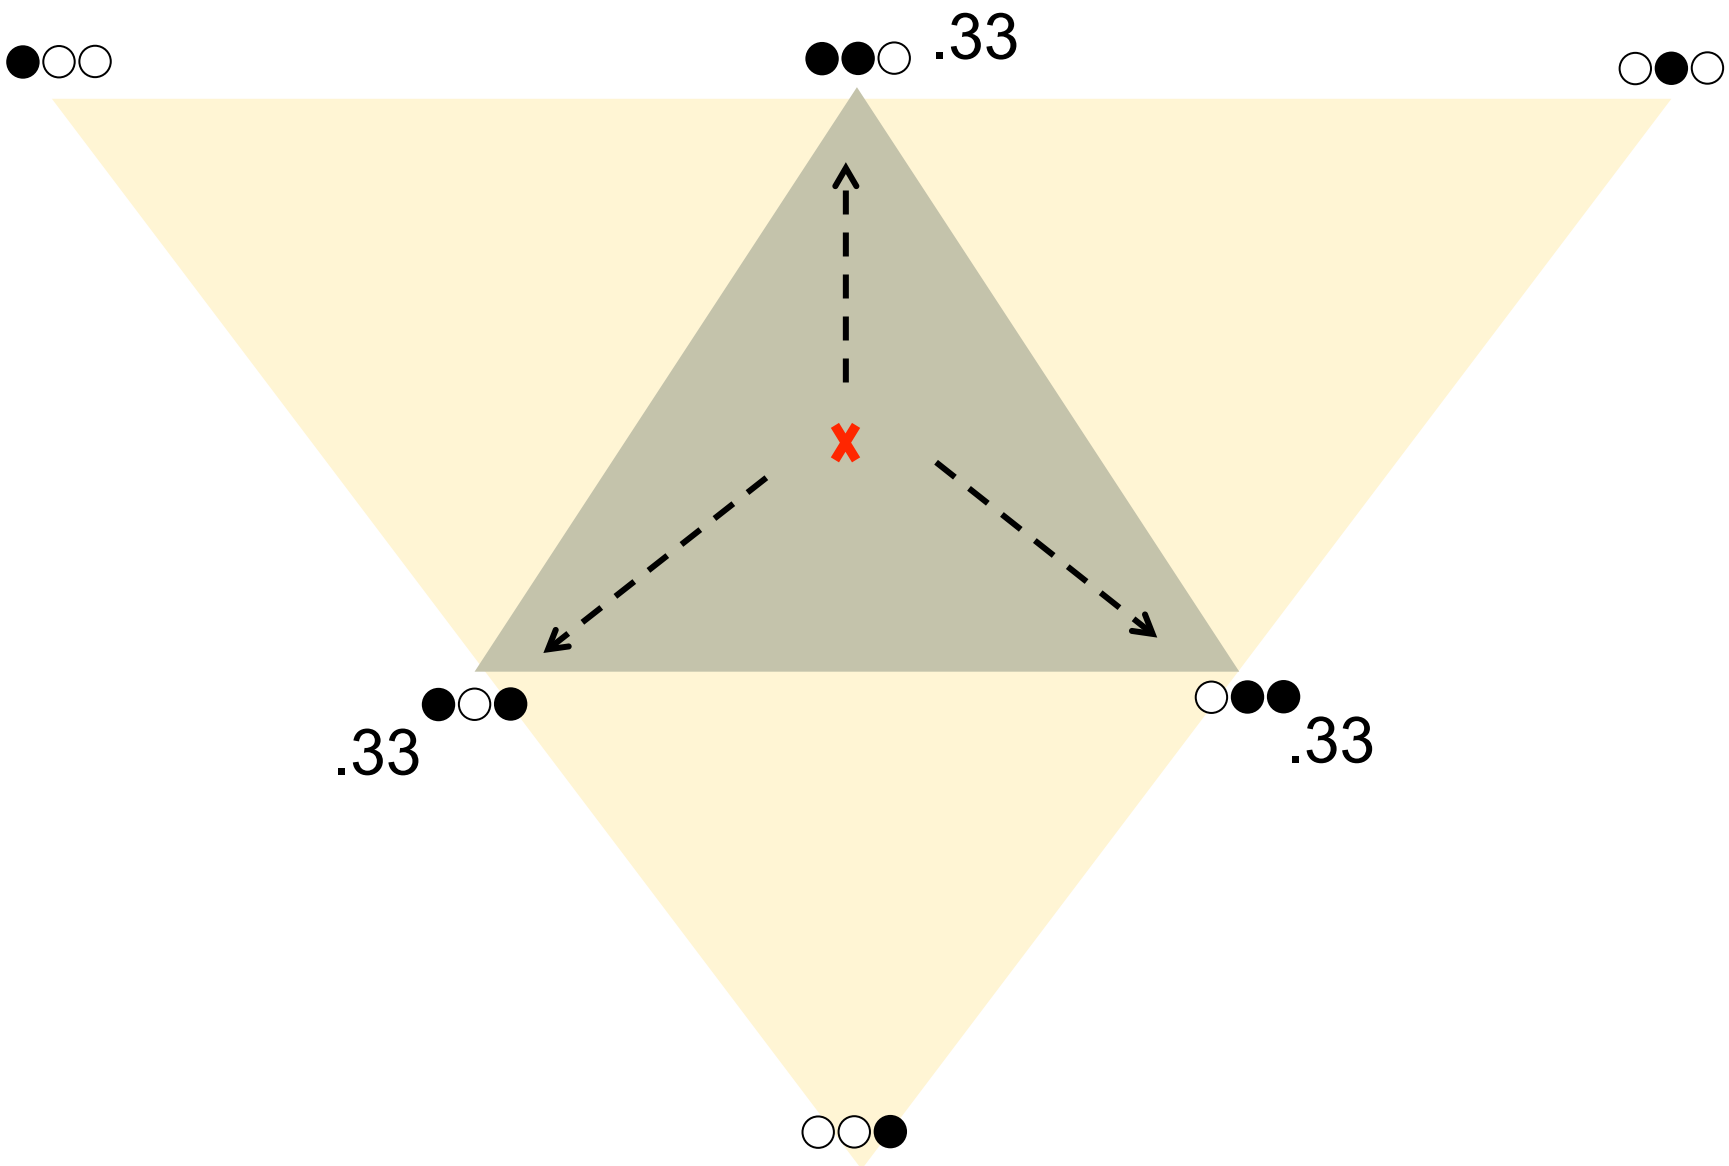
 **Apple loses more global smartphone market-share to Samsung**
San Jose Merc... | 1 hour ago | Written by  Pat May
Once again, Samsung's smartphone successes have come at Apple's expense. In its latest report, research firm IDC revealed Friday that even though robust iPhone 5 sales helped goose Apple's total smartphone sales in its most recent quarter by 6.6 percent ...

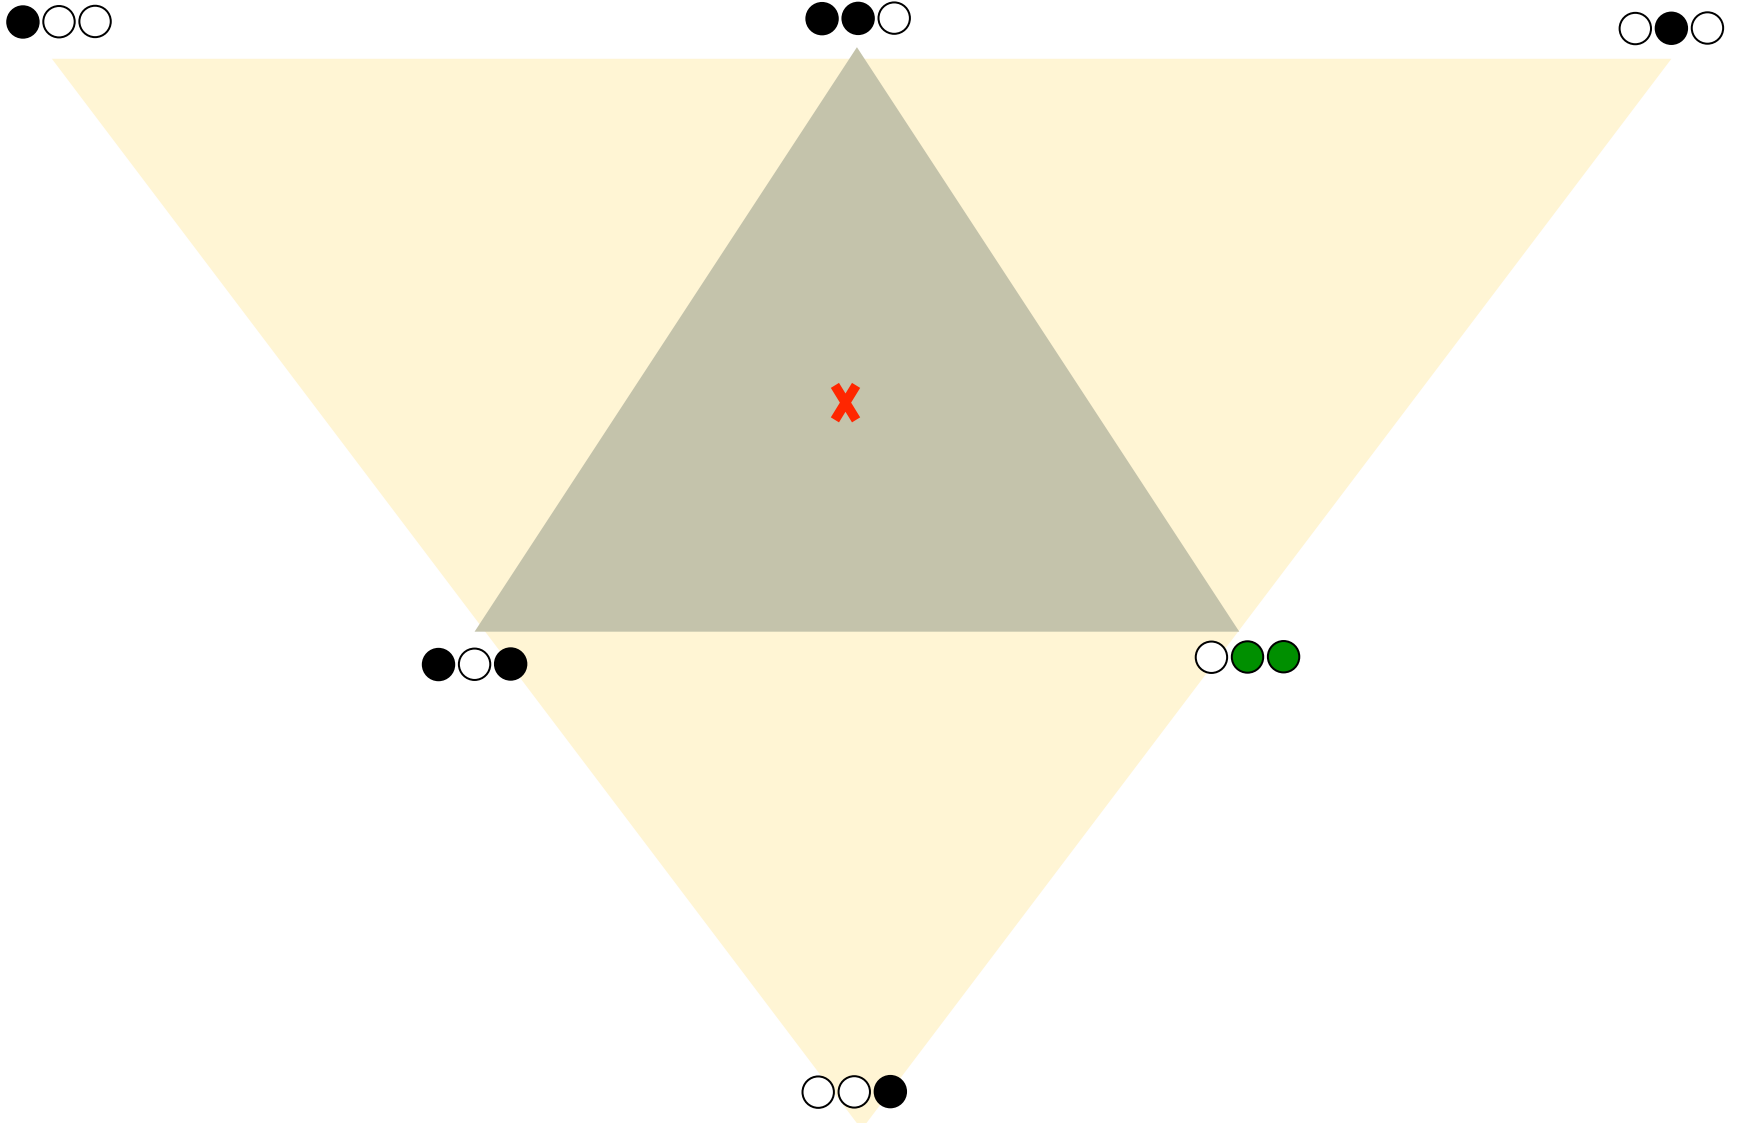
 **Monkeys imitate local food norms, study finds**
Christian Scien... | 1 hour ago | Written by  Mai Ngoc Chau
The tendency to adapt to cultural behaviors in a new place is not unique to us, a new study suggests. Skip to next paragraph. In Pictures: Monkeying around!

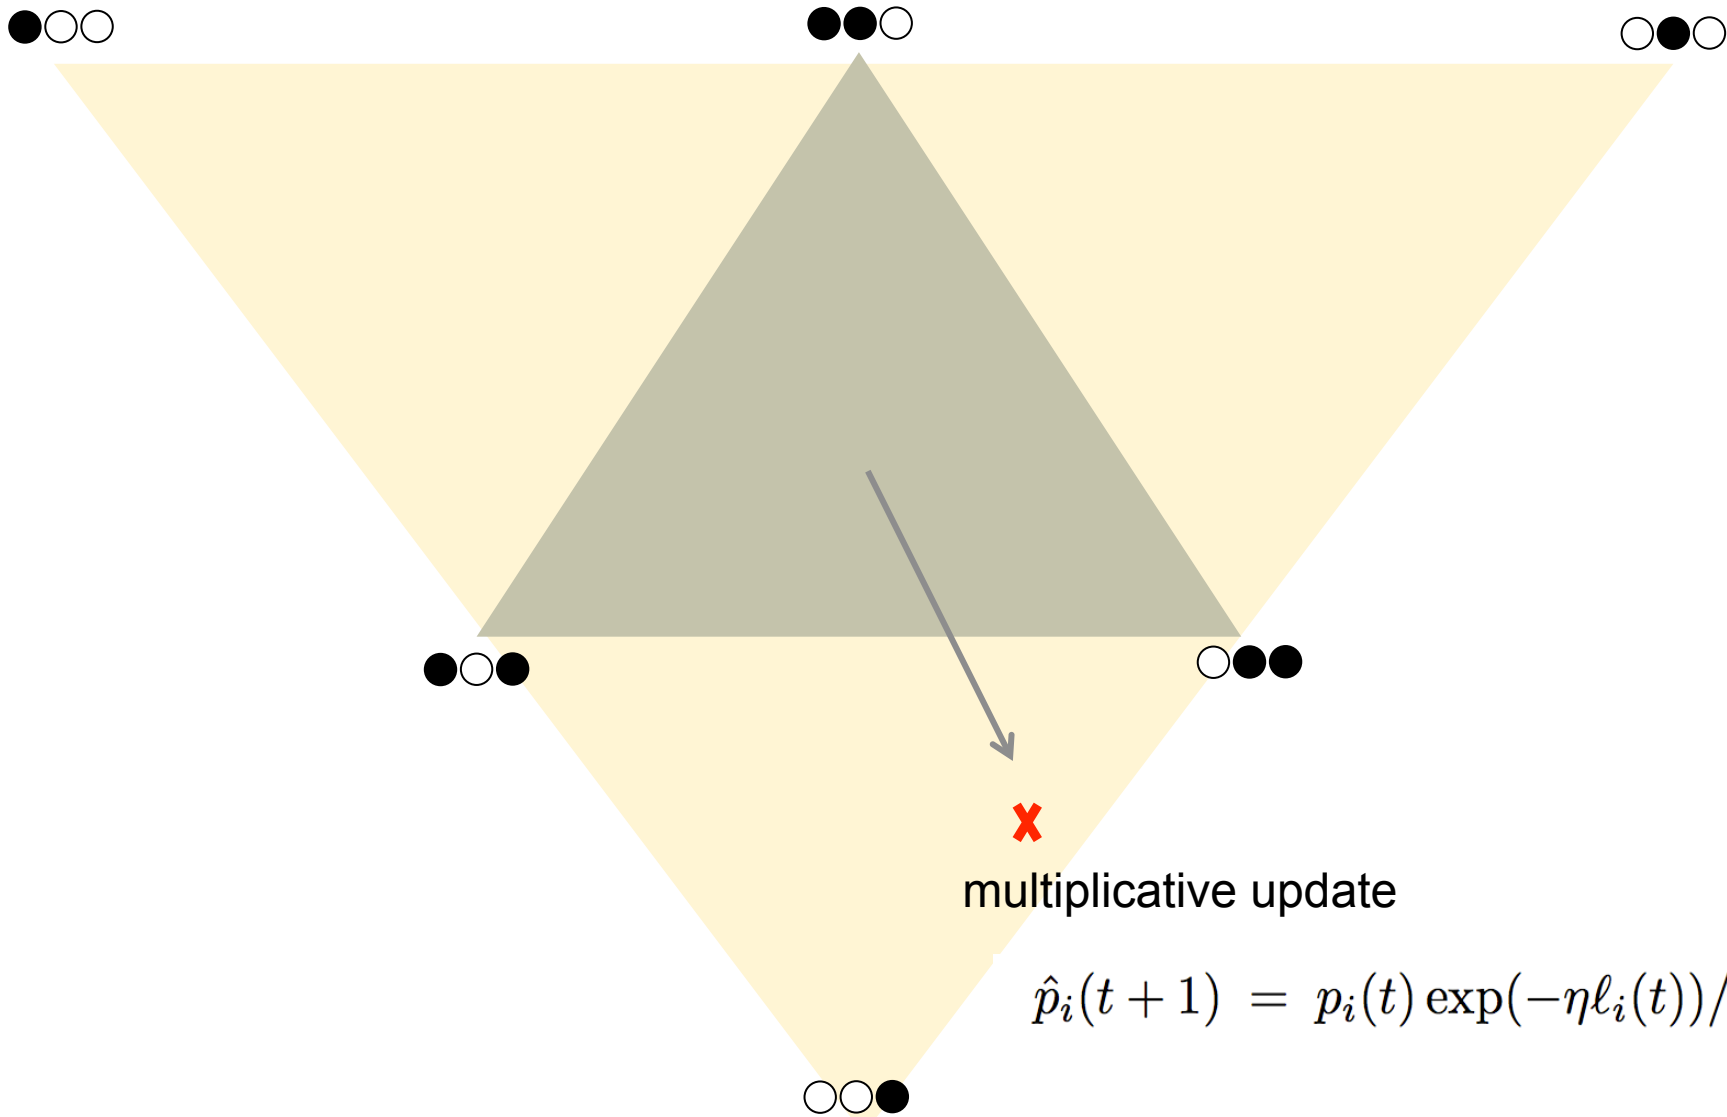
 **Dutch Man Said to Be Held in Powerful Internet Attack**
New York Times | 1 hour ago | Written by  Nicole Perloth
Dutch authorities say police officials in Spain have arrested a man believed to be connected to an online attack on a spam-fighting site that snarled the Internet last month.

Instead of selecting one article, we need to **select $s \geq 1$** , articles (possibly ranked). The motivation is web ads where a search engine shows multiple articles at once.



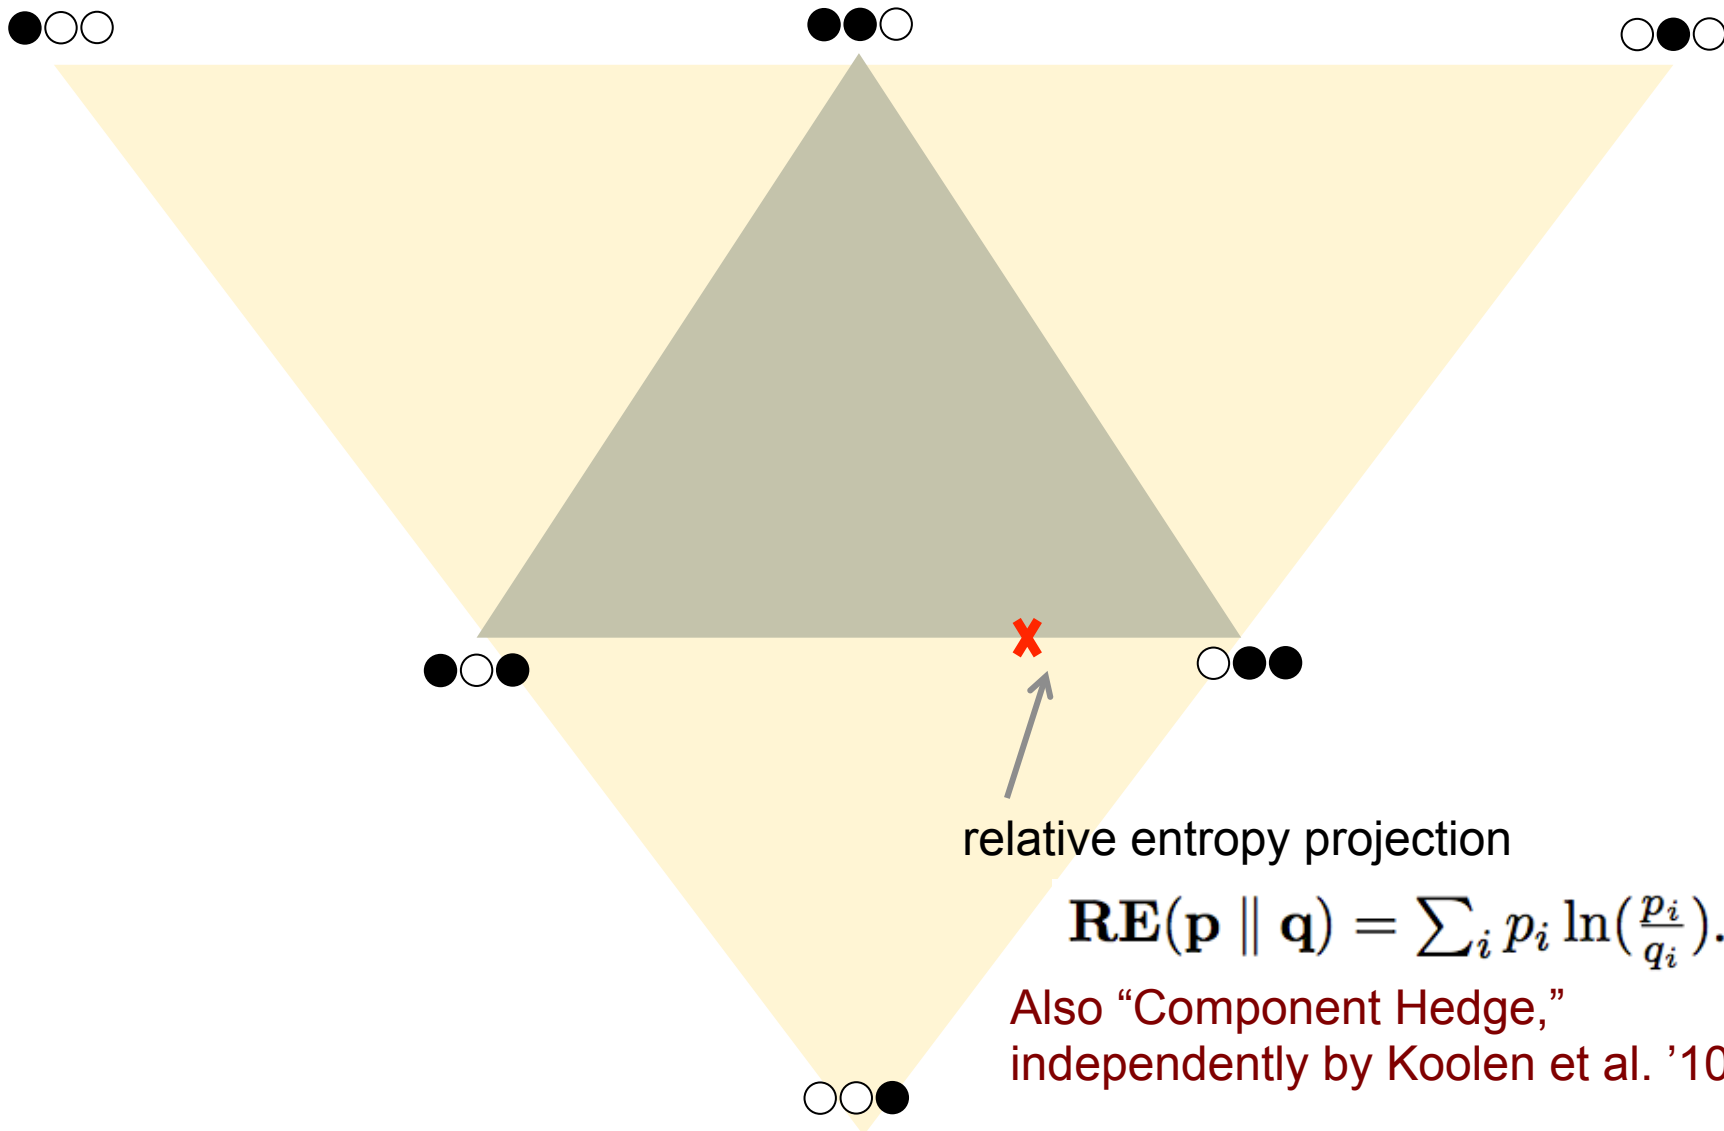







multiplicative update

$$\hat{p}_i(t+1) = p_i(t) \exp(-\eta \ell_i(t)) / Z(t)$$





Ads

Google Lev Reyzin 0 + Share

Web Images Maps Shopping More Search tools

160 personal results. 1,810,000,000 other results.

Ad related to **ipad**

[Official Apple® Store - Order iPad with Retina display now](#)
[store.apple.com/ipad](#)

★★★★★ 209 reviews for store.apple.com
Free shipping and free engraving.

241 people +1'd this page

10 hour battery life - Free In-Store Pickup - Choose Cellular or Wi-Fi - iPad 2

[Shop for ipad on Google](#)

Sponsored



[Apple iPad with Retina...](#)
\$479.99
AT&T



[iPad with Retina displa...](#)
\$499.00
Apple Store



[Apple iPad 2 Mc954ll/a Wif...](#)
\$389.00
eBay



[Apple iPad mini 16GB wit...](#)
\$299.00
Walmart



[Apple - Ipad Mini...](#)
\$329.99
Best Buy

Shop by cellular connectivity: [Wi-Fi Only](#) [3G](#)

[Apple - iPad](#)

[www.apple.com/ipad/](#)

iPad is a magical window where nothing comes between you and what you love. And it comes in two sizes.

[Shop iPad](#)

iPad 2 - Compare iPads - iPad mini -

[Compare iPad models.](#)

Whether you choose iPad mini, iPad

Ads

[iPad at Walmart](#)

[www.walmart.com/Apple_iPad](#)

Save On iPad at Walmart.
Free Shipping Site to Store.

[iPad Apple at Amazon](#)

[www.amazon.com/iPad+Apple](#)

★★★★★ 736 reviews for amazon.com
Save Big on New Gear at Amazon!
Free 2-Day Shipping w/Amazon Prime.

[50% Off iPad](#)

[www.comparedstores.com/iPad](#)

iPad Huge Discounts
2012 Clearance Sale, Free Shipping!


[Overstock iPads Sale](#)

[www.dealday.com/iPad](#)


Want Apple iPads At Cheap Prices?
Signup And Save On All Apple iPads!

[See your ad here »](#)

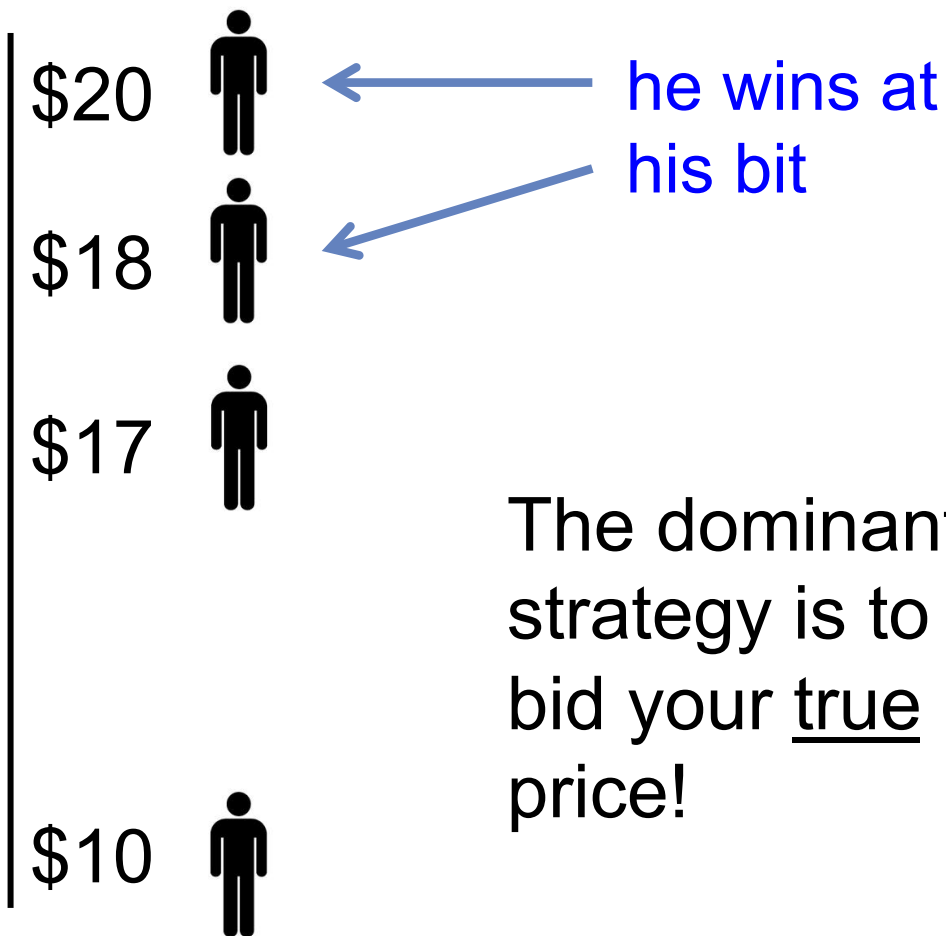
SECOND PRICE, TRUTHFUL BIDDING



[Back to search results](#) | Listed as Machine Learning t



[Click to view larger image](#)



CONTEXTUAL ADVERTISING

1. **For an ad to be shown, it must have high **expected earnings**.**
Earnings = clickthrough rate (CTR) x expected charged price
2. **CTR must be **learned****
a classic contextual bandits problem
3. **Charged prices are functions of the bids of advertisers.**
e.g. Can't ever charge more than an advertiser's bid
4. **Ads must be shown so that CTR is learned quickly, but the auction should be **truthful**.**

SUMMARY

When dealing with many customers/subscribers and many options, a smart automated strategy needs to be employed.

This is becoming true of nearly every company presenting content online.

Presents many important mathematical challenges, most of which are wide open.

Now is a great time for applied mathematicians and computer scientists to have an impact on the real world!