Counterfactual Machine Learning

CS 7792 - Fall 2016

Thorsten Joachims

Department of Computer Science & Department of Information Science

Cornell University

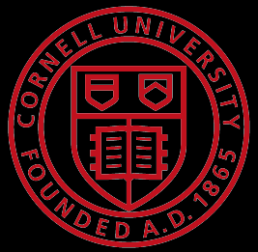

Outline of Today

- Introduction
  - Thorsten Joachims
- Overview of Class Topics
  - Machine Learning with Humans in the Loop
  - Counterfactual Model and Machine Learning
  - Challenges
- Administrivia
  - Goals for the Class
  - Pre-Requisites
  - Credit Options and Format
  - Course Material
  - Contact Info

User Interactive Systems

Examples
- Search engines
- Entertainment media
- E-commerce
- Smart homes, robots, etc.

User Behavior as Data for
- Evaluating system performance
- Learning improved systems and gathering knowledge
- Personalization

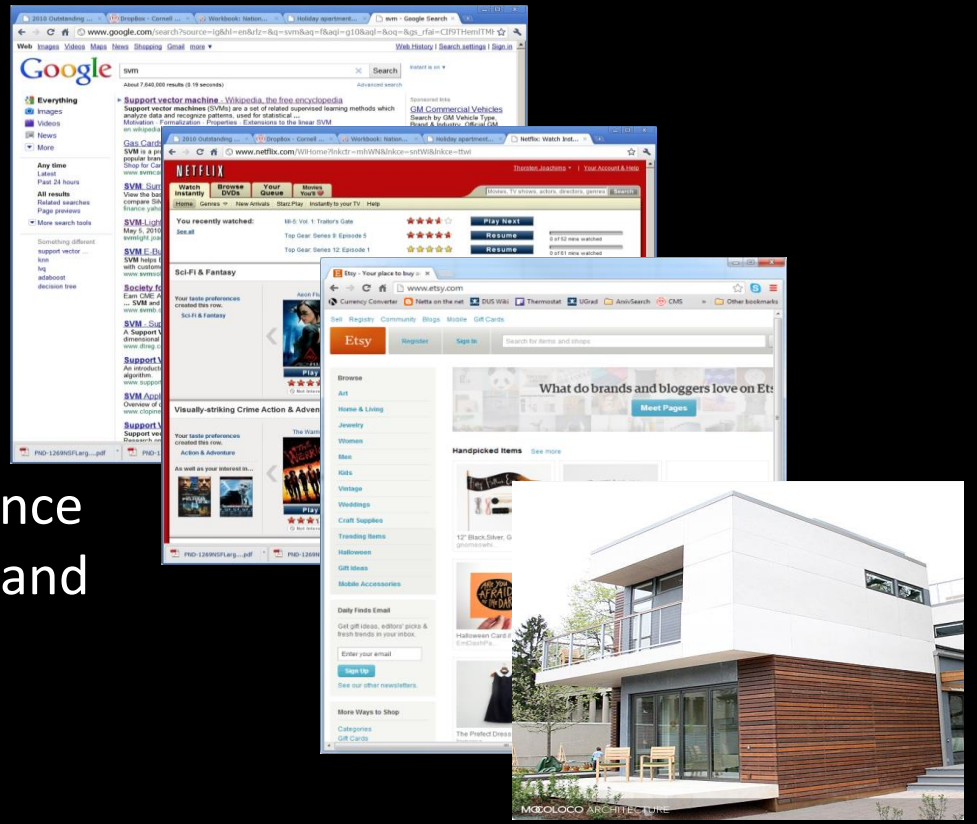

Implicit Feedback in Web Search

- Observable actions
  - Queries / reformulations
  - Clicks
  - Order, dwell time
  - Etc.
- Implicit feedback
  - Personalized
  - Democratic
  - Timely
  - Cheap
  - Abundant

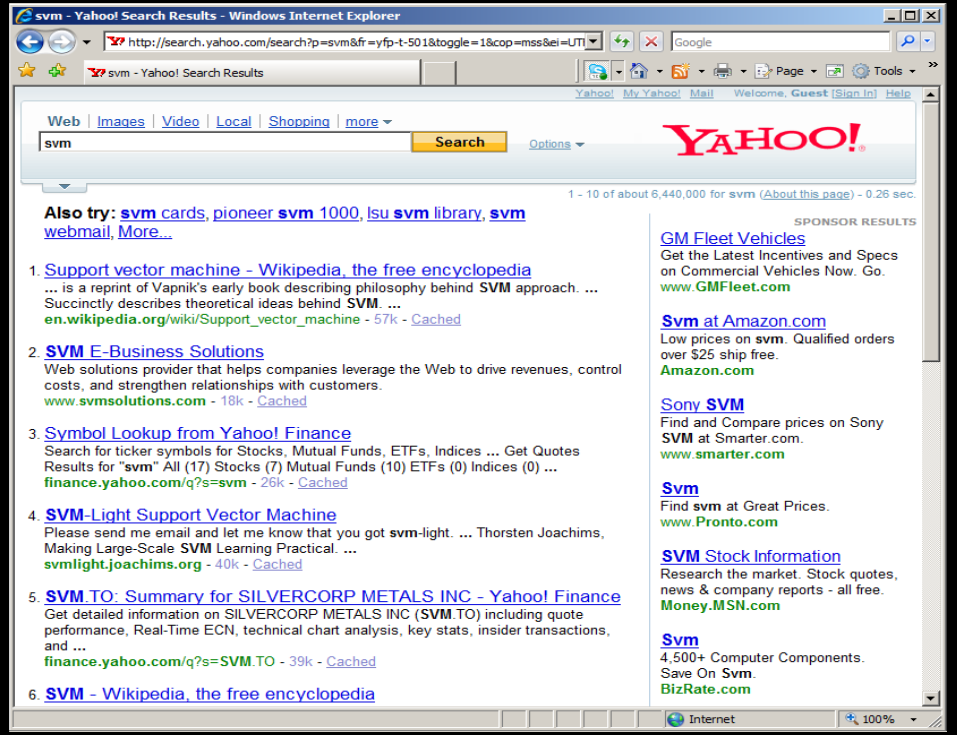

Which Ranking Function is Better?

Distribution P(u,q)
of users u, queries q

(tj,"SVM")

Retrieval Function 1
$f_1(u,q) \rightarrow r_1$

Which one is better?

Retrieval Function 2
$f_2(u,q) \rightarrow r_2$

1. Kernel Machines
   http://svm.first.gmd.de/
2. SVM-Light Support Vector Machine
   http://svmlight.joachims.org/
3. School of Veterinary Medicine at UPenn
   http://www.vet.upenn.edu/
4. An Introduction to Support Vector Machines
   http://www.support-vector.net/
5. Service Master Company
   http://www.servicemaster.com/

$U(tj,"SVM",r_1)$

1. School of Veterinary Medicine at UPenn
   http://www.vet.upenn.edu/
2. Service Master Company
   http://www.servicemaster.com/
3. Support Vector Machine
   http://jbolivar.freeservers.com/
4. Archives of SUPPORT-VECTOR-MACHINES
   http://www.jiscmail.ac.uk/lists/SUPPORT...
5. SVM-Light Support Vector Machine
   http://ais.gmd.de/~thorsten/svm light/

$U(tj,"SVM",r_2)$

Measuring Utility

Name	Description	Aggre-gation	Hypothesized Change with Decreased Quality
Abandonment Rate	% of queries with no click	N/A	Increase
Reformulation Rate	% of queries that are followed by reformulation	N/A	Increase
Queries per Session	Session = no interruption of more than 30 minutes	Mean	Increase
Clicks per Query	Number of clicks	Mean	Decrease
Click@1	% of queries with clicks at position 1	N/A	Decrease
Max Reciprocal Rank*	1/rank for highest click	Mean	Decrease
Mean Reciprocal Rank*	Mean of 1/rank for all clicks	Mean	Decrease
Time to First Click*	Seconds before first click	Median	Increase
Time to Last Click*	Seconds before final click	Median	Decrease

(*) only queries with at least one click count

ArXiv.org: User Study

User Study in ArXiv.org

- Natural user and query population
- User in natural context, not lab
- Live and operational search engine
- Ground truth by construction

ORIG ≻ SWAP2 ≻ SWAP4

- ORIG: Hand-tuned fielded
- SWAP2: ORIG with 2 pairs swapped
- SWAP4: ORIG with 4 pairs swapped

ORIG ≻ FLAT ≻ RAND

- ORIG: Hand-tuned fielded
- FLAT: No field weights
- RAND : Top 10 of FLAT shuffled

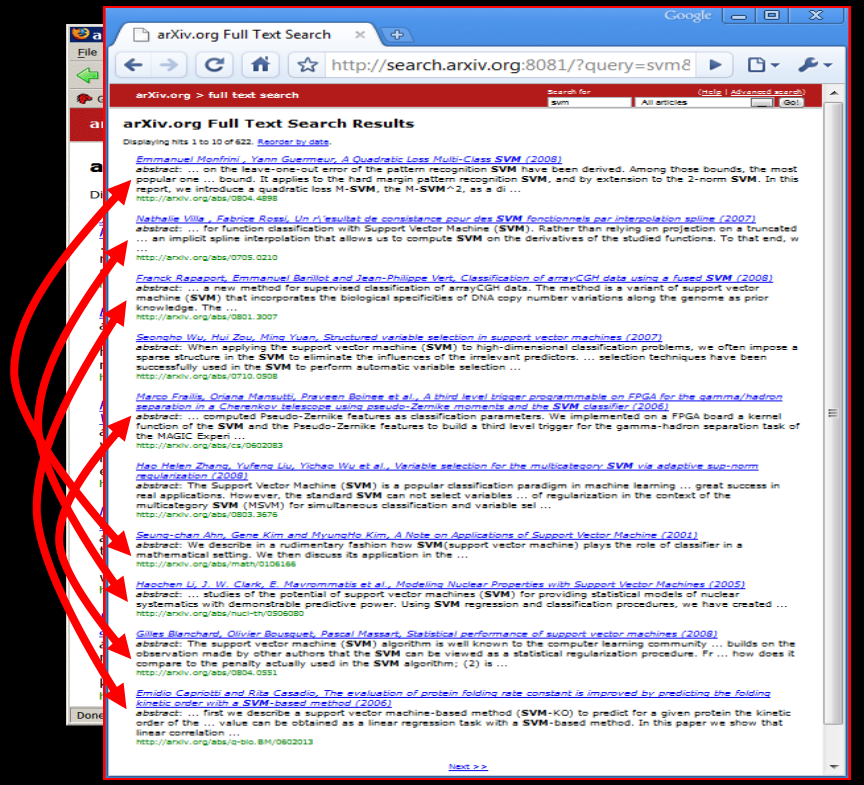

[Radlinski et al., 2008]

ArXiv.org: Experiment Setup

- Experiment Setup
  - Phase I: 36 days
    - Users randomly receive ranking from Orig, Flat, Rand
  - Phase II: 30 days
    - Users randomly receive ranking from Orig, Swap2, Swap4
  - User are permanently assigned to one experimental condition based on IP address and browser.
- Basic Statistics
  - ~700 queries per day / ~300 distinct users per day
- Quality Control and Data Cleaning
  - Test run for 32 days
  - Heuristics to identify bots and spammers
  - All evaluation code was written twice and cross-validated

Arxiv.org: Results

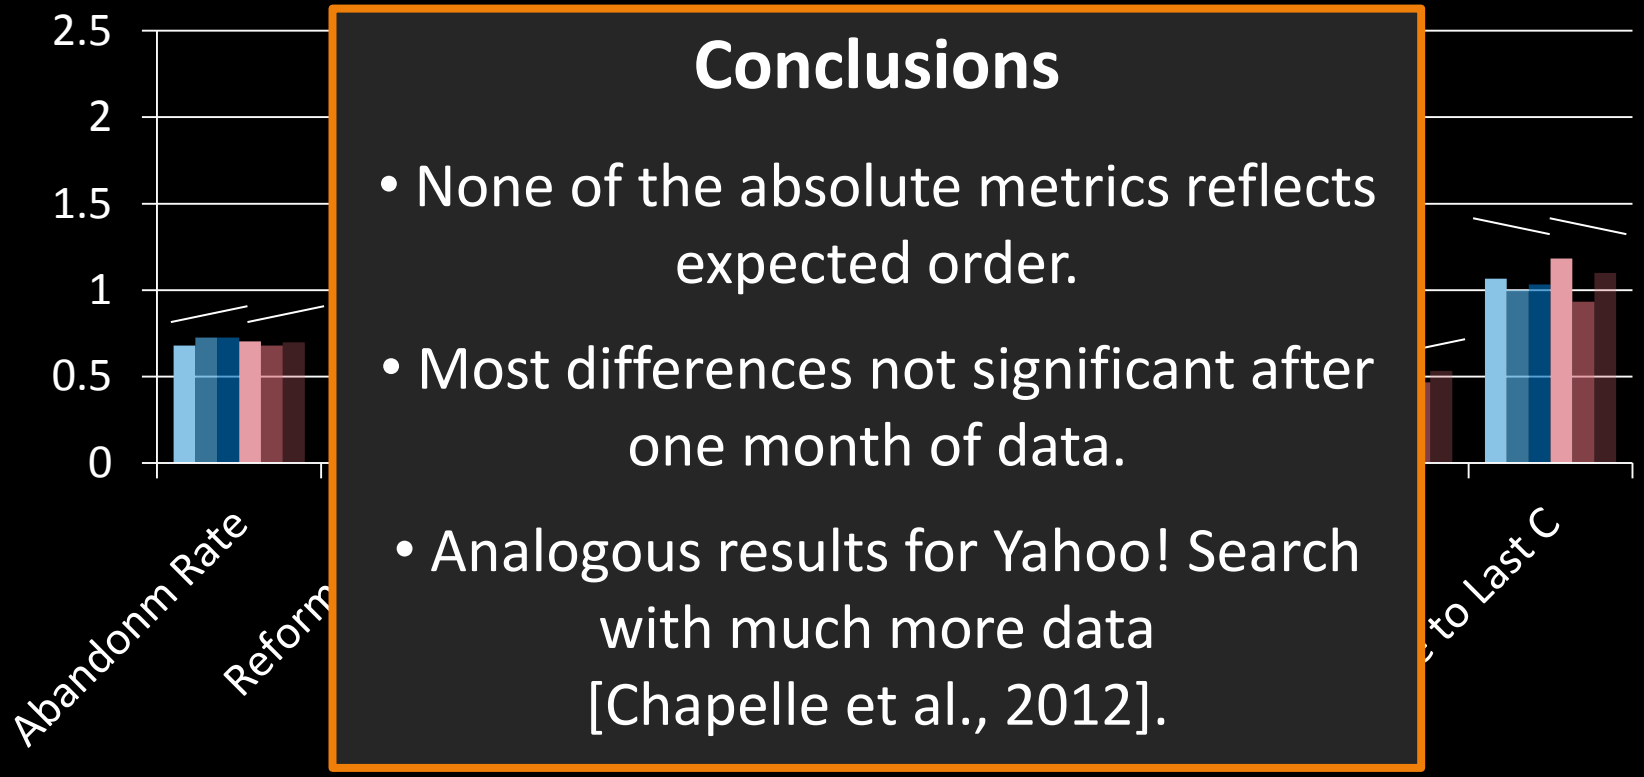

Conclusions

- None of the absolute metrics reflects expected order.
- Most differences not significant after one month of data.
- Analogous results for Yahoo! Search with much more data [Chapelle et al., 2012].

[Radlinski et al., 2008]

Economic Models of Decision Making

- Rational Choice
  - Alternatives Y
  - Utility function $U(y)$
  - Decision $\bar{y} = \underset{y \in Y}{\operatorname{argmax}}\{U(y)\}$
- Bounded Rationality
  - Time constraints
  - Computation constraints
  - Approximate $U(y)$
- Behavioral Economics
  - Framing
  - Fairness
  - Loss aversion

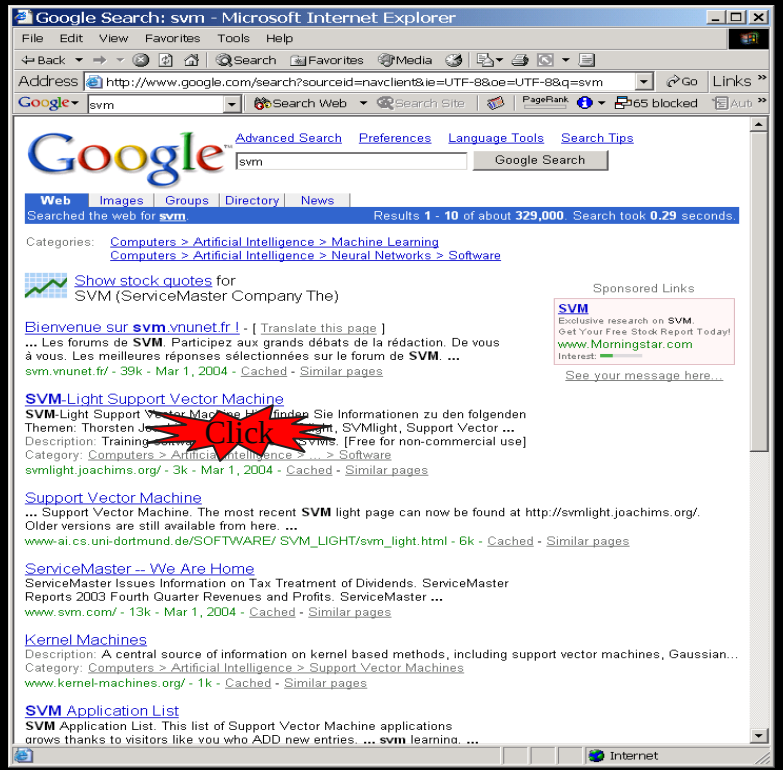

A Model of how Users Click in Search

- Model of clicking:
  - Users explore ranking to position k
  - Users click on most relevant (looking) links in top k
  - Users stop clicking when time budget up or other action more promising (e.g. reformulation)
  - Empirically supported by [Granka et al., 2004]

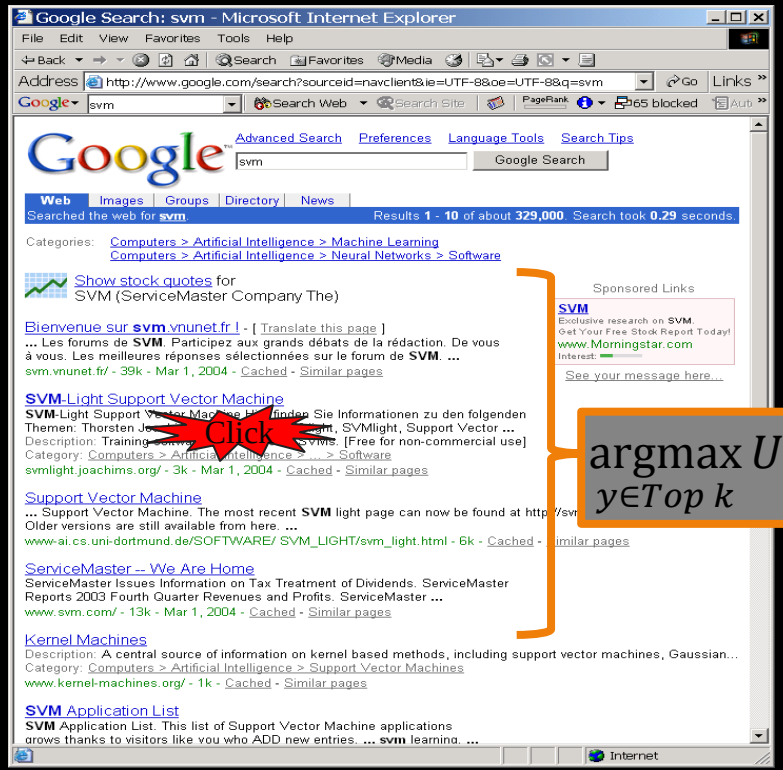

Balanced Interleaving

(u=tj, q="svm")

$f_1(u,q) \rightarrow r_1$ ⟷ $f_2(u,q) \rightarrow r_2$

1. Kernel Machines *http://svm.first.gmd.de/*	
2. Support Vector Machine *http://jbolivar.freeservers.com/*	
3. An Introduction to Support Vector Machines *http://www.support-vector.net/*	
4. Archives of SUPPORT-VECTOR-MACHINES ... *http://www.jiscmail.ac.uk/lists/SUPPORT...*	
5. SVM-Light Support Vector Machine *http://ais.gmd.de/~thorsten/svm light/*	

1. Kernel Machines *http://svm.first.gmd.de/*	
2. SVM-Light Support Vector Machine *http://ais.gmd.de/~thorsten/svm light/*	
3. Support Vector Machine and Kernel ... References *http://svm.*research.bell-labs.com/SVMrefs.html	
4. Lucent Technologies: SVM demo applet *http://svm.*research.bell-labs.com/SVT/SVMsvt.html	
5. Royal Holloway Support Vector Machine *http://svm.dcs.rhbnc.ac.uk*	

Interleaving(r_1,r_2)

1. Kernel Machines *http://svm.first.gmd.de/*	1
2. Support Vector Machine *http://jbolivar.freeservers.com/*	2
3. SVM-Light Support Vector Machine *http://ais.gmd.de/~thorsten/svm light/*	2
4. An Introduction to Support Vector Machines *http://www.support-vector.net/*	3
5. Support Vector Machine and Kernel ... References *http://svm.*research.bell-labs.com/SVMrefs.html	3
6. Archives of SUPPORT-VECTOR-MACHINES ... *http://www.jiscmail.ac.uk/lists/SUPPORT...*	4
7. Lucent Technologies: SVM demo applet *http://svm.*research.bell-labs.com/SVT/SVMsvt.html	4

Model of User:
Boundedly rational choice among top k observed results.

Invariant:
For all k, top k of balanced interleaving is union of top k_1 of r_1 and top k_2 of r_2 with $k_1=k_2 \pm 1$.

Interpretation: $(r_1 \succ r_2) \leftrightarrow \text{clicks}(\text{topk}(r_1)) > \text{clicks}(\text{topk}(r_2))$

→ see also [Radlinski, Craswell, 2012] [Hofmann, 2012]

[Joachims, 2001] [Radlinski et al., 2008]

Arxiv.org: Interleaving Experiment

- Experiment Setup
  - Phase I: 36 days
    - Balanced Interleaving of (Orig,Flat) (Flat,Rand) (Orig,Rand)
  - Phase II: 30 days
    - Balanced Interleaving of (Orig,Swap2) (Swap2,Swap4) (Orig,Swap4)
- Quality Control and Data Cleaning
  - Same as for absolute metrics

Arxiv.org: Interleaving Results

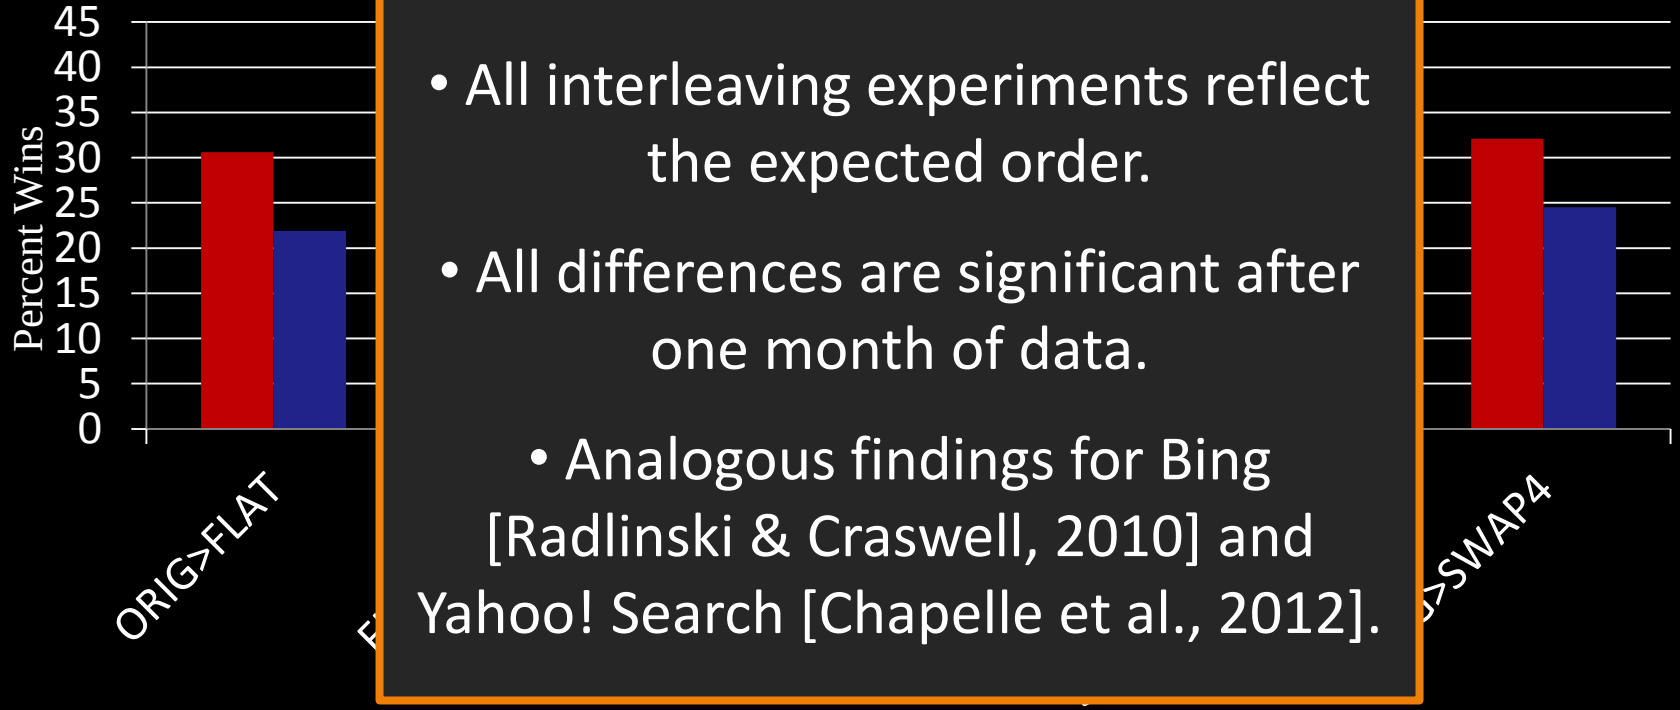

Conclusions

- All interleaving experiments reflect the expected order.
- All differences are significant after one month of data.
- Analogous findings for Bing [Radlinski & Craswell, 2010] and Yahoo! Search [Chapelle et al., 2012].

Using Behavior as Feedback

- Measuring User Satisfaction
  - Need behavioral model to get accurate training data out of biased feedback
  - Use experimental control to collect unbiased data

→ Data comes from experiment, not omniscient teacher

Interactive Learning System

Context x

Feedback $\delta(x, y)$

System π_0

Utility: $U(\pi_0)$

Action y for x

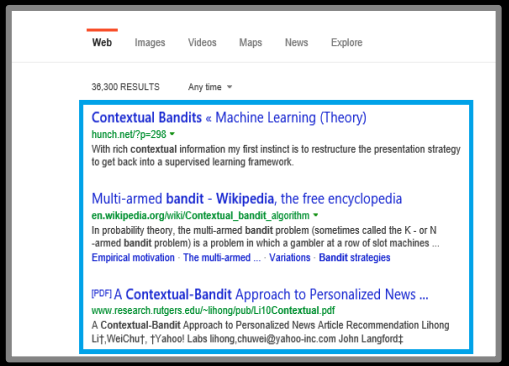

Ad Placement

- Context x:
  - User and page
- Action y:
  - Ad that is placed
- Feedback $\delta(x, y)$:
  - Click / no-click

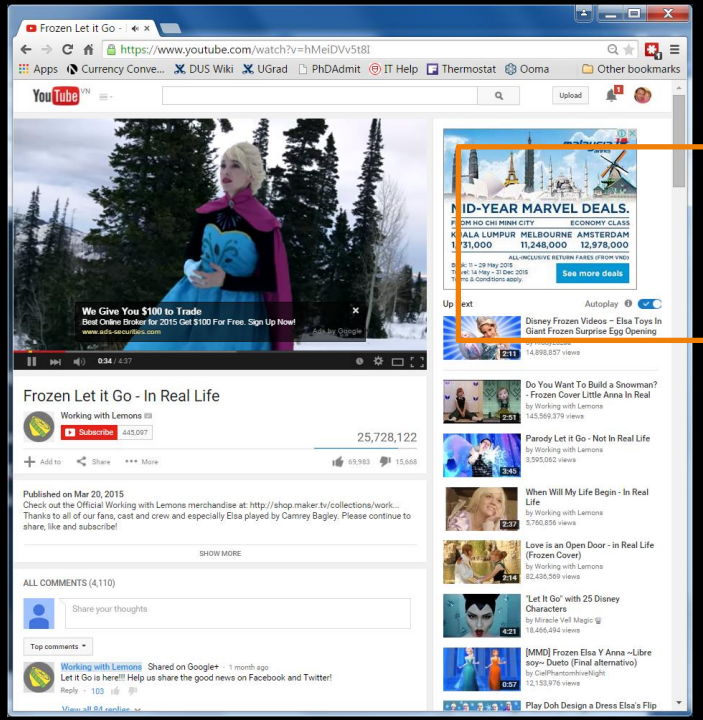

News Recommender

- Context x:
  - User
- Action y:
  - Portfolio of newsarticles
- Feedback $\delta(x, y)$:
  - Reading time in minutes

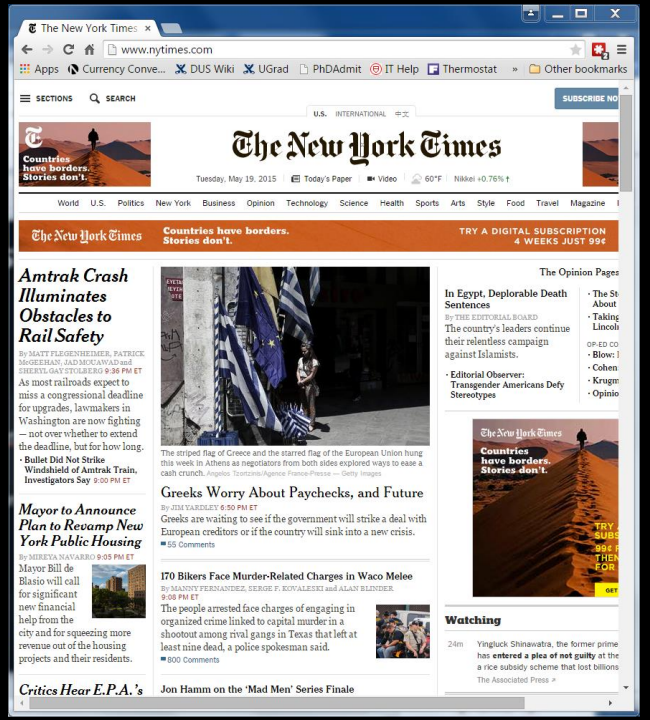

Search Engine

- Context x:
  - Query
- Action y:
  - Ranking
- Feedback $\delta(x, y)$:
  - win/loss against baseline in interleaving

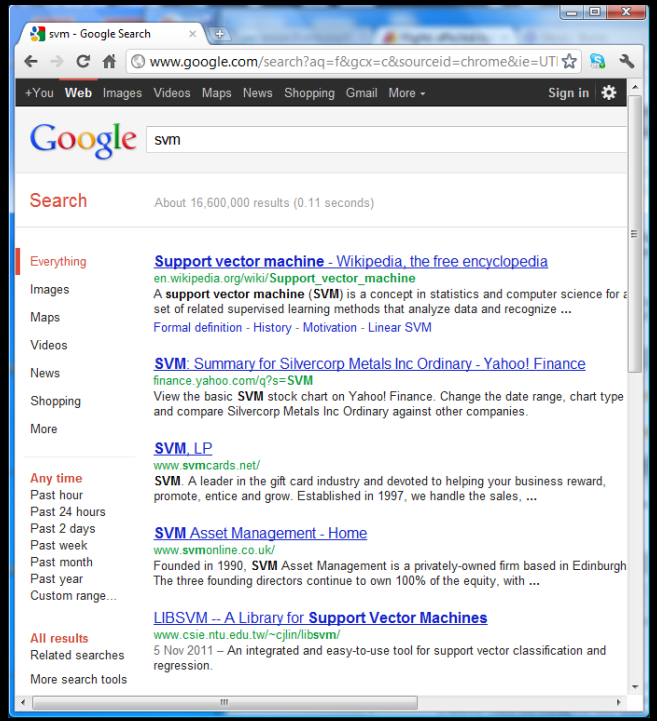

Log Data from Interactive Systems

- Data

$$S = \left((x_1, y_1, \delta_1), \ldots, (x_n, y_n, \delta_n)\right)$$

→ Partial Information (aka "Contextual Bandit") Feedback

- Properties
  - Contexts x_i drawn i.i.d. from unknown $P(X)$
  - Actions y_i selected by existing system π_0: $X \rightarrow Y$
  - Feedback δ_i from unknown function $\delta: X \times Y \rightarrow \Re$

[Zadrozny et al., 2003] [Langford & Li], [Bottou, et al., 2014]

Learning from User Behavior

- Data dependent on system actions
  - Not full information, but partial information feedback
  - Data comes from experiments, not teacher
- Ability to run interactive experiments with users
  - Adaptive vs. stationary experiment control
  - Exploration/exploitation trade-offs
- Reusing existing log data
  - Observational vs. experimental data
  - Stochastic vs. deterministic logging systems

Overall Goals for this Class

- Deeply explore one active research area in ML.
  - Batch Learning from Bandit Feedback
  - Learning under selection bias and MNAR data
  - ML algorithms based in counterfactual model
  - Behavioral feedback models

  → Incredibly narrow focus.
- Practice being a successful academic

→ Class targeted towards current PhD students with research interests in this area!

Pre-Requisites

- This is not an introductory Machine Learning class!
- You need to satisfy one of the following ML pre-reqs:
  - Successfully taken CS4780 “Machine Learning”
  - Successfully taken CS6780 “Advanced Machine Learning”
  - Successfully taken a comparable “Intro to ML” class (*)
  - Acquired the equivalent ML knowledge in some other way (e.g. strong background in Statistics + ML textbook) (*)
- Currently doing or planning to do research in this area of ML
- Basic probability, basic statistics, general mathematical maturity

(*) means talk to me

Format of Class

- Lectures (by TJ)
  - Background material
- Research paper presentations (by students)
  - Explore current state of the art
- Peer reviewing

Research Paper Presentations

- Students present the paper in class
  - Slide presentation
  - Create critique, extended bibliography, examples, demo software, experiments etc. that help understand the paper
  - Prepare discussion topics / group activity
  - Prepare quiz
- Everybody reads the paper in preparation for class
  - Quiz about each paper
- All students give feedback afterwards.

Credit Options and Grades

- Pass/Fail: Need to get at least 50% of points on each of following to pass.
  - paper presentation
  - in-class quizzes (lowest grades replaced by second lowest grade)
  - peer reviewing (lowest grades replaced by second lowest grade)
  - in-class participation
- Letter grade:
  - not allowed
- Audit:
  - not allowed, unless you have very good arguments

Course Material

- Reference Books
  - Imbens, Rubin, "Causal Inference for Statistics, Social, and Biomedical Sciences", Cambridge University Press, 2015. (online via Cornell Library)
  - Morgan, Winship "Counterfactuals and Causal Inference", Cambridge University Press, 2007.
- Background Reading
  - K. Murphy, "Machine Learning - a Probabilistic Perspective", MIT Press, 2012. (online via Cornell Library)
  - B. Schoelkopf, A. Smola, "Learning with Kernels", MIT Press, 2001. (online)
  - C. Bishop, “Pattern Recognition and Machine Learning”, Springer, 2006.
  - R. Duda, P. Hart, D. Stork, "Pattern Classification“, Wiley, 2001.
  - T. Hastie, R. Tishirani, and J. Friedman, "The Elements of Statistical Learning“, Springer, 2001.
- Slides, Notes and Papers
  - Slides available on course homepage
  - Papers on course homepage

How to Get in Touch

- Course Web Page
  - http://www.cs.cornell.edu/Courses/cs7792/2016fa/
- Email
  - Thorsten Joachims: tj@cs.cornell.edu
- Office Hours
  - Fridays 11:10pm – 12:10pm, 236 Gates Hall
- Piazza
  - https://piazza.com/cornell/fall2016/cs7792
- Peer reviewing platform
  - https://cmt.research.microsoft.com/CS2016