Privacy, AI, and the AI Enterprise

Eric Horvitz

IAPP 2015
Washington DC

Hot Commodities

Centrality of Machine Learning

Machine Learning

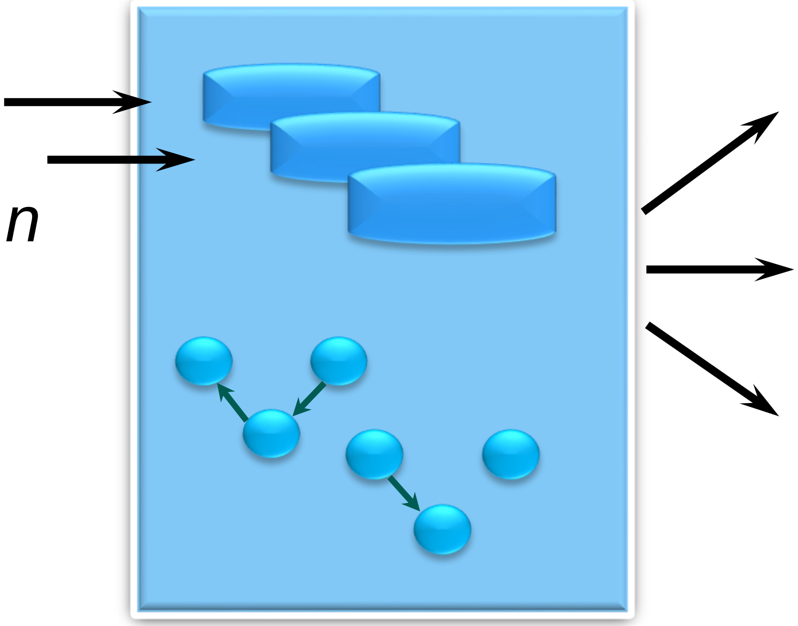

Observe 1

Observe *n*

Preferences
Goals
Actions
Content
Gender
Identity
Locations
Health

Consent. Terms of services: declaration of policy, opt-out

Centrality of Machine Learning

Machine Learning

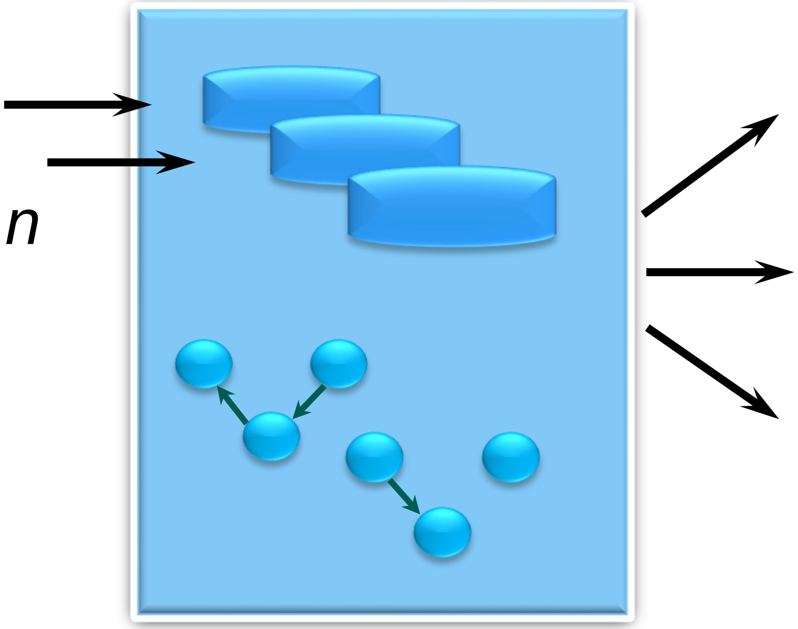

Observe 1

Observe *n*

Preferences
Goals
Actions
Content
Gender
Identity
Locations
Health

Consent. Terms of services: declaration of policy, opt-out

"May I access your location to enhance services?"

"Umm…I guess so.

AI for Minimally-Invasive Sensing

Needs → *Consider information value & sensitivity*

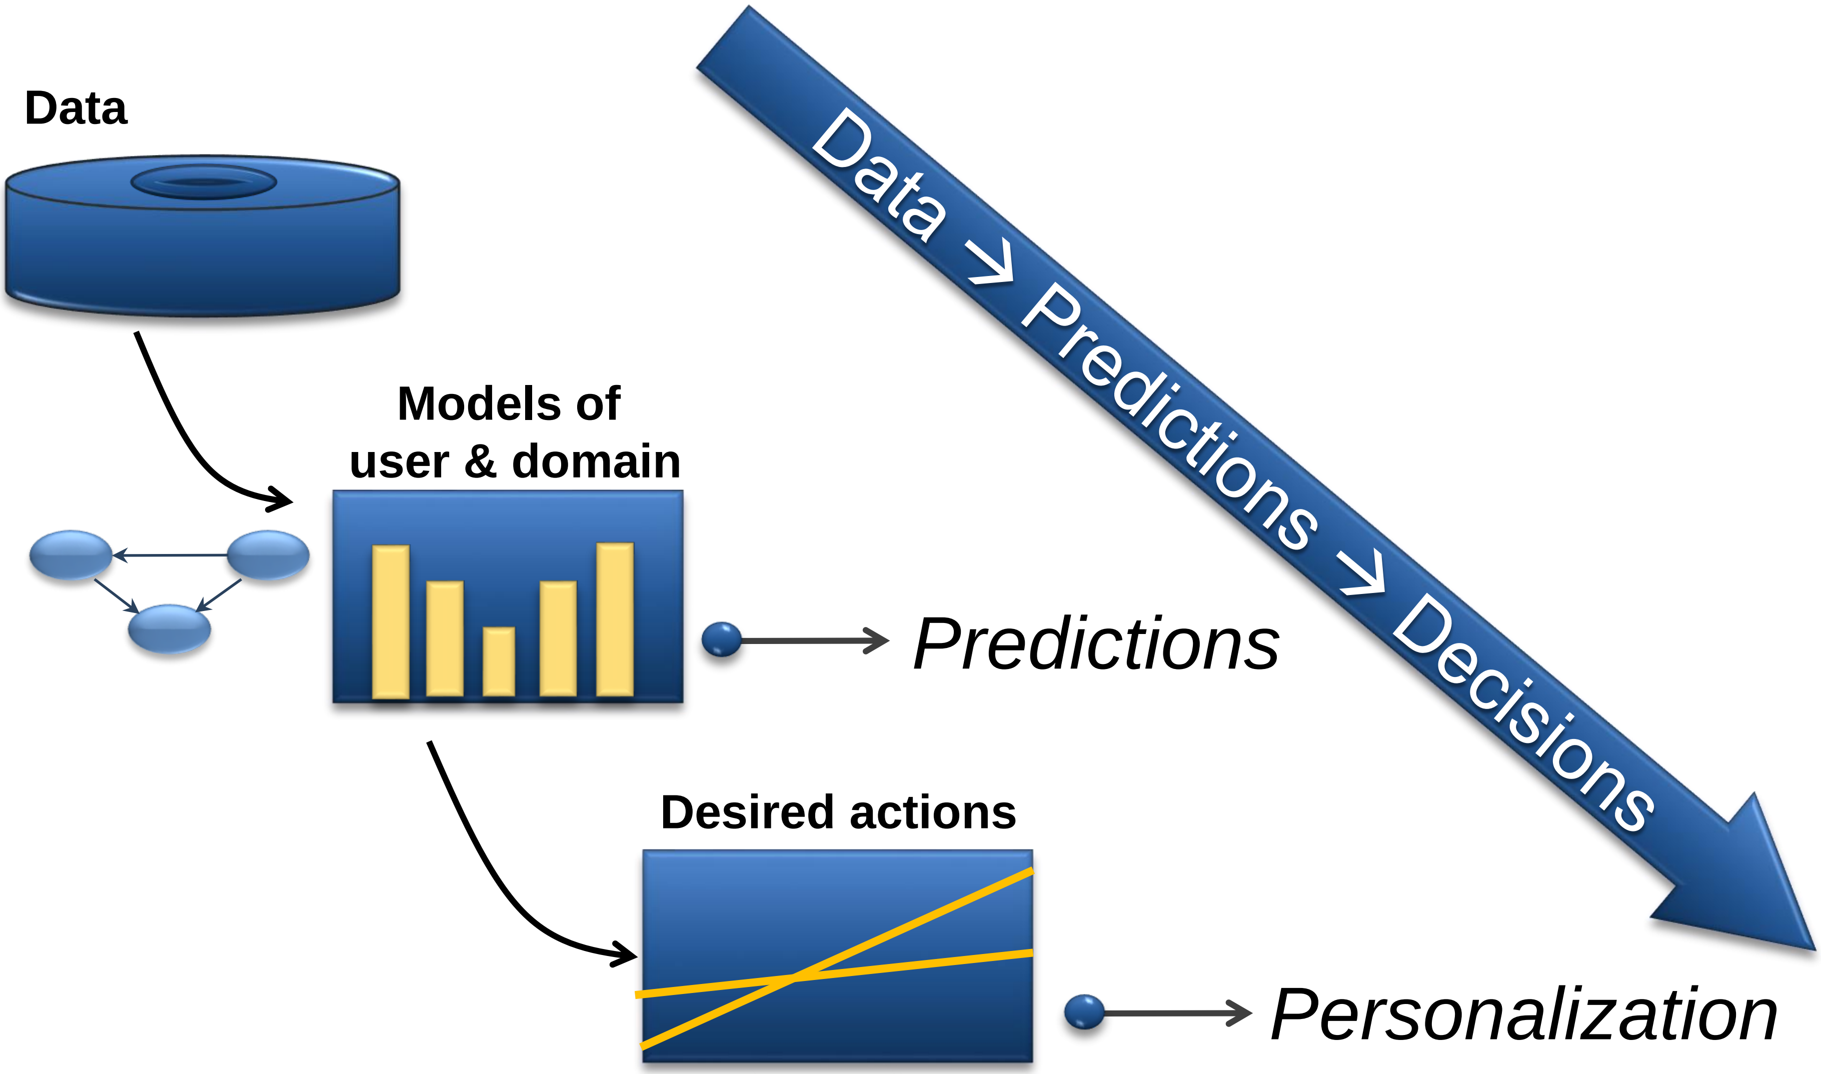

Data

Models of user & domain

Predictions

Desired actions

Personalization

Data → Predictions → Decisions

AI for Minimally-Invasive Sensing

Needs → *Consider information value & sensitivity*

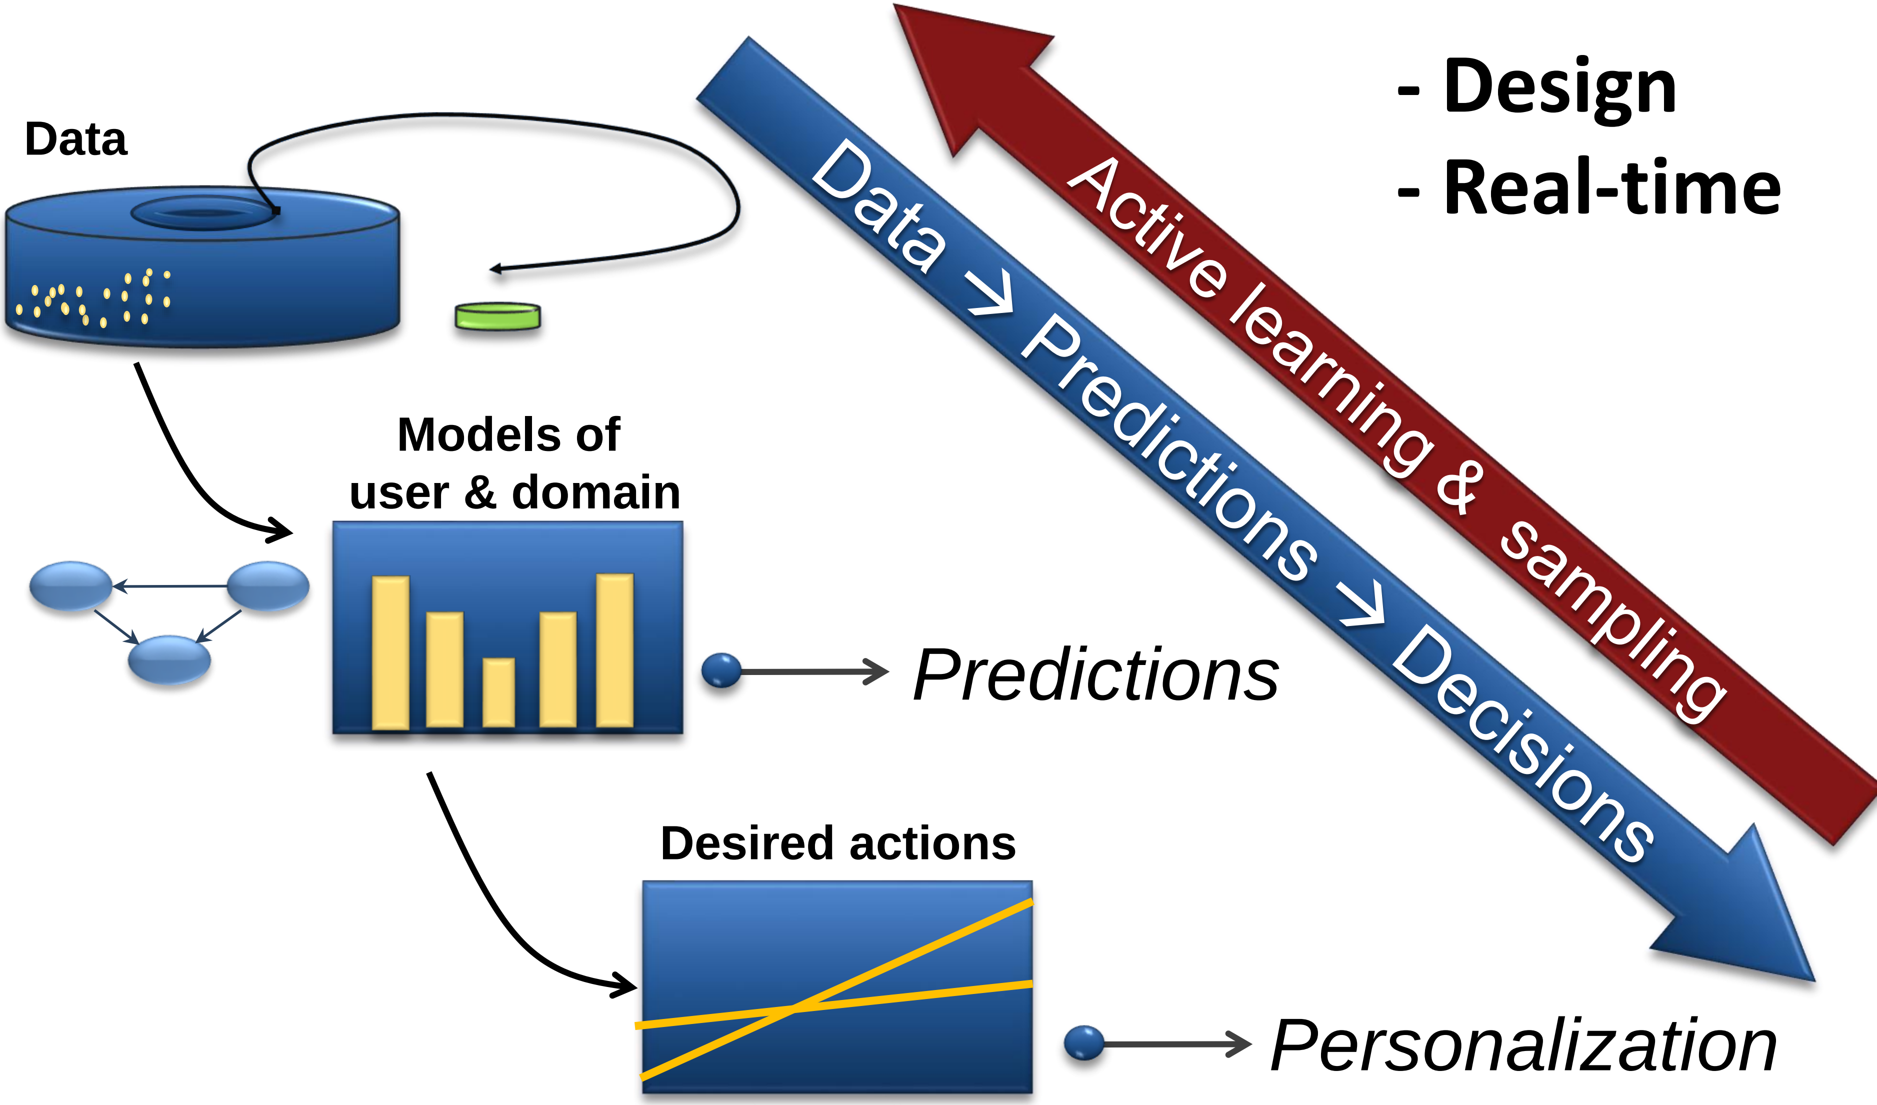

Data

Models of user & domain

Predictions

Desired actions

Personalization

Data → Predictions → Decisions

Active learning & sampling

- **Design**
- **Real-time**

I. Personalization—Privacy Tradeoffs

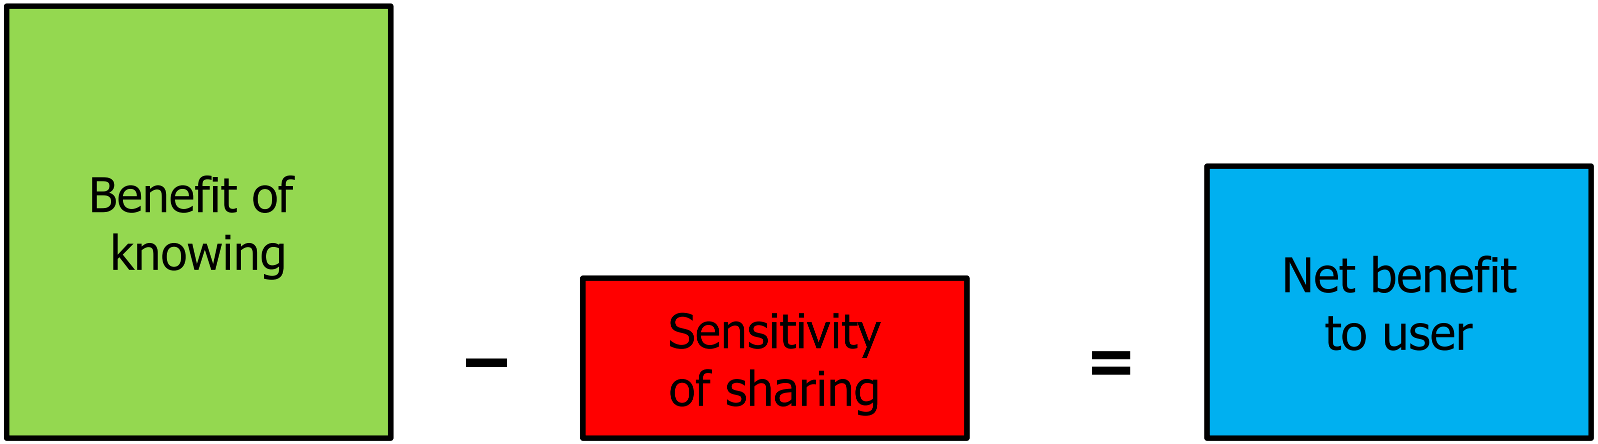

Sharing personal data (demographics, interests, activity)

with Andreas Krause

Access paper

I. Personalization—Privacy Tradeoffs

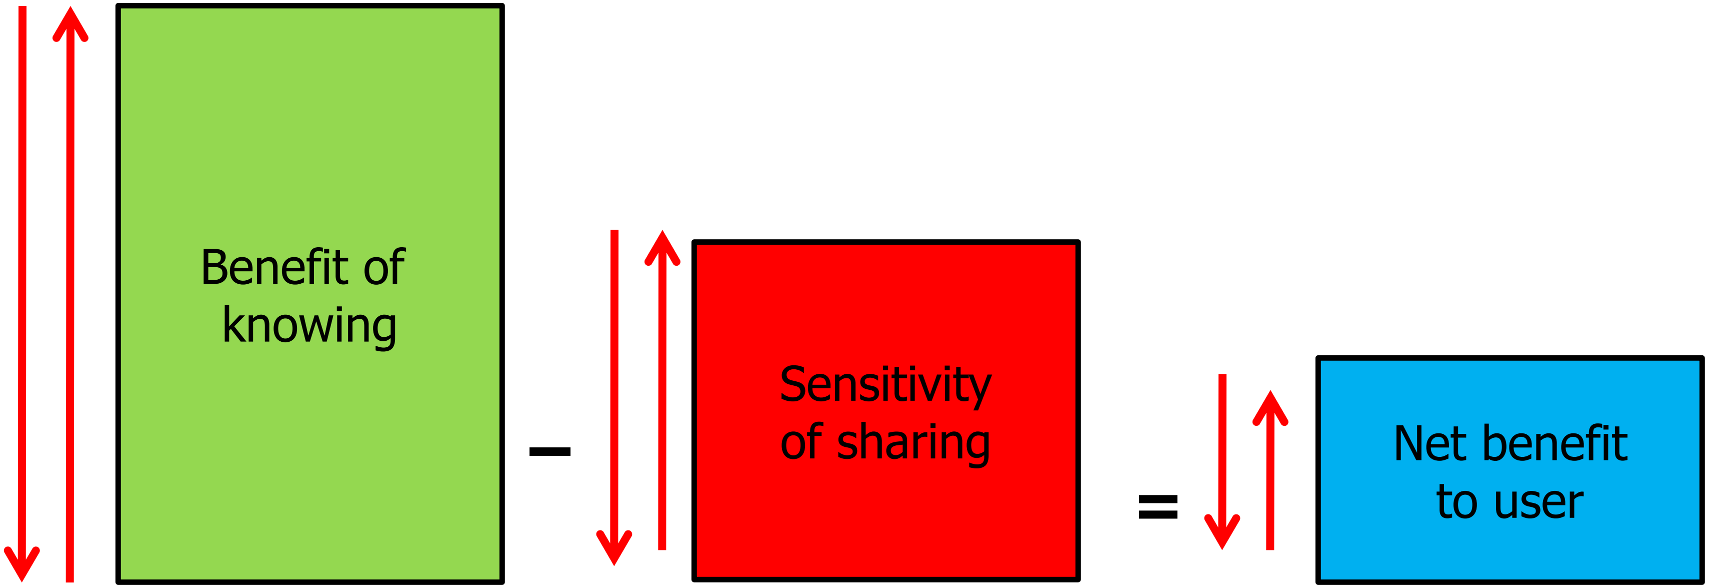

Sharing more information might decrease net benefit

with Andreas Krause

Access paper

Personalization—Privacy Study

Web search: ~15,000 users, ~250,000 queries

User data can reduce uncertainty about info needs

Query: "*sports*"

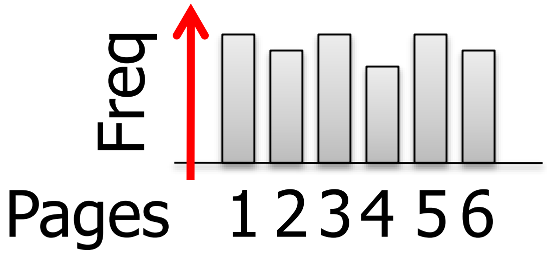

Entropy H = 2.6

Country: USA

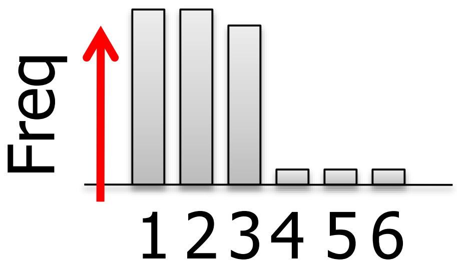

H = 1.7

Uncertainty reduction: 0.9

Label	Type	bits	Description
DGDR	Demographic	1	Gender
DAGE	Demographic	2	Age group (<18, 18-50, >50)
DOCC	Demographic	3	Occupation (6 groups of related jobs)
DREG	Demographic	2	Region (4 geographic regions)
DMTL	Demographic	1	Marital status (*)
DCHD	Demographic	1	Whether the searcher has children or not (*)
AQRY	Activity	1	Performed same query before
ACLK	Activity	1	Visited same website before
AFRQ	Activity	1	User performs at least 1 query per day on average
AZIP	Activity	1	User performed queries from at least 2 different zip codes
ACTY	Activity	1	User performed queries from at least 2 different cities
ACRY	Activity	1	User performed queries from at least 2 different countries
AWHR	Activity	1	Current query performed during working hours
AWDY	Activity	1	Current query performed during workday / weekend
ATLV	Activity	2	Top-level domain of query IP address (.com, .net, .org, .edu)
TART	Topic	1	User previously visited arts related webpage
TADT	Topic	1	User previously visited webpage with adult content
TBUS	Topic	1	User previously visited business related webpage
TCMP	Topic	1	User previously visited compute related webpage
TGMS	Topic	1	User previously visited games related webpage
THEA	Topic	1	User previously visited health related webpage
THOM	Topic	1	User previously visited home related webpage
TKID	Topic	1	User previously visited kids / teens related webpage

Understanding Sensitivities

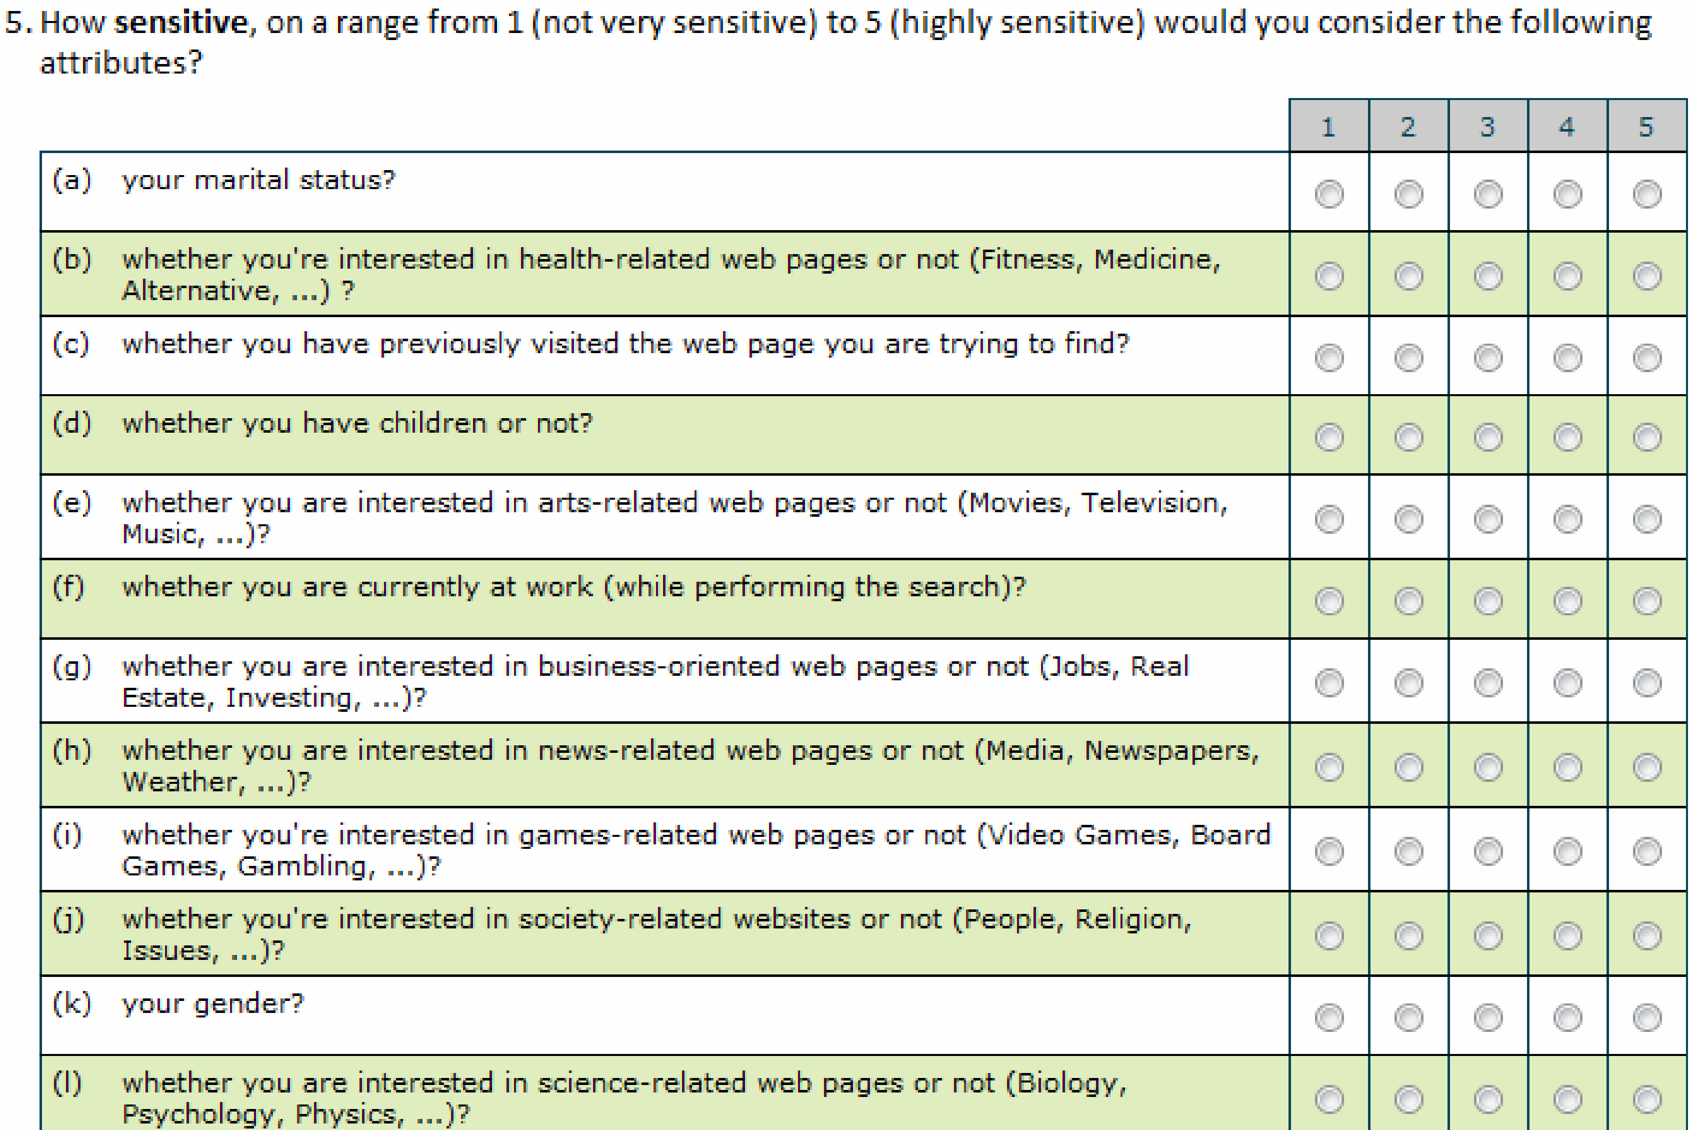

5. How **sensitive**, on a range from 1 (not very sensitive) to 5 (highly sensitive) would you consider the following attributes?

		1	2	3	4	5
(a)	your marital status?	○	○	○	○	○
(b)	whether you're interested in health-related web pages or not (Fitness, Medicine, Alternative, ...) ?	○	○	○	○	○
(c)	whether you have previously visited the web page you are trying to find?	○	○	○	○	○
(d)	whether you have children or not?	○	○	○	○	○
(e)	whether you are interested in arts-related web pages or not (Movies, Television, Music, ...)?	○	○	○	○	○
(f)	whether you are currently at work (while performing the search)?	○	○	○	○	○
(g)	whether you are interested in business-oriented web pages or not (Jobs, Real Estate, Investing, ...)?	○	○	○	○	○
(h)	whether you are interested in news-related web pages or not (Media, Newspapers, Weather, ...)?	○	○	○	○	○
(i)	whether you're interested in games-related web pages or not (Video Games, Board Games, Gambling, ...)?	○	○	○	○	○
(j)	whether you're interested in society-related websites or not (People, Religion, Issues, ...)?	○	○	○	○	○
(k)	your gender?	○	○	○	○	○
(l)	whether you are interested in science-related web pages or not (Biology, Psychology, Physics, ...)?	○	○	○	○	○

Understanding Sensitivities

Working hour query

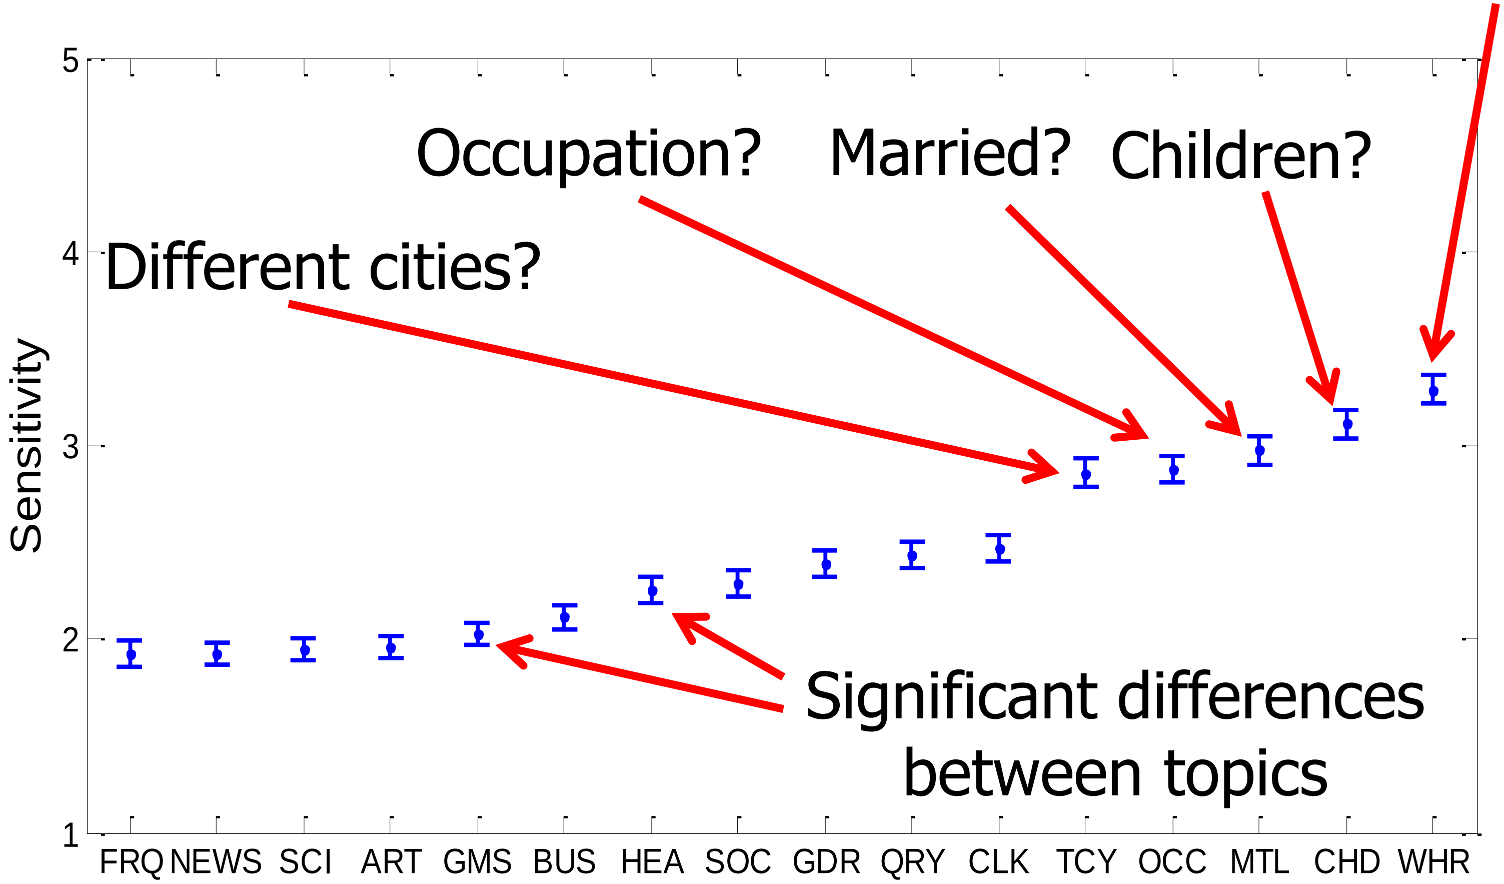

Occupation? Married? Children?

Different cities?

Significant differences between topics

Sensitivity about Location Resolution

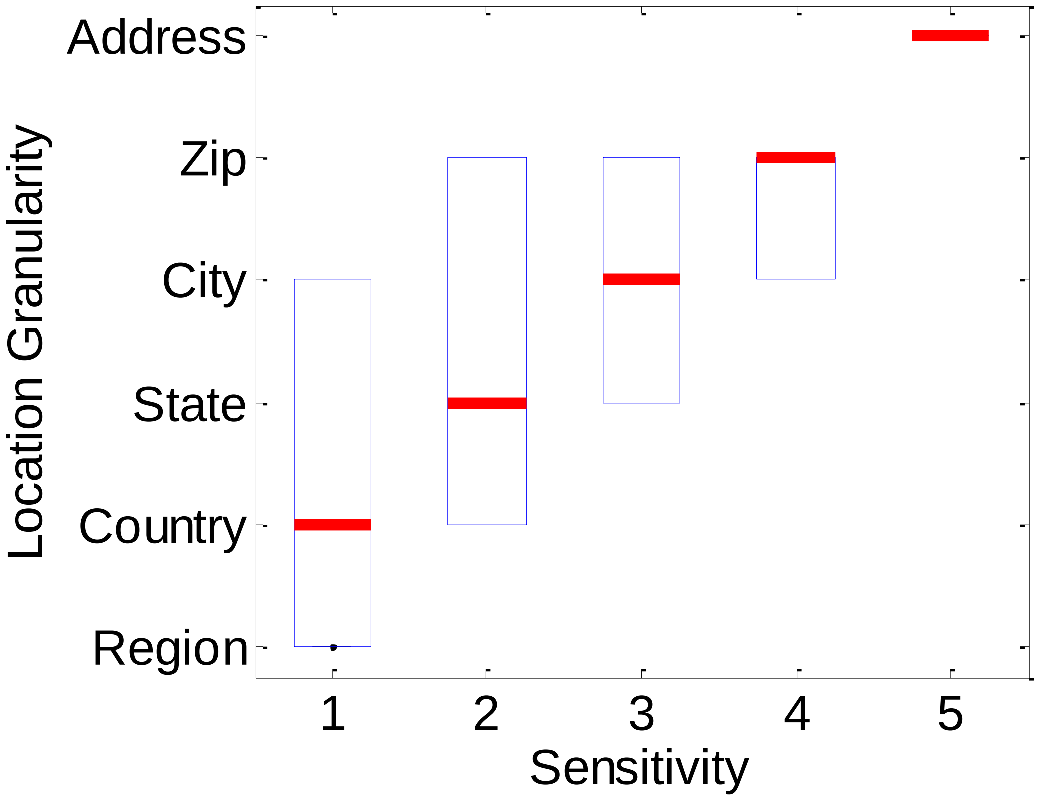

Sensitivity vs utility of enhanced service

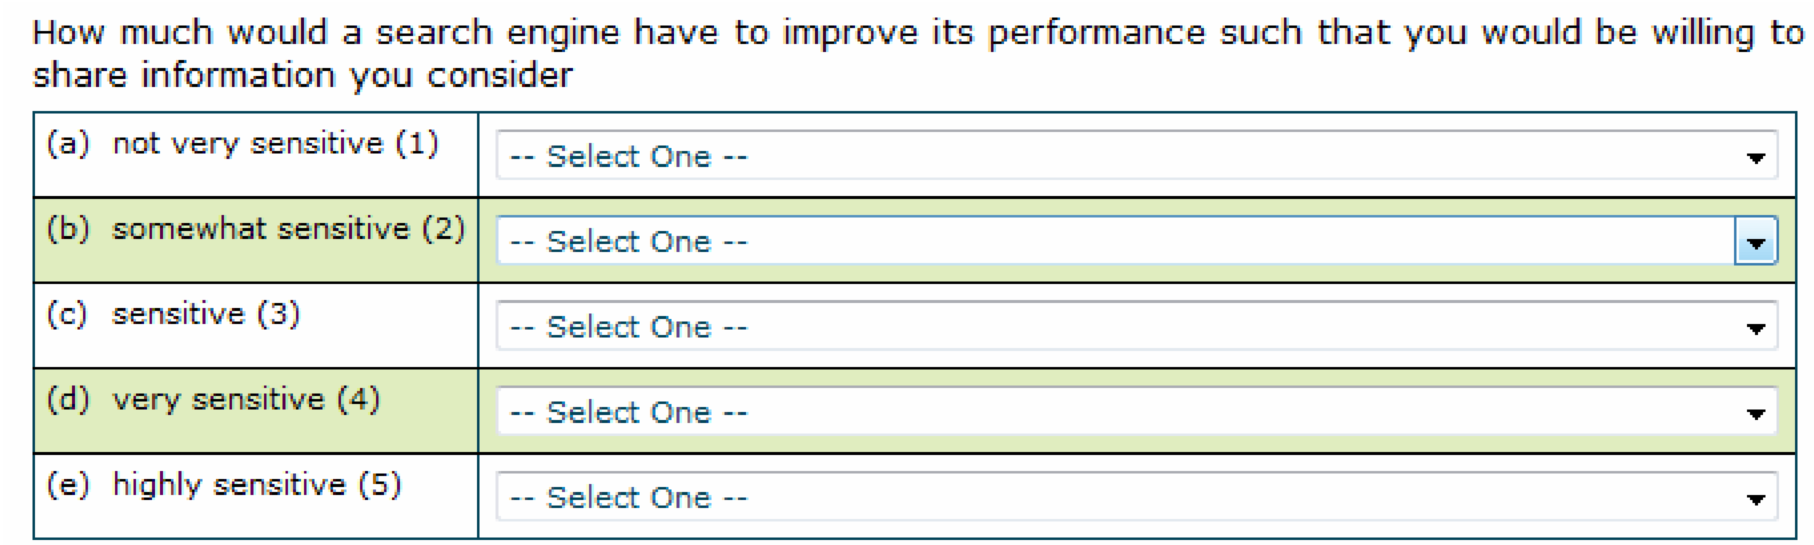

How much would a search engine have to improve its performance such that you would be willing to share information you consider

(a) not very sensitive (1)	-- Select One -- ▾
(b) somewhat sensitive (2)	-- Select One -- ▾
(c) sensitive (3)	-- Select One -- ▾
(d) very sensitive (4)	-- Select One -- ▾
(e) highly sensitive (5)	-- Select One -- ▾

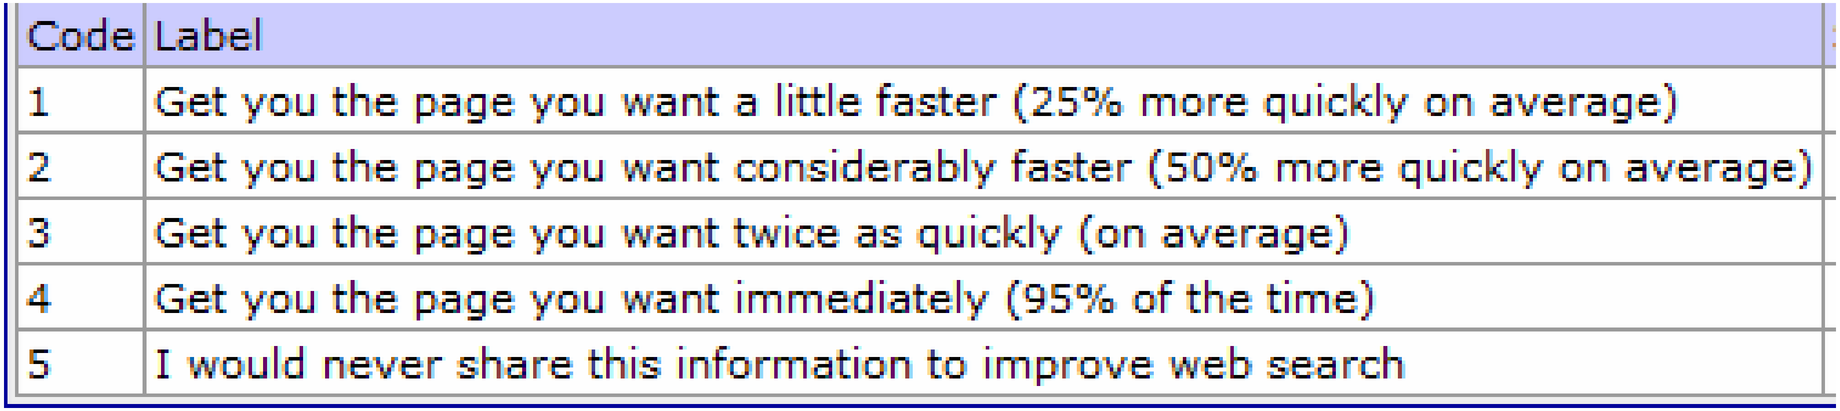

Code	Label
1	Get you the page you want a little faster (25% more quickly on average)
2	Get you the page you want considerably faster (50% more quickly on average)
3	Get you the page you want twice as quickly (on average)
4	Get you the page you want immediately (95% of the time)
5	I would never share this information to improve web search

User data and personalization

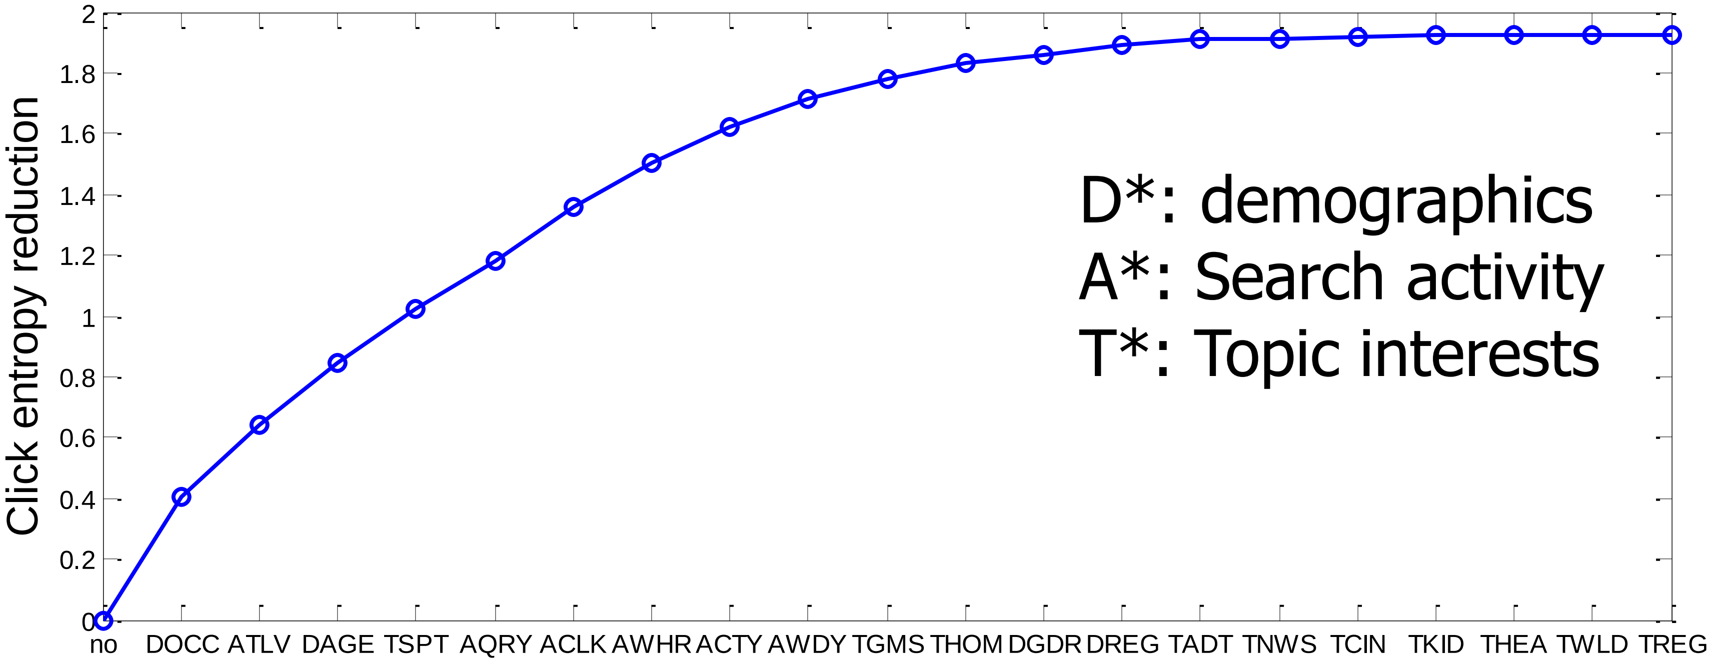

D*: demographics
A*: Search activity
T*: Topic interests

Web search study: ~15,000 users, ~250,000 queries

Cost of increasing identifiability

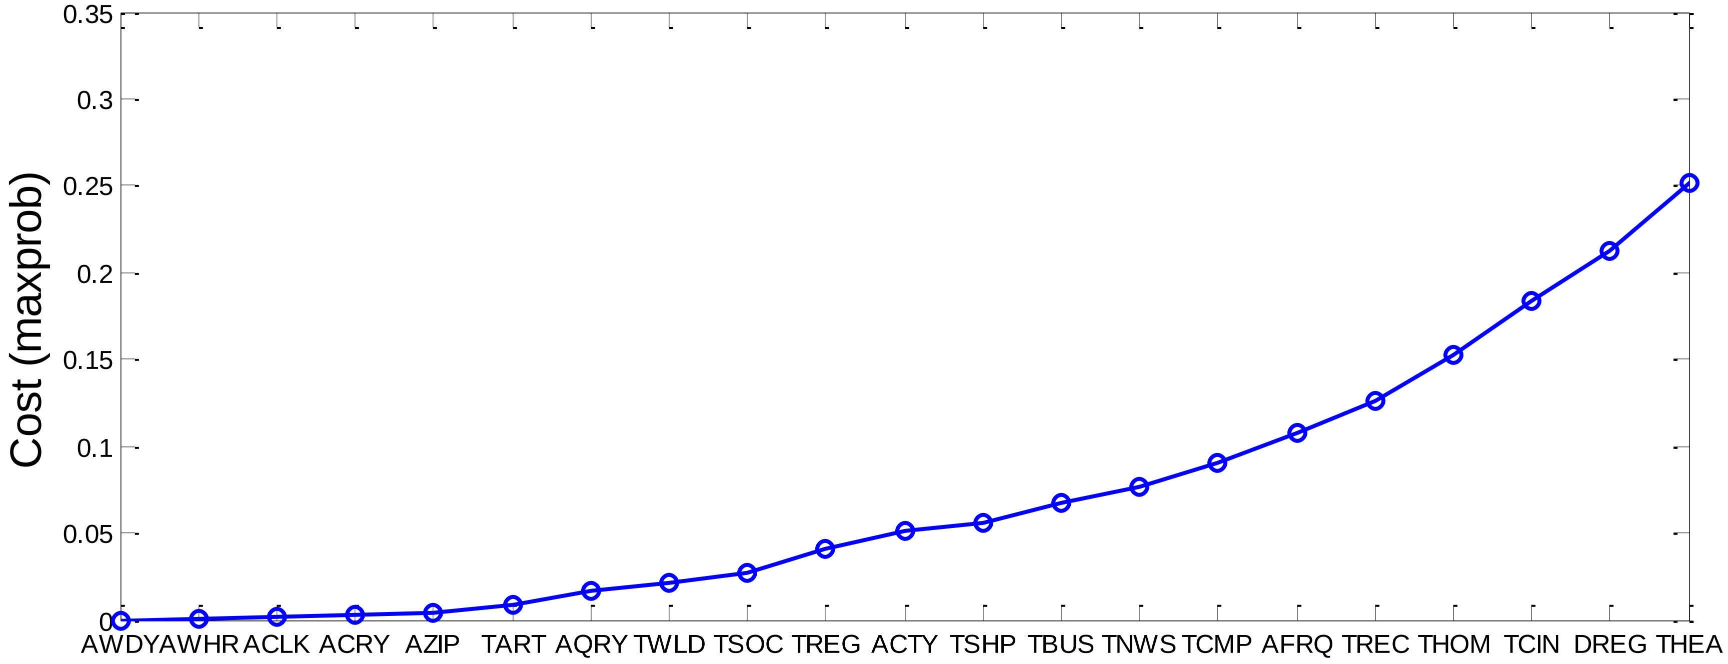

Web search study: ~15,000 users, ~250,000 queries

Optimization

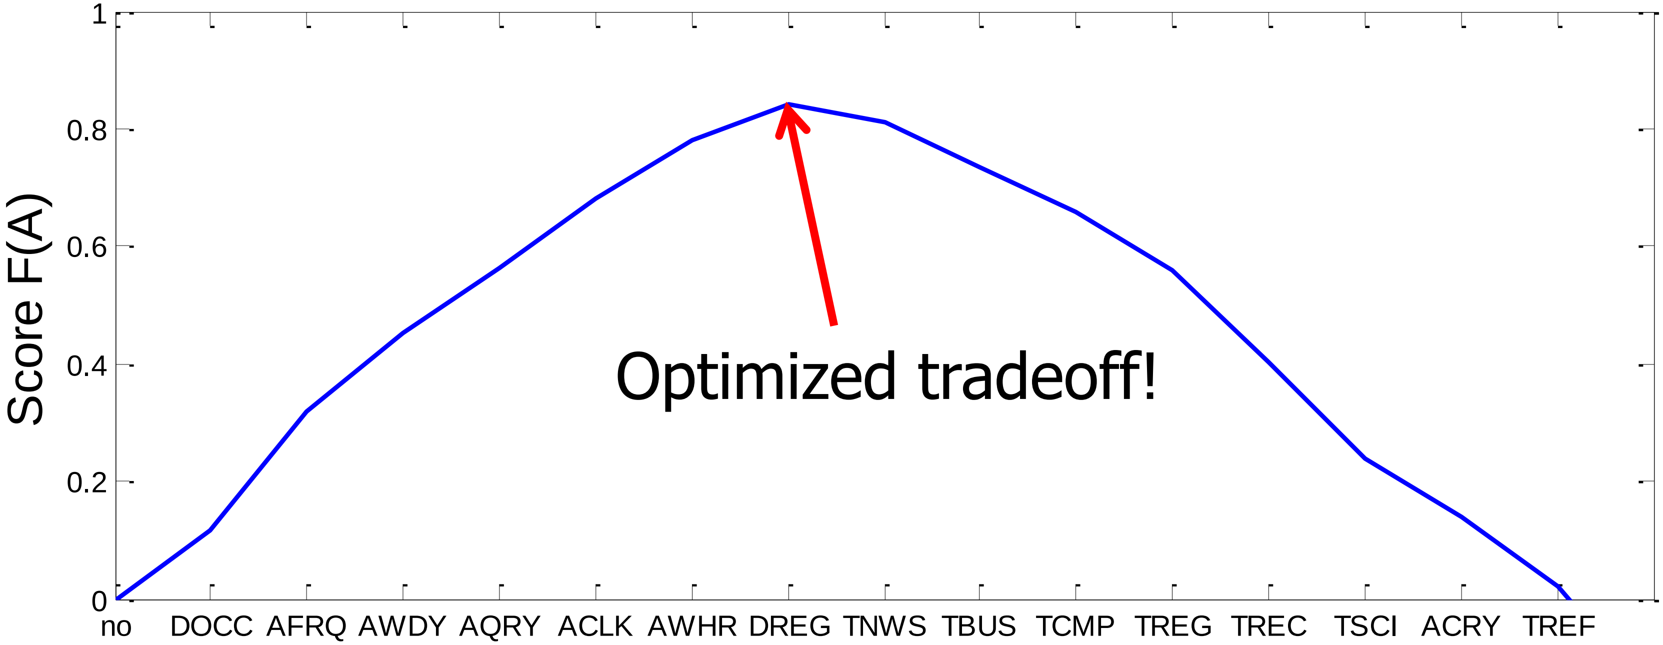

Optimized tradeoff!

Decisions and Tradeoffs

Value:
Diminishing returns

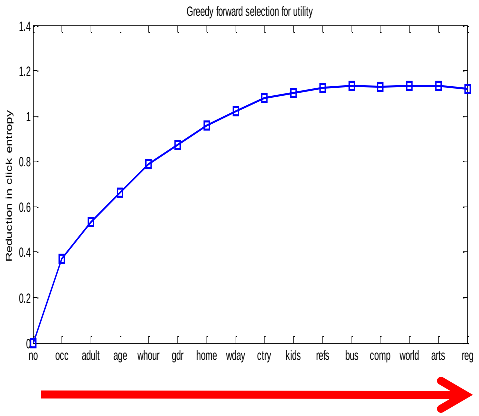

More observations

- λ

Cost:
Accelerating

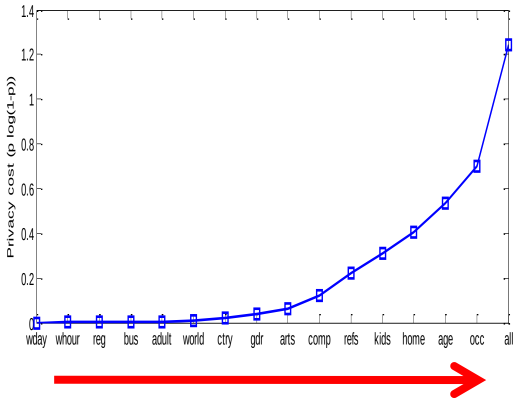

More observations

=

Optimization

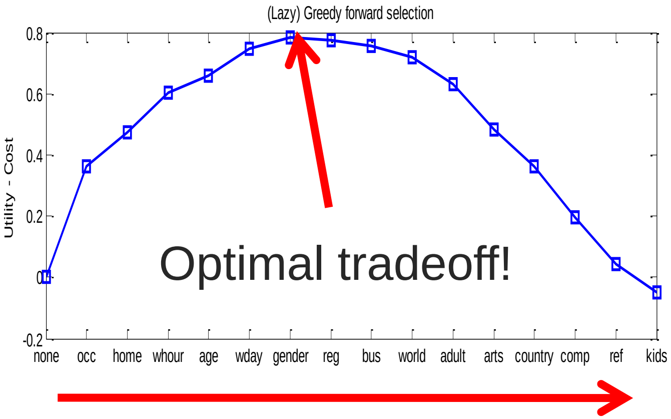

Optimal tradeoff!

More observations

Optimization

- Repeated visit
- Query workday/weekend
- Query working hour
- Country
- Top-level domain
- Avg. queries per day

II. Community Sensing

Anonymized data from volunteers

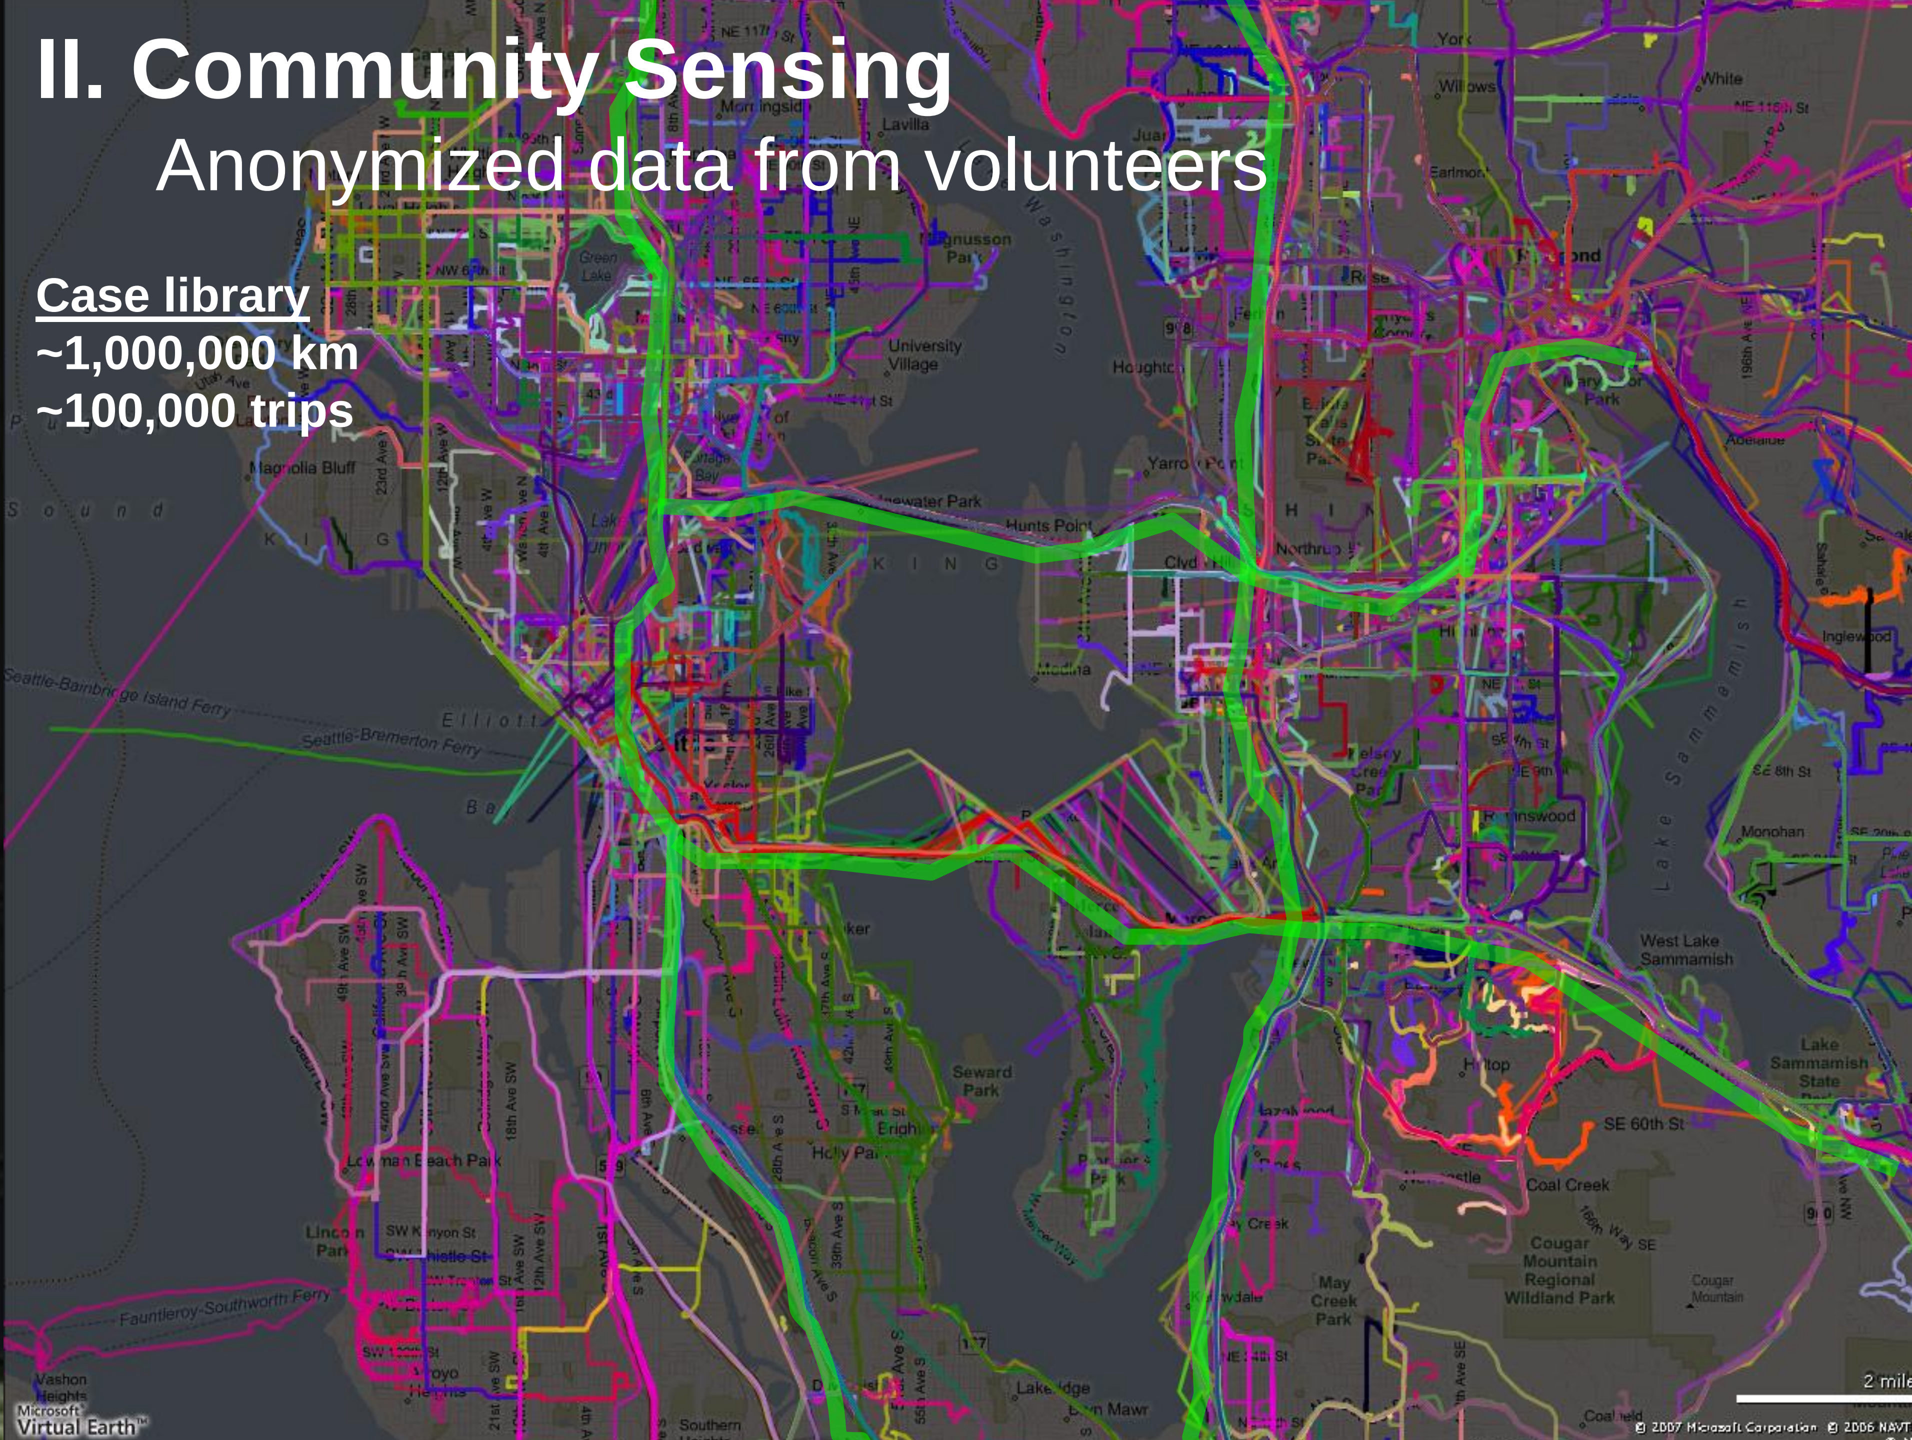

Case library

~1,000,000 km

~100,000 trips

Community Sensing

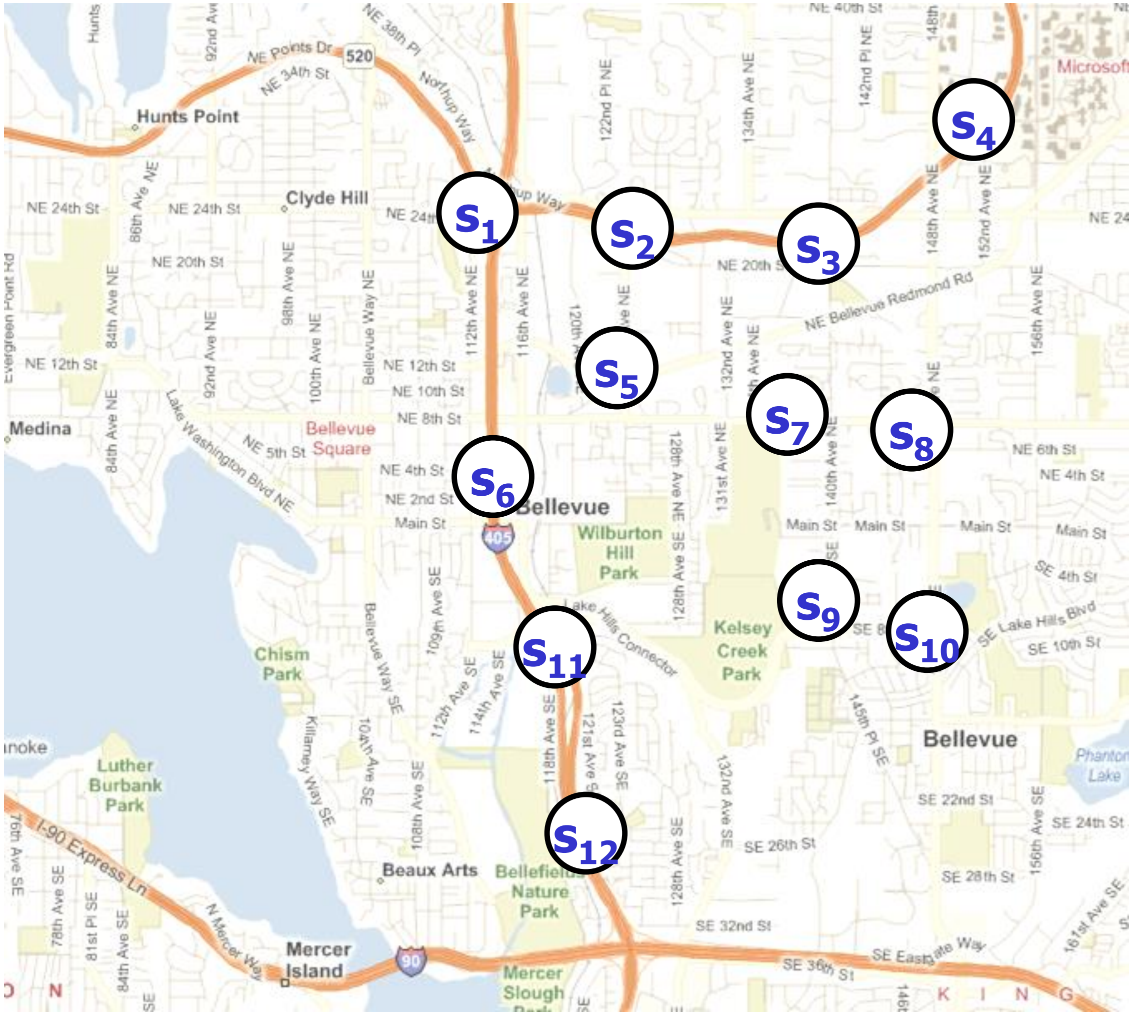

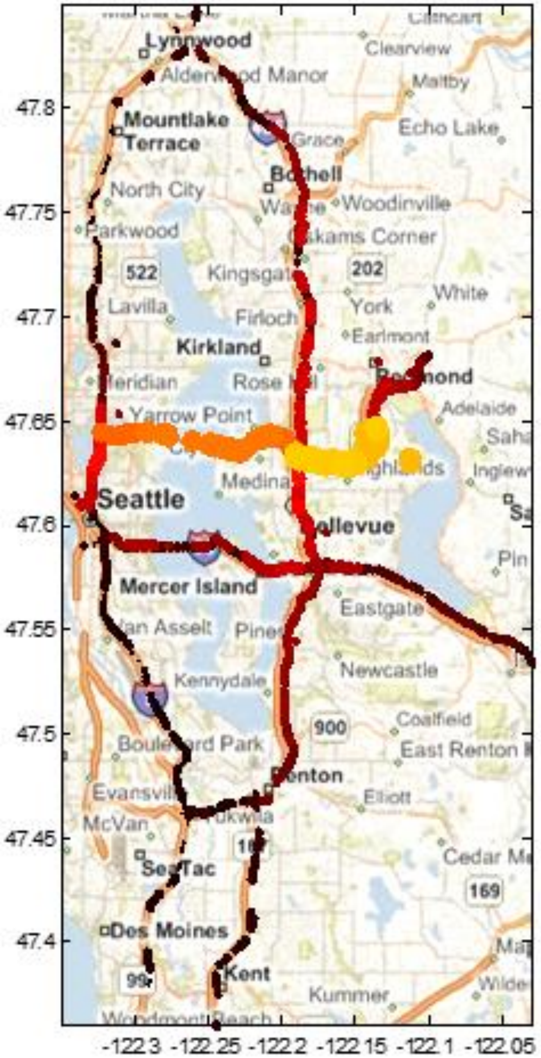

with A. Krause, A. Kansal, F. Zhao

Access paper

Community Sensing

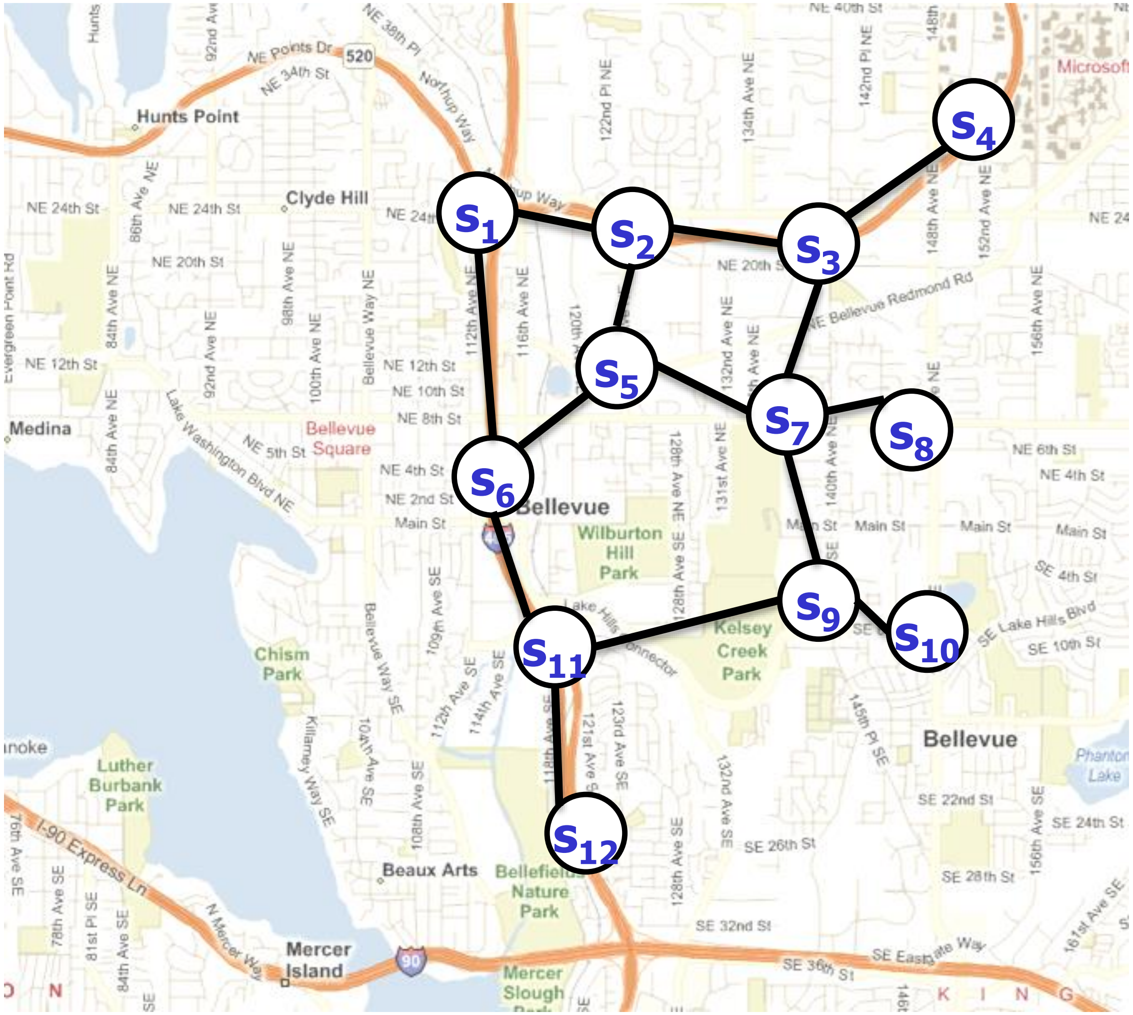

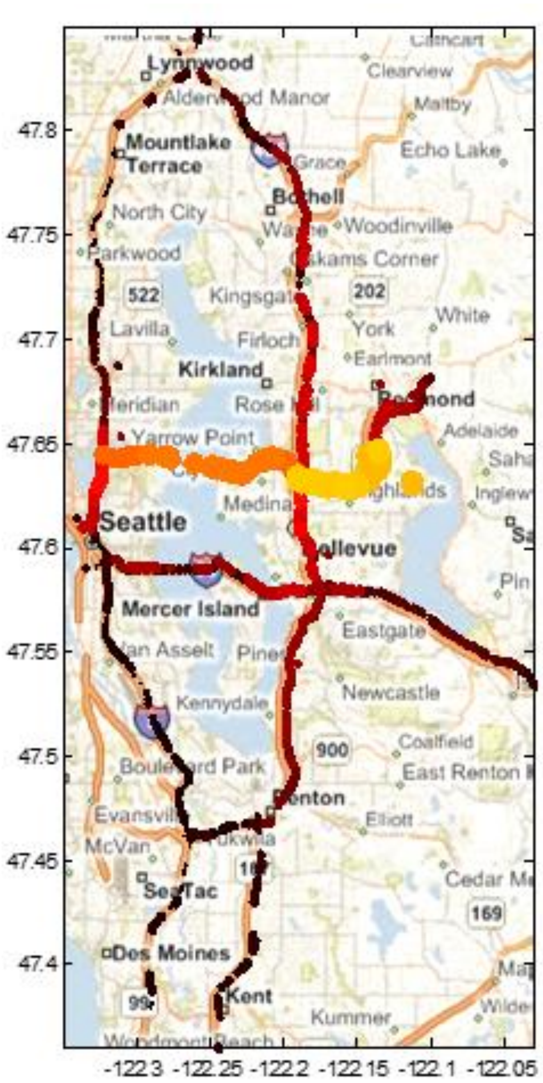

with A. Krause, A. Kansal, F. Zhao

Access paper

Community Sensing

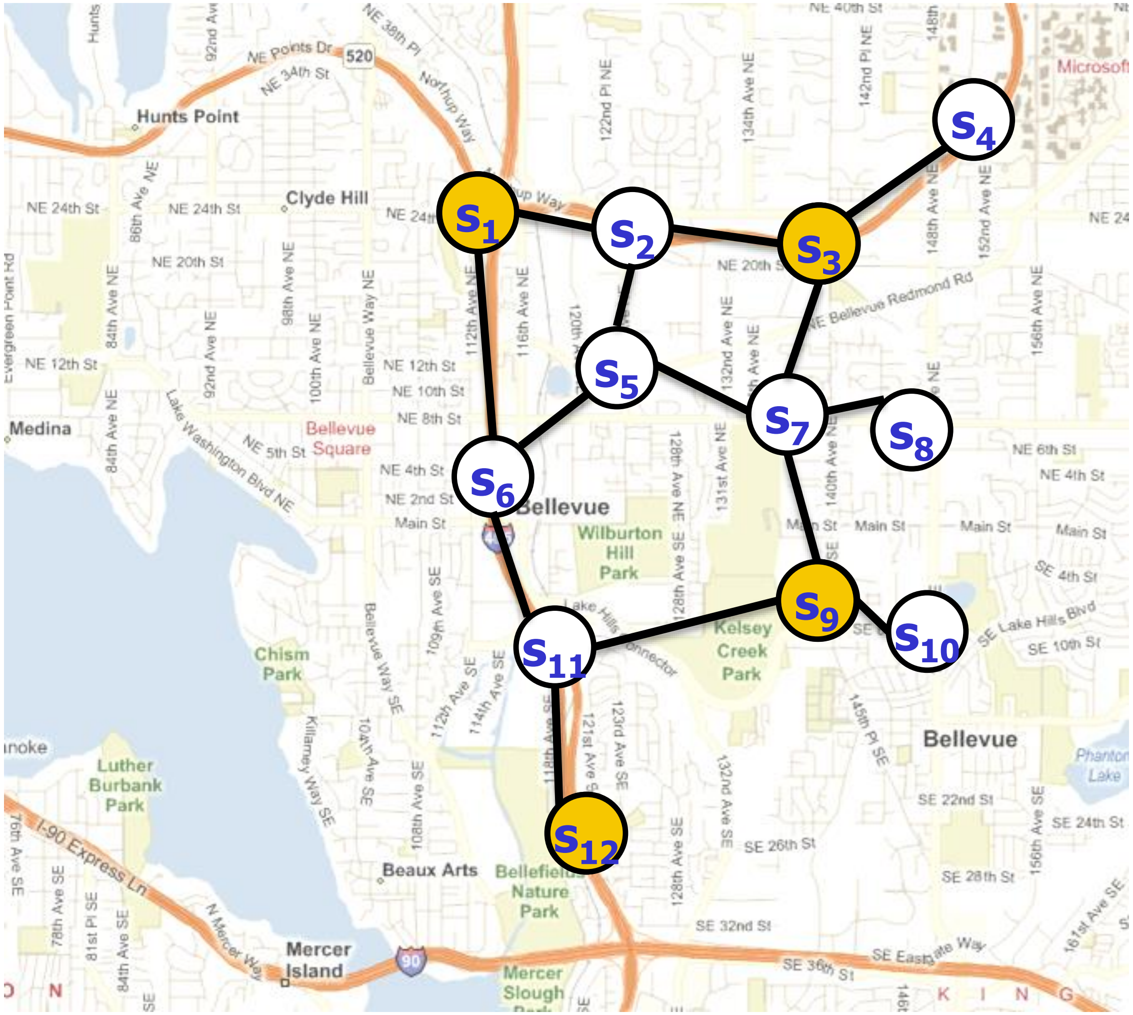

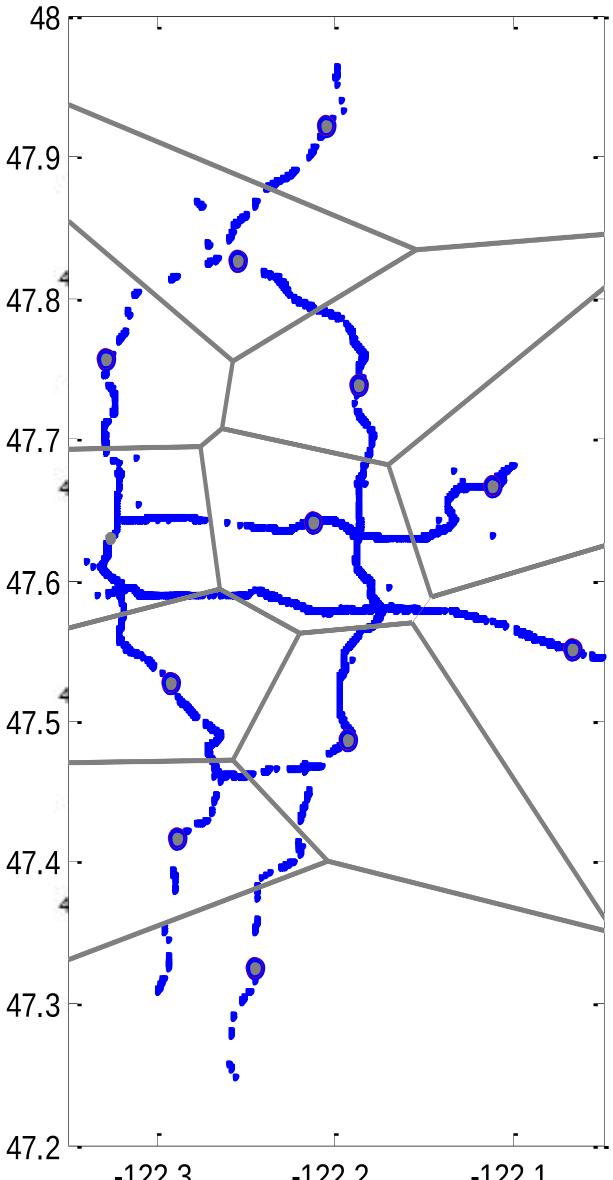

with A. Krause, A. Kansal, F. Zhao

Community Sensing

Utilitarian: Contribute for good of larger population

Phenomenon Model

Spatiotemporal process

Uncertainties, value of sensing

Demand Model

Population needs

Distribution of demand

Preference Model

Avail. of observations

Preferences on sharing

Access paper

Community Sensing

Utilitarian: Contribute for good of larger population

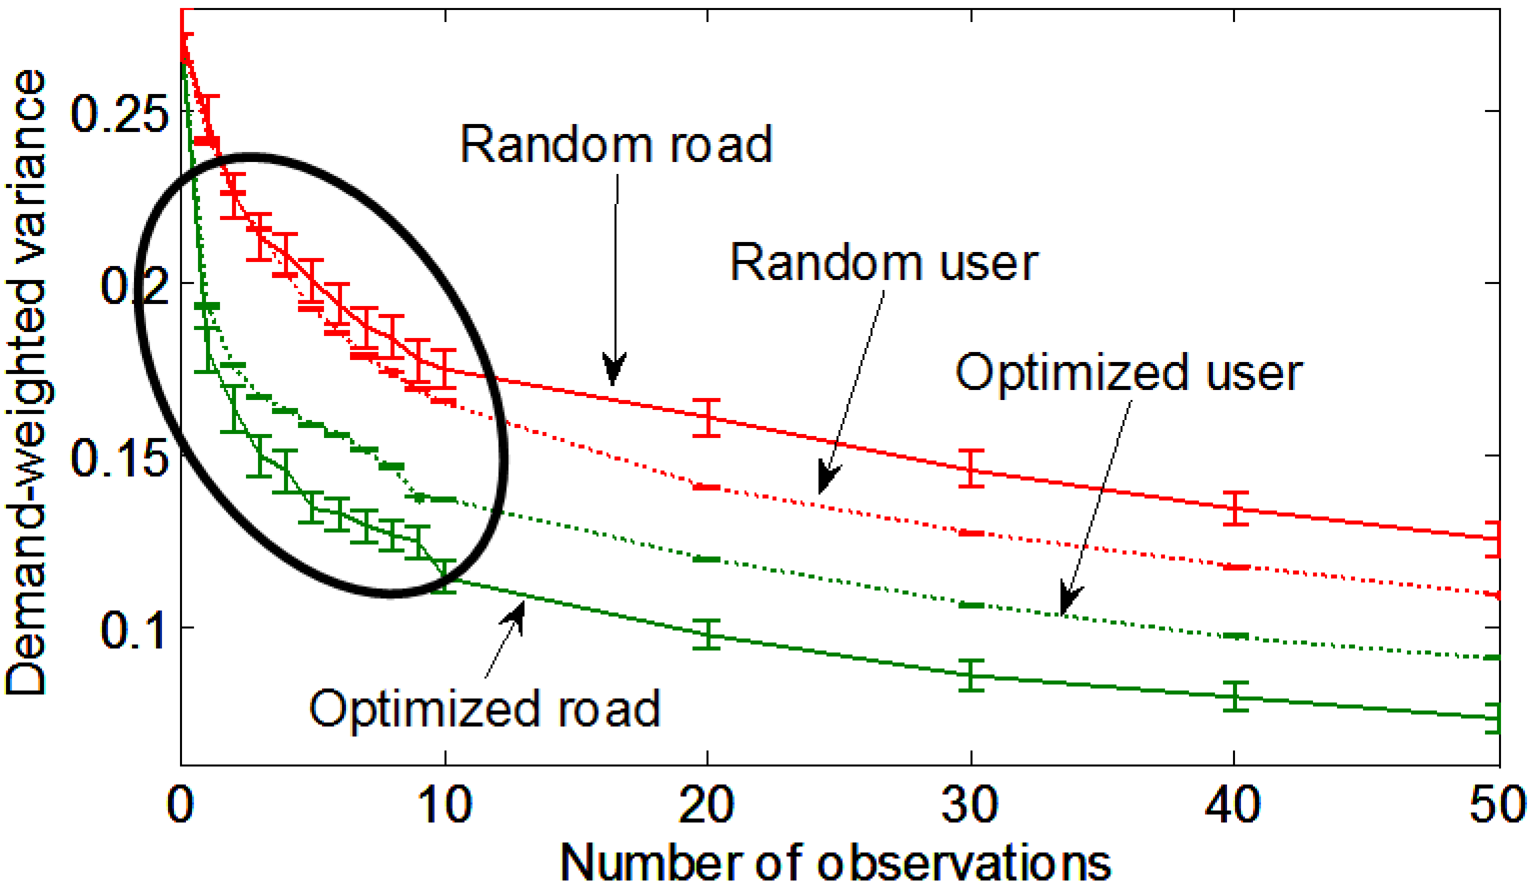

III. Stochastic Privacy

Provide bounds on small "privacy risk"

<u>System request</u>: "Please accept small *privacy risk*."→

Privacy risk: probability that some data is accessed

<u>System responsibility</u>: "We'll work within that promise."

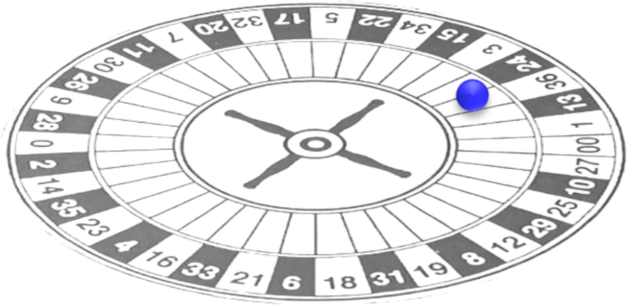

with A. Singla, E. Kamar, R. White

Access paper

Stochastic Privacy

Guaranteed bound on likelihood that data is accessed

- User’s agree to small ***privacy risk*** *r* (*e.g*, *p* < 0.000001)
- Small probabilities may be tolerable to users

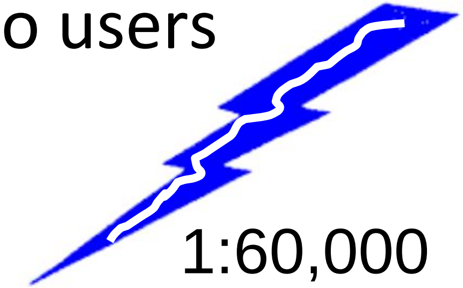

Large design space

- e.g., User’s trade higher privacy risk for incentives

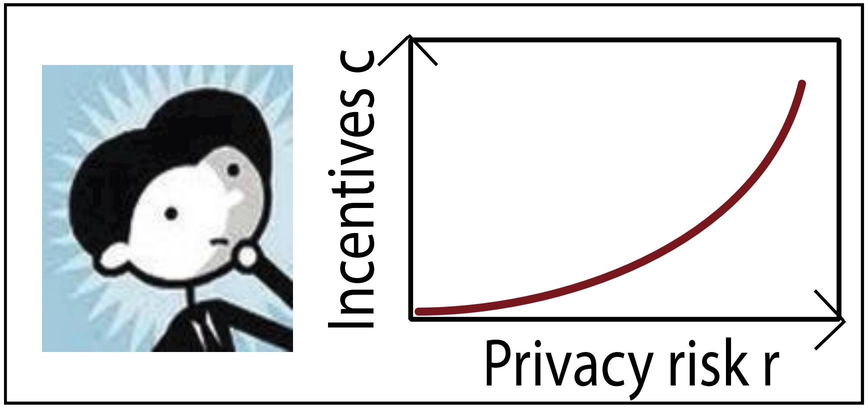

Access paper

Approach

We can identify most valuable sources of data

We can sample to guarantee bound on risk

Random sampling → ***Ideal selection***

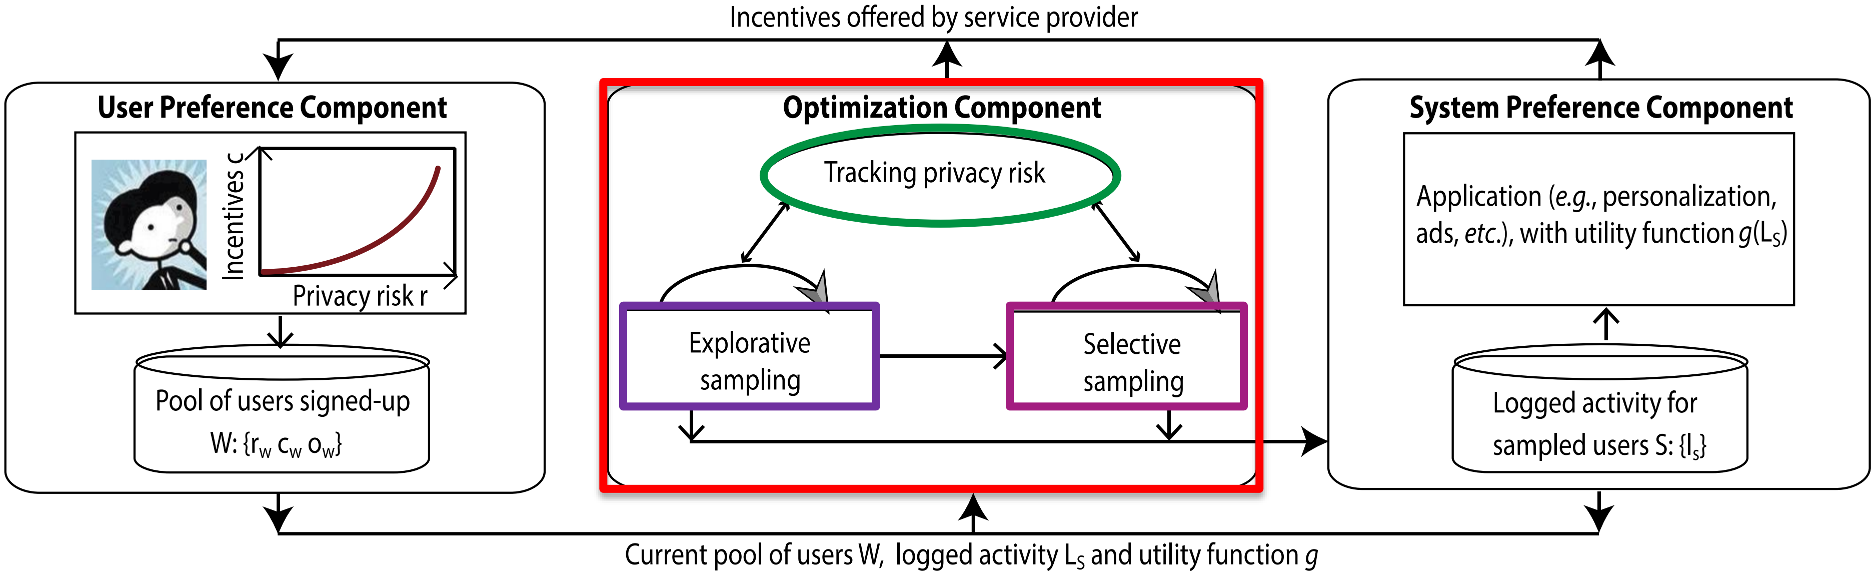
Incentives offered by service provider
User Preference Component
Incentives c
Privacy risk r
Pool of users signed-up
W: {r_w c_w o_w}
Optimization Component
Tracking privacy risk
Explorative sampling
Selective sampling
System Preference Component
Application (*e.g.*, personalization, ads, *etc.*), with utility function $g(L_S)$
Logged activity for sampled users S: {l_s}
Current pool of users W, logged activity L_S and utility function *g*

Approach

We can identify most valuable sources of data

We can sample to guarantee bound on risk

Random sampling ⟲ ***Ideal selection***

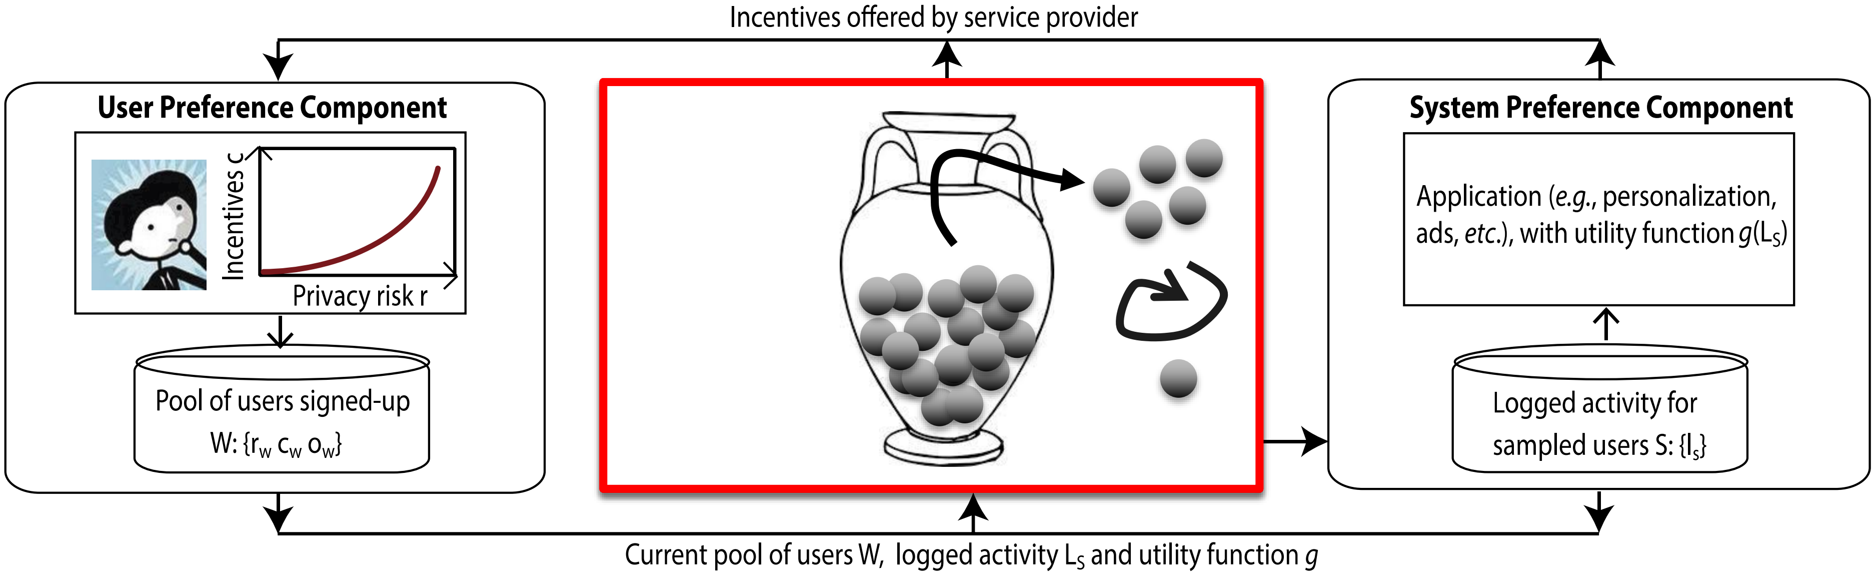

Incentives offered by service provider

User Preference Component

Pool of users signed-up

W: $\{r_W\ c_W\ o_W\}$

System Preference Component

Application (*e.g.*, personalization, ads, *etc.*), with utility function $g(L_S)$

Logged activity for sampled users S: $\{l_s\}$

Current pool of users W, logged activity L_S and utility function g

Random Greedy: Random Sample→Select Best

1. Random sample to manage privacy risk
2. Select most informative source
3. Remove others from further analysis
4. Repeat.

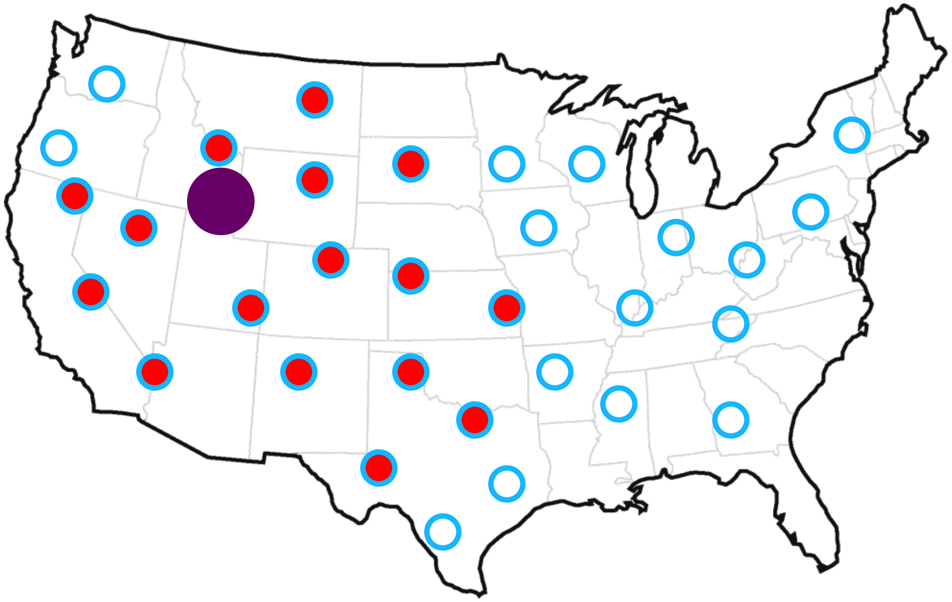

Random Greedy: Random Sample→Select Best

1. Random sample to manage privacy risk
2. Select most informative source
3. Remove others from further analysis
4. Repeat.

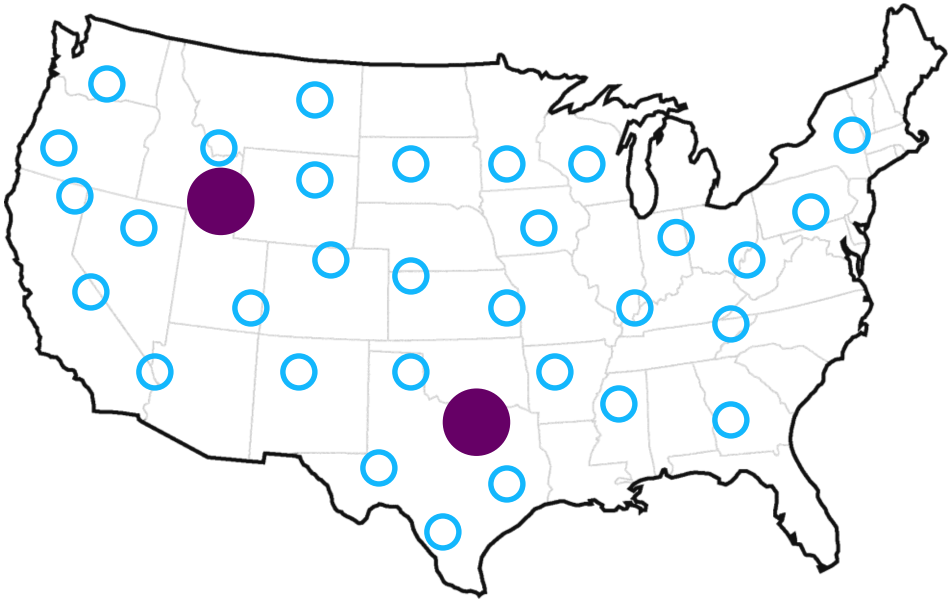

SPGreedy: Select Best→Expand→Random Sample

1. Select most informative source
2. Identify set of similar users
3. Sample single user randomly from set.
4. Repeat.

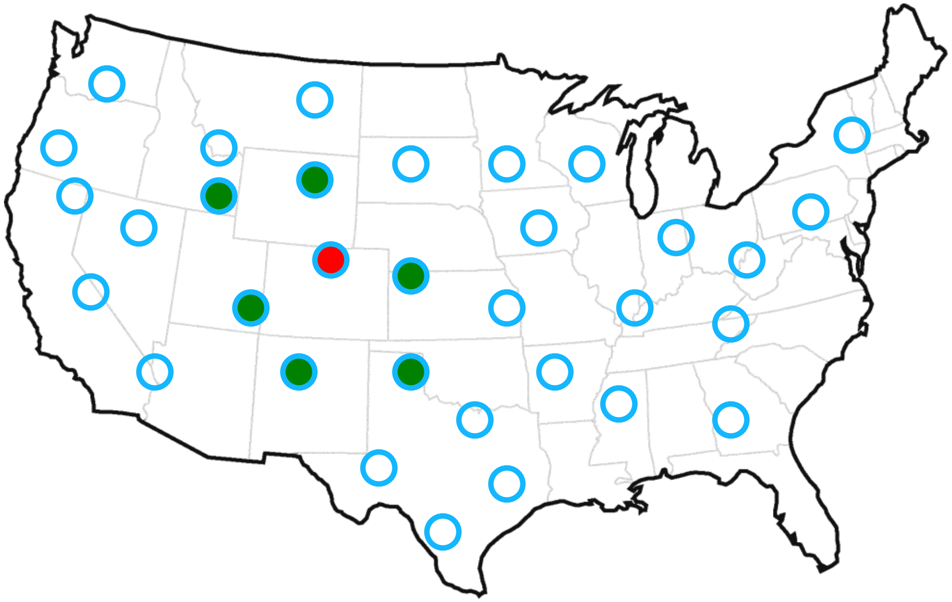

SPGreedy: Select Best→Expand→Random Sample

1. Select most informative source
2. Identify set of similar users
3. Sample single user randomly from set.
4. Repeat.

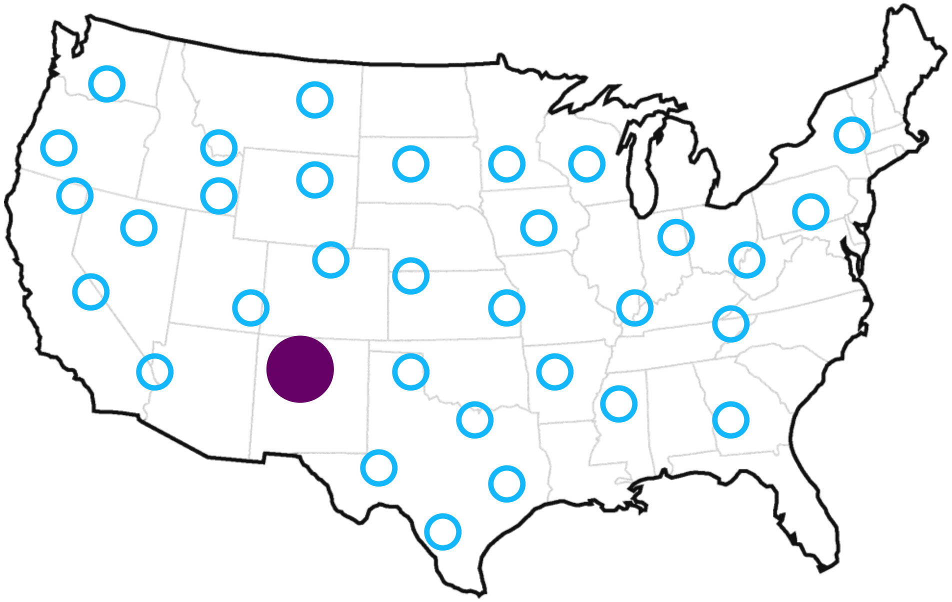

Study: Location-Based Personalization

Web search logs: Oct’2013, 10 US states

→ 7 million users

Access attributes of users prior to sampling

- Topic area: Business
- Use location data
- Last 20 result clicks (to infer expertise profile)

Results: Varying Budget

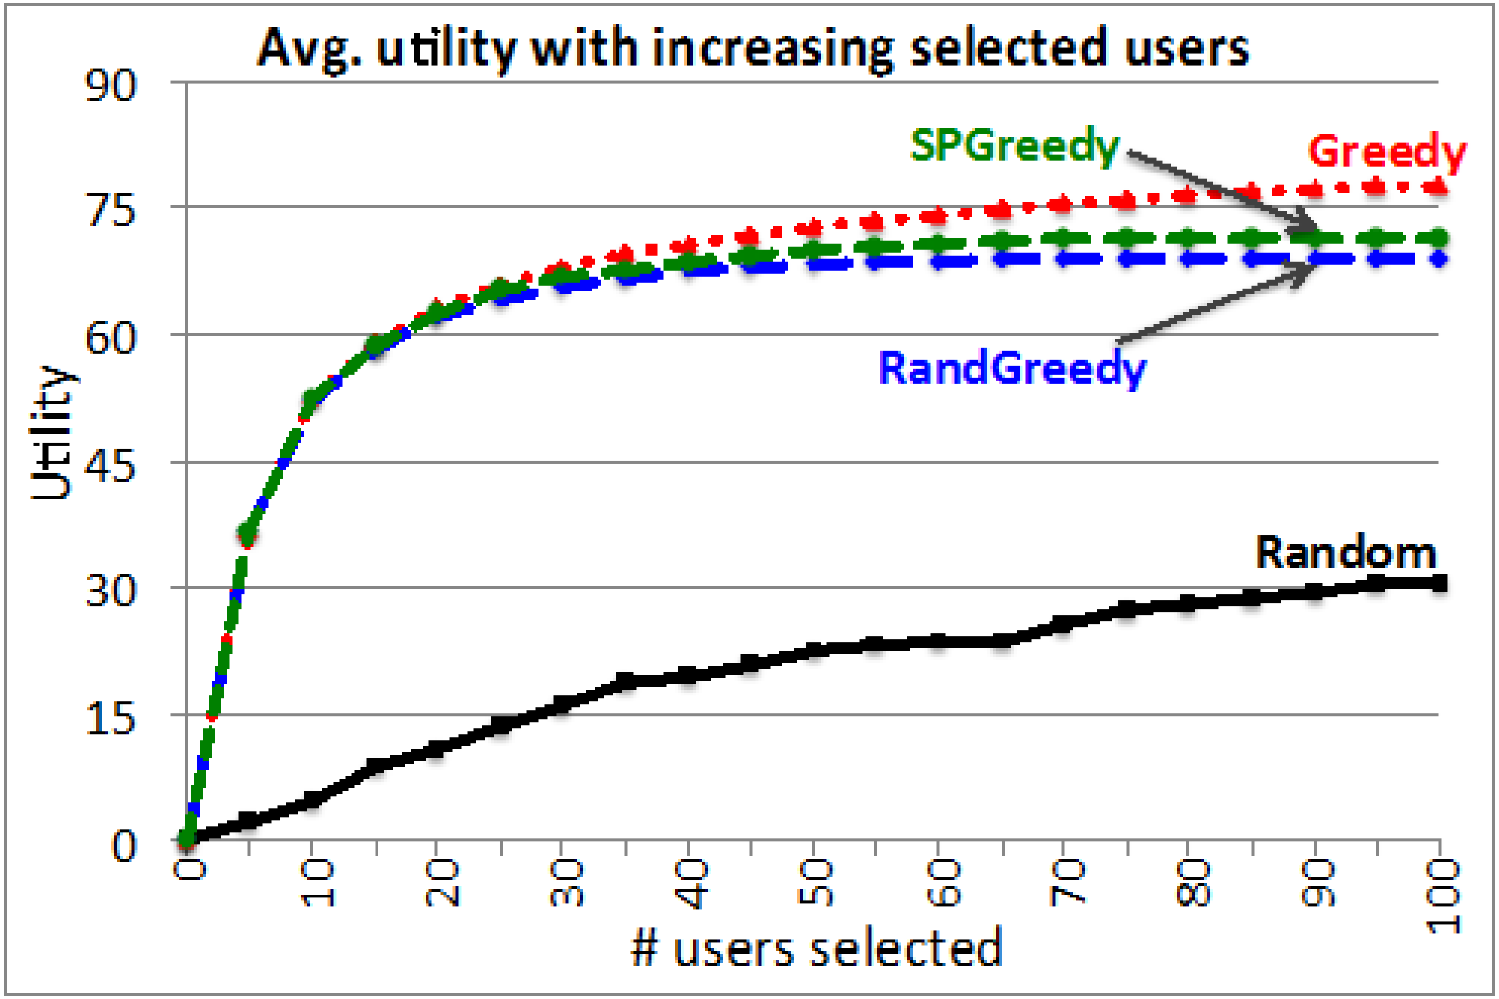

- Both RANDGREEDY and SPGREEDY are competitive w.r.t. GREEDY
- Naïve baseline RANDOM perform poorly

Results: Varying Privacy Risk

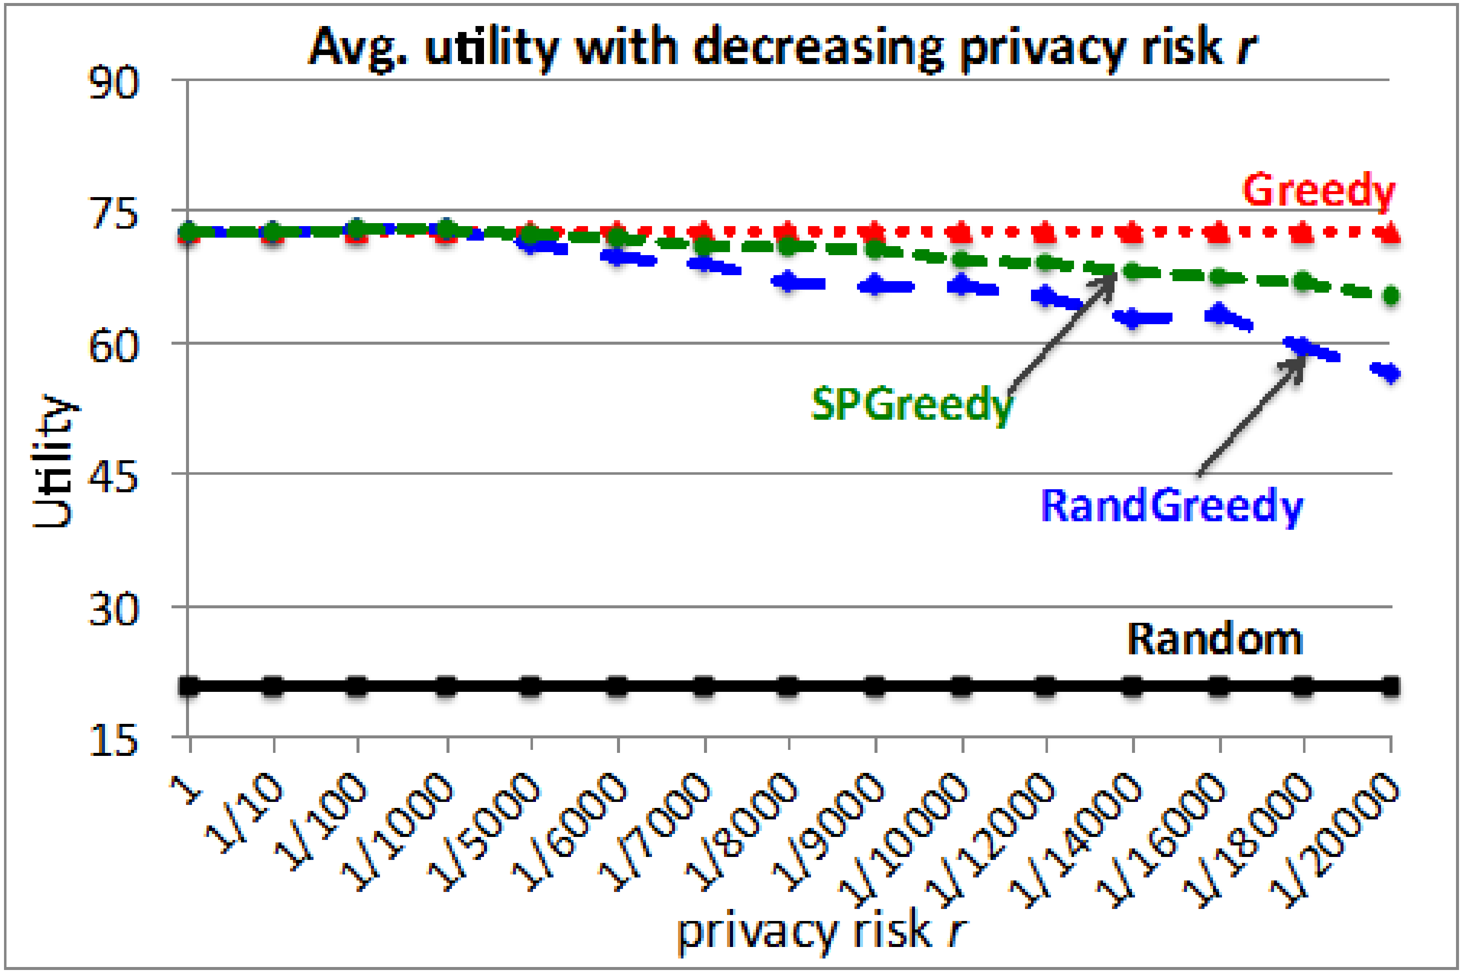

- Performance of both RANDGREEDY and SPGREEDY degrades smoothly with decreasing privacy risk (i.e. tighter sampling constraint)

Studies of Preferences

Opportunity to assess and understand conceptions about privacy—and preferences about privacy mechanisms.

e.g.,

Understanding *privacy risk*

Comfort with increasing privacy risk

Studies of Pref

Opportunity to a s about privacy ns.

(

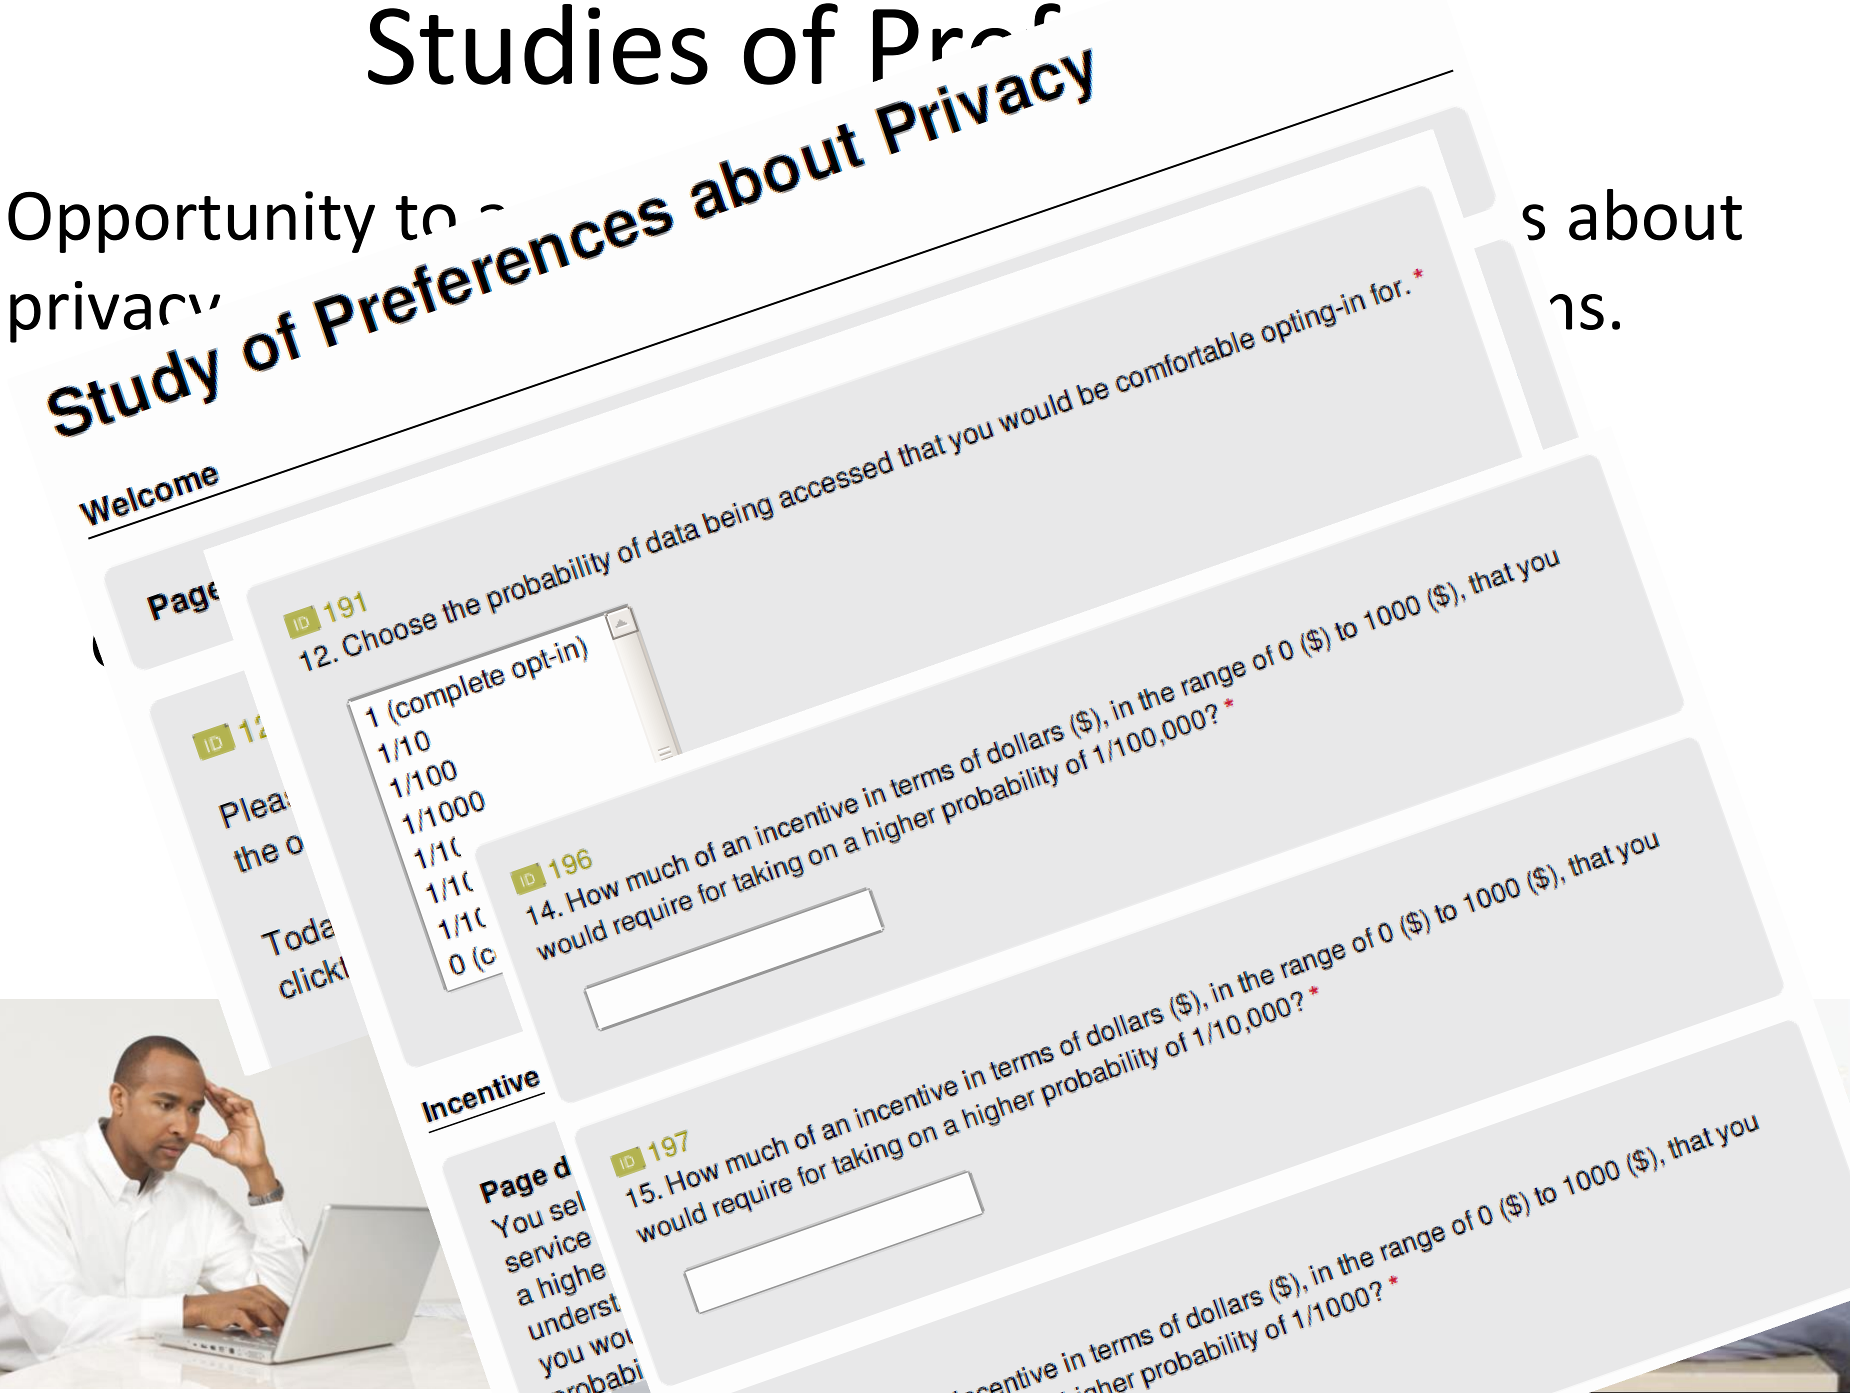

Study of Preferences about Privacy

Welcome

Page

ID 12

Plea the o

Toda click

ID 191

12. Choose the probability of data being accessed that you would be comfortable opting-in for. *

1 (complete opt-in)
1/10
1/100
1/1000
1/1(
1/1(
1/1(
0 (c

ID 196

14. How much of an incentive in terms of dollars ($), in the range of 0 ($) to 1000 ($), that you would require for taking on a higher probability of 1/100,000? *

Incentive

Page d
You sel
service
a highe
underst
you wo
probabi

ID 197

15. How much of an incentive in terms of dollars ($), in the range of 0 ($) to 1000 ($), that you would require for taking on a higher probability of 1/10,000? *

ncentive in terms of dollars ($), in the range of 0 ($) to 1000 ($), that you
igher probability of 1/1000? *

Studies of Pref

Opportunity to a s about privacy ns.

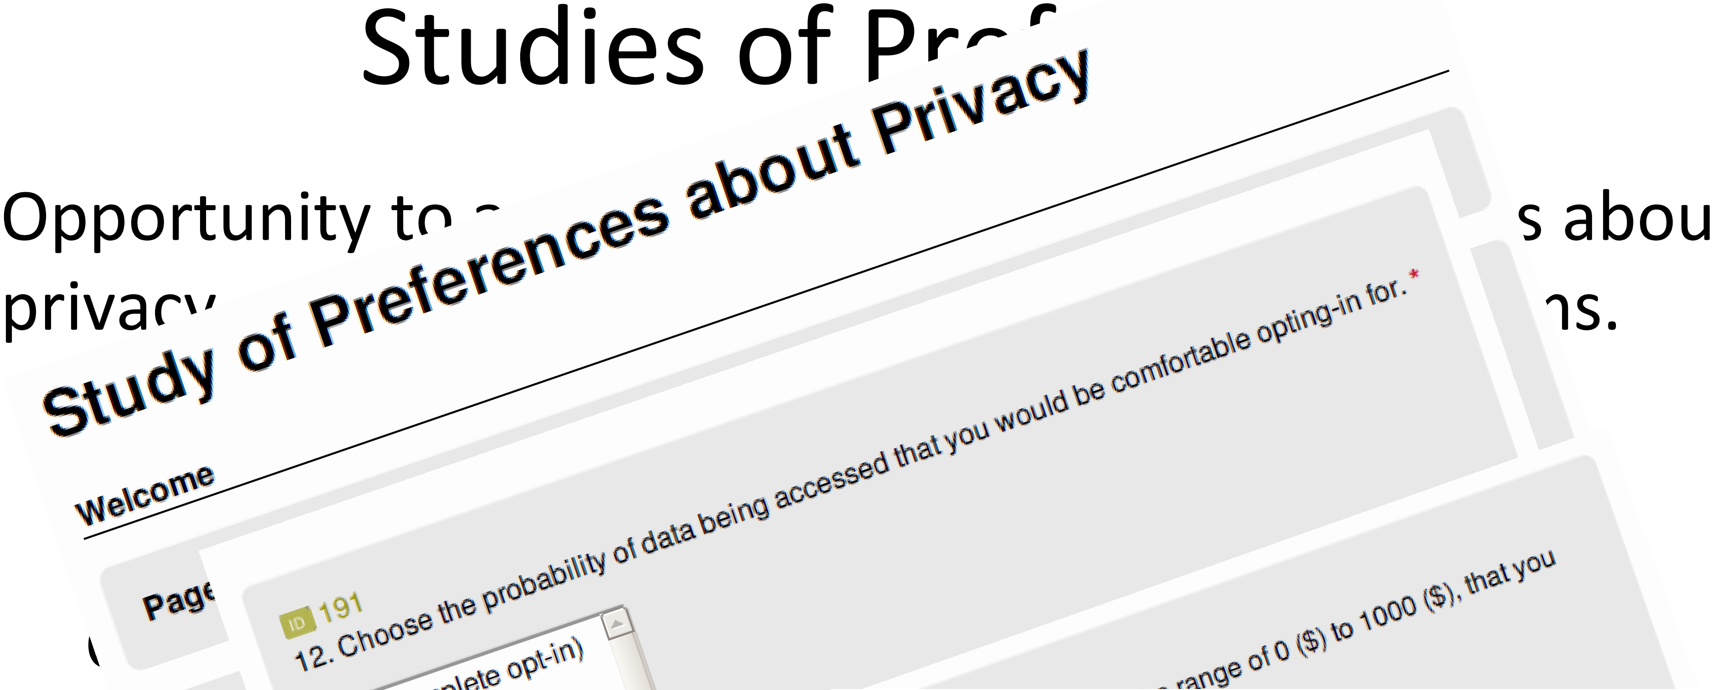

Study of Preferences about Privacy

Welcome

Page

ID 12

Plea the o

Toda click

ID 191
12. Choose the probability of data being accessed that you would be comfortable opting-in for. *

1 (complete opt-in)
1/10
1/100
1/1000
1/1(
1/1(
1/1(
0 (c

range of 0 ($) to 1000 ($), that you

ID 196
14. How much of an incenti would require for taking on

Incentive

Page d
You sel service a highe underst you wou probabi

ID 197
15. How much of a would require for

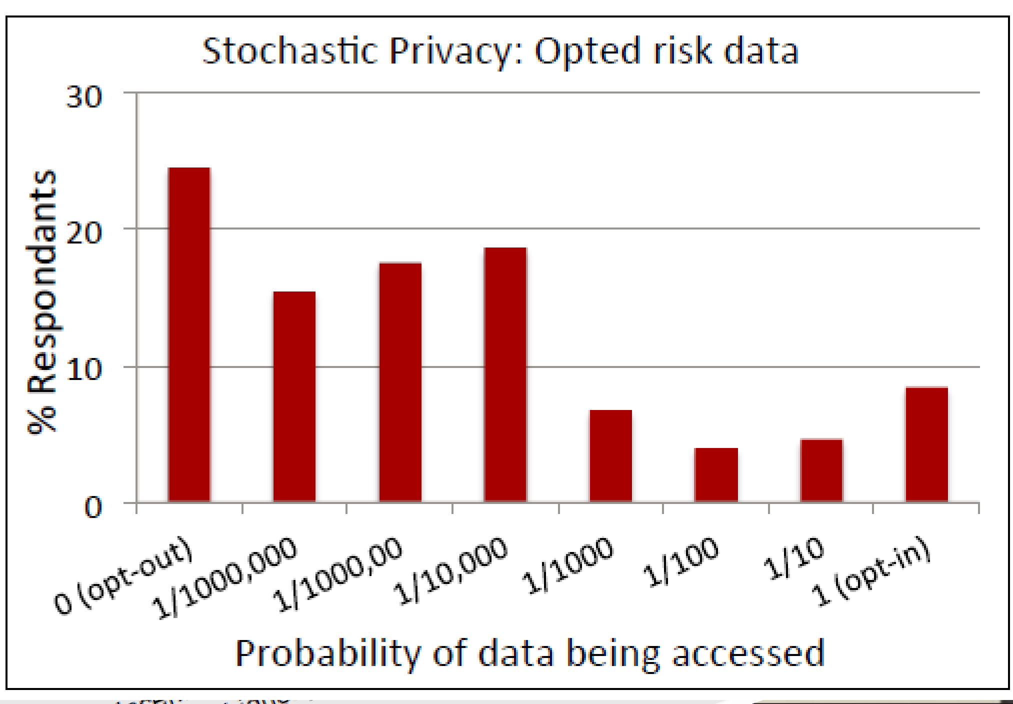

Stochastic Privacy: Opted risk data

Studies of Pre

Opportunity to a s about privacy ns.

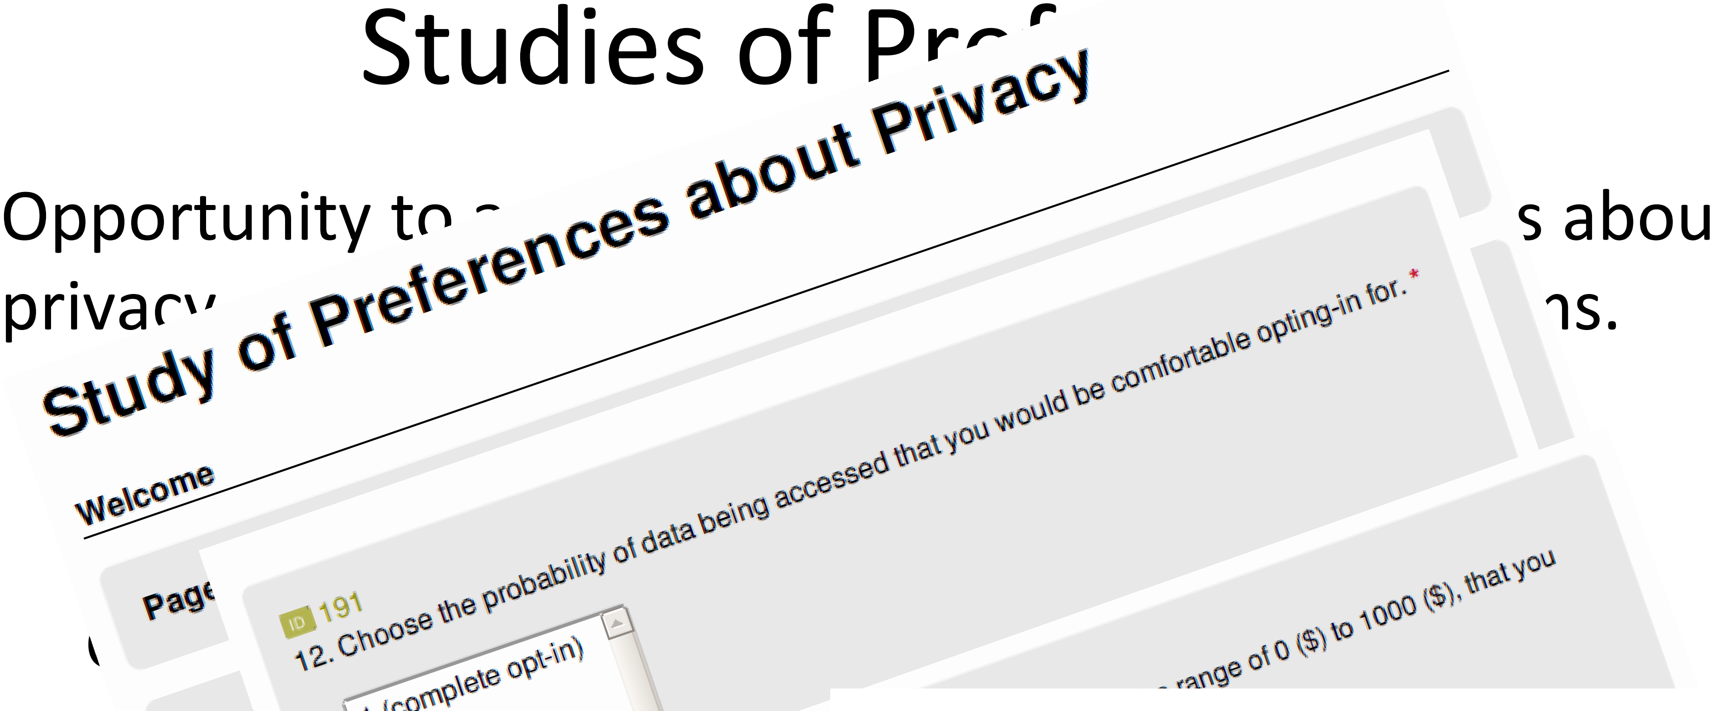

Study of Preferences about Privacy

Welcome

Page

ID 12

Plea the o

Toda click

ID 191
12. Choose the probability of data being accessed that you would be comfortable opting-in for. *

1 (complete opt-in)
1/10
1/100
1/1000
1/1(
1/1(
1/1(
0 (c

range of 0 ($) to 1000 ($), that you

ID 196
14. How much of an incenti would require for taking on

Incentive

Page d
You sel service a highe underst you wo probabi

ID 197
15. How much of a would require for

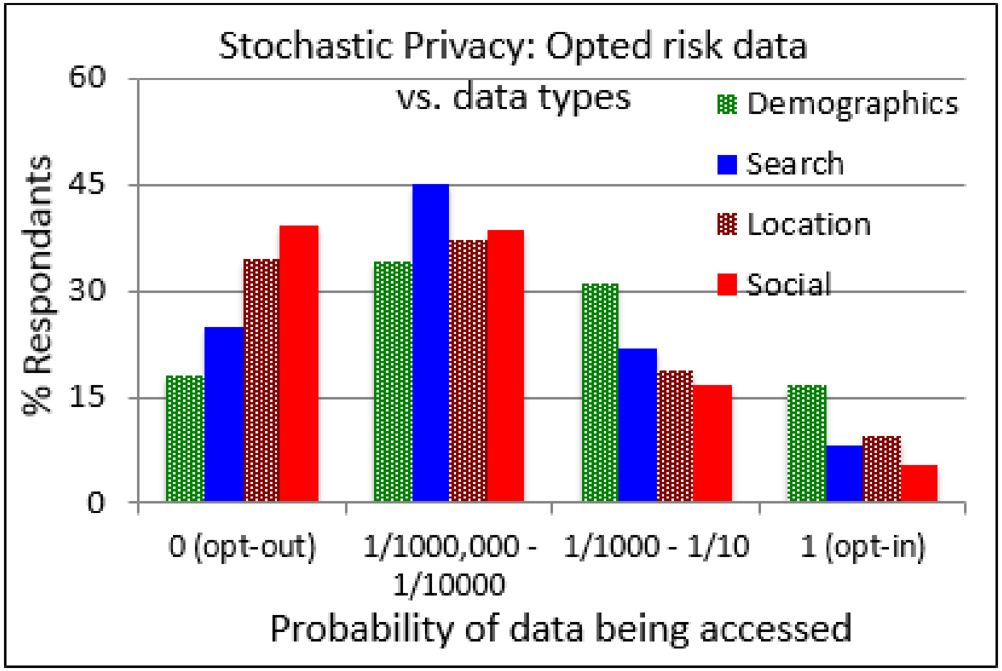

Harness AI for Privacy

Toward minimally-invasive sensing

AI methods for balancing sensitivity & value

Tradeoffs & optimization: QoS, revenues

Understand & assess user preferences