


Parallel Task Routing for Crowdsourcing

Jonathan Bragg
UW

Andrey Kolobov
MSR

Mausam
IIT Delhi

Daniel S. Weld
UW



UW
CrowdLab

HCOMP November 3, 2014

Microsoft
Research


iit

Task routing

Workers




Questions




Task routing

Workers



Questions



Prior work

- Offline
- Sequential routing
- Workers are willing to wait

[Donmez et al. 2009; Yan et al. 2011; Wauthier and Jordan 2011; Karger et al. 2013; Ho, Jabbari, Vaughan 2013; Chen et al. 2013]

Parallel routing

1. A pool of workers arrives
2. Assign a task to each worker
3. Workers take the same time to answer


Will relax this
assumption




Applications

- Volunteer (e.g., Zooniverse)
- Pay by task (e.g., Amazon Mechanical Turk)
- Pay by time (e.g., oDesk)


Problem space



Problem space

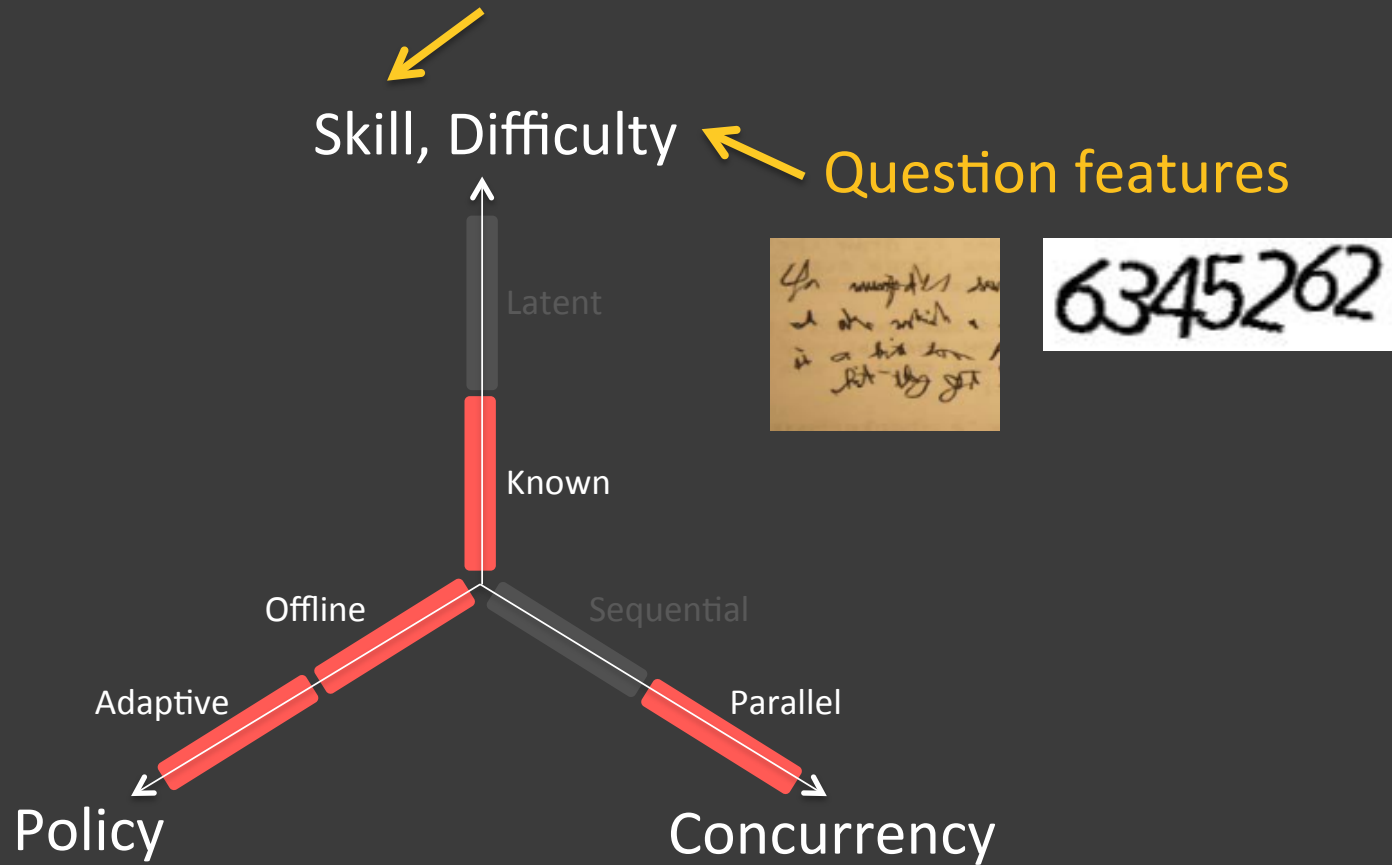


Problem space



Problem space


Community-based aggregation
[Venanzi et al. 2014]



Overview



- Workers not willing to wait
- Offline algorithm Juggler_{OFF}
- Adaptive algorithm Juggler_{AD}
 - Scalability
 - Variable response times

Problem definition



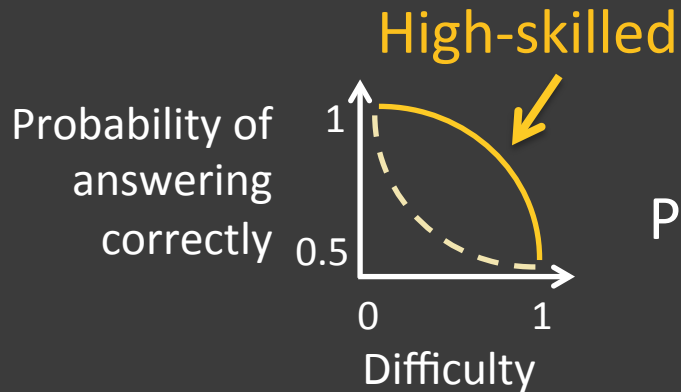
Problem definition

[Dai, Mausam, Weld 2010]

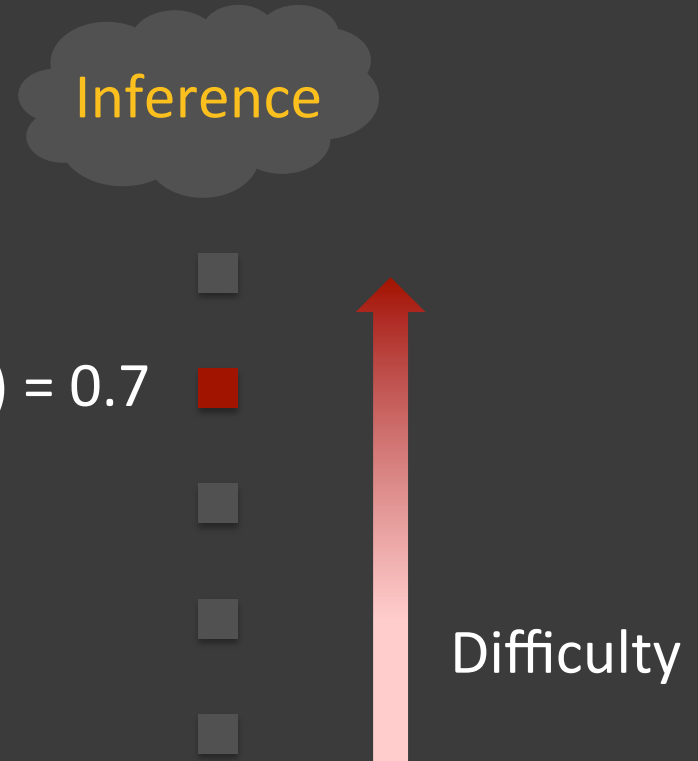


Problem definition


[Dai, Mausam, Weld 2010]



$$P(\text{True} \mid \text{votes}) = 0.7$$



Problem definition



Problem definition

Answers Assigned questions

Goal: Minimize entropy $H(\mathcal{A} \mid \mathcal{X}_S)$

$P(\text{True} \mid \text{votes}) = 0.7$ ■

$= 0.7$ ■

$= 0.8$ ■


$= 0.9$ ■

$= 0.9$ ■





Difficulty



Offline routing



Hardness



Hardness



Theorem

NP-hard, by reduction from the PartitionProblem

Juggler_{OFF}

Submodular for
known skill and difficulty

Goal: Minimize entropy $H(\mathcal{A} \mid \mathcal{X}_S)$

The Juggler_{OFF} algorithm

$\mathcal{X}_S \leftarrow \emptyset$


while not_finished:

$\mathcal{X}_S \leftarrow \mathcal{X}_S \cup \left\{ \arg \max_{X \in \mathcal{X} \setminus \mathcal{X}_S} H(\mathcal{A} \mid \mathbf{x}) - H(\mathcal{A} \mid X, \mathcal{X}_S, \mathbf{x}) \right\}$

Remaining
assignments


Observed votes

Juggler_{OFF}



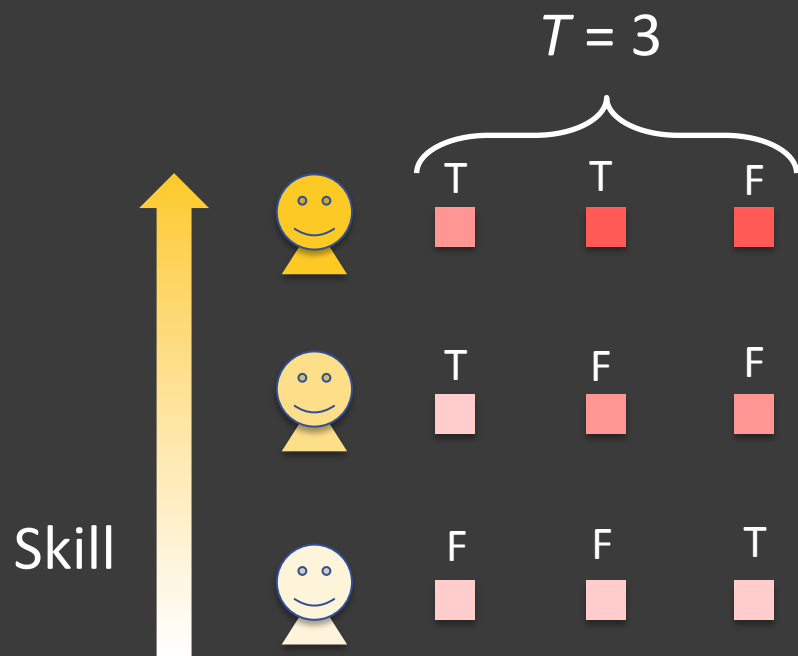
Juggler_{OFF}

Reverse skill!

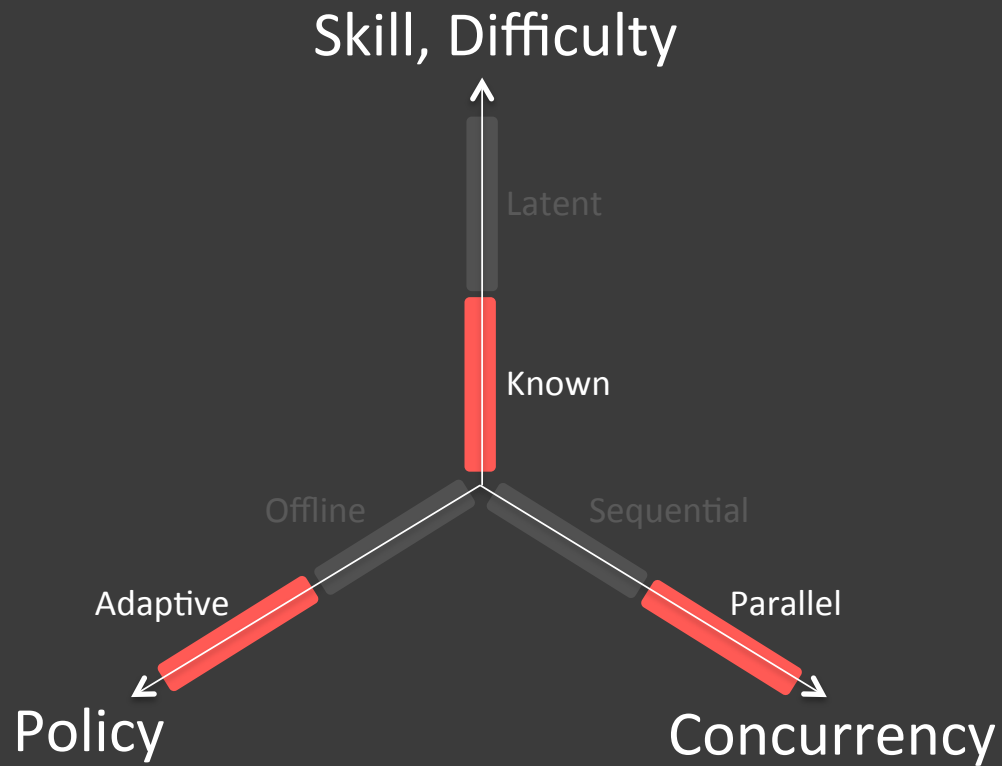


Juggler_{OFF}

Reverse skill!



Adaptive routing



Juggler_{AD}


The Juggler_{AD} algorithm

for $t = 1$ to T :

$\mathcal{X}_{S_t} \leftarrow$ call Juggler_{OFF} with horizon 1

Observe votes and update prior

Juggler_{AD}



Live experiment

12 workers

oDesk

198 questions

named-entity
linking

Tollner's dismissal was announced after his team lost to **Notre Dame** and U.C.L.A., but he was allowed to coach through the bowl game.

Which of the following Wikipedia articles defines "**Notre Dame**" in exactly the way it is used in the above sentence? [Tell me more](#)

● [Notre Dame Fighting Irish football](#)

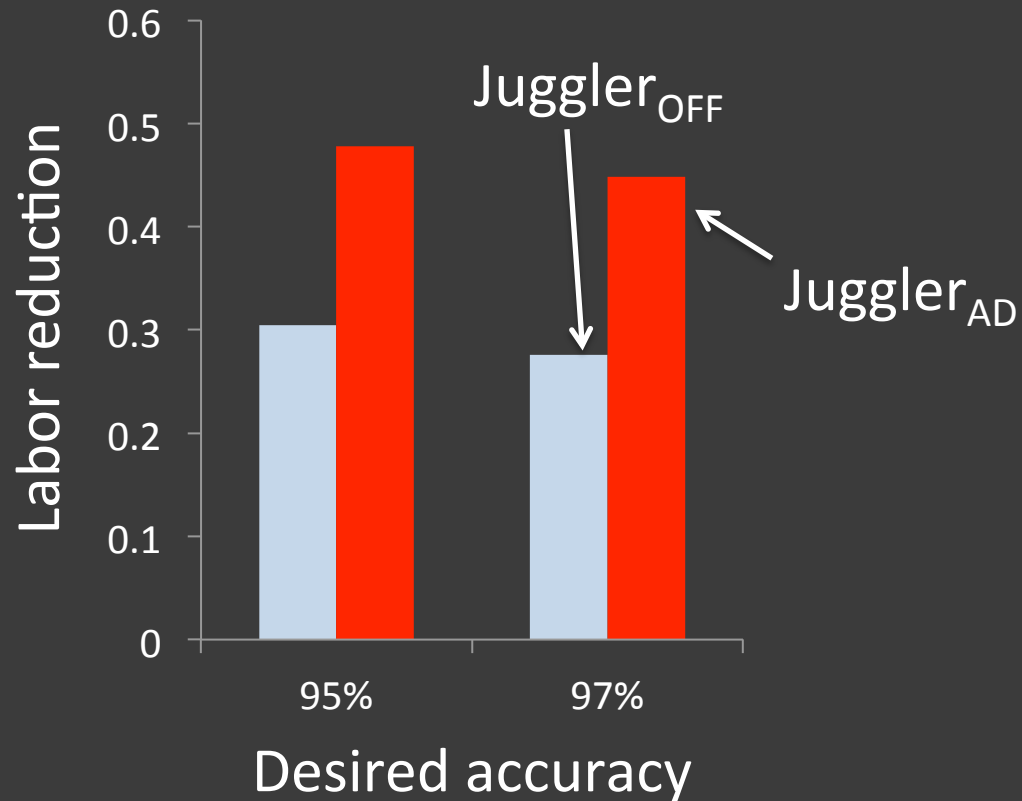
The Notre Dame Fighting Irish football team is the intercollegiate football team of the University of Notre Dame. The team is currently coached by Brian Kelly and play home games at the campus' Notre Dame Stadium, with a capacity of 80,795. Notre Dame competes as an Independent at the NCAA Football Bowl Subdivision level and is a founding member of the Bowl Championship Series coalition (BCS). The Fighting Irish hold the highest winning percentage among college football programs and have 13 national championships recognized by the NCAA, tied for first out of all FBS schools in the post-1900 er...

● [University of Notre Dame](#)

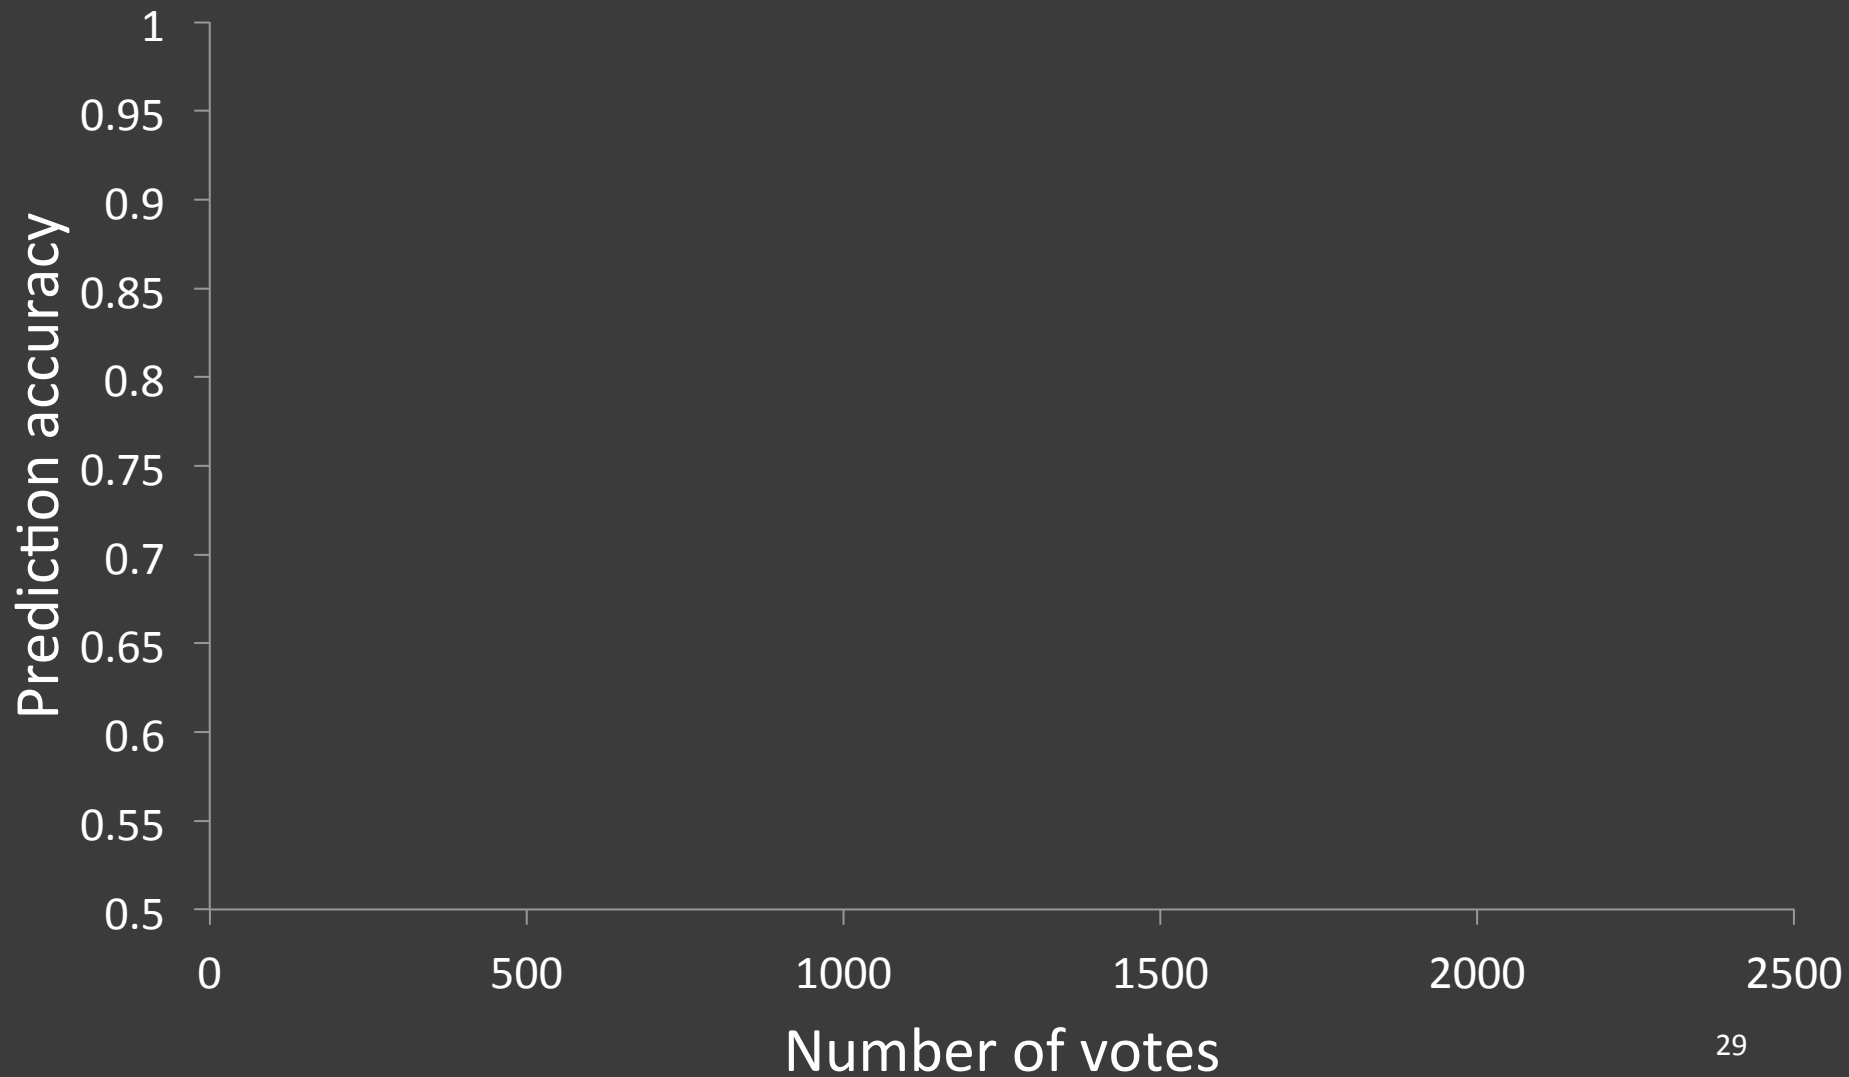
The University of Notre Dame du Lac (or simply Notre Dame NOH-tr-DAYM) is a Catholic research university located near South Bend, Indiana, in the United States. "Notre Dame" means "Our Lady" in French and refers to the university's patron saint, the Virgin Mary....

Submit

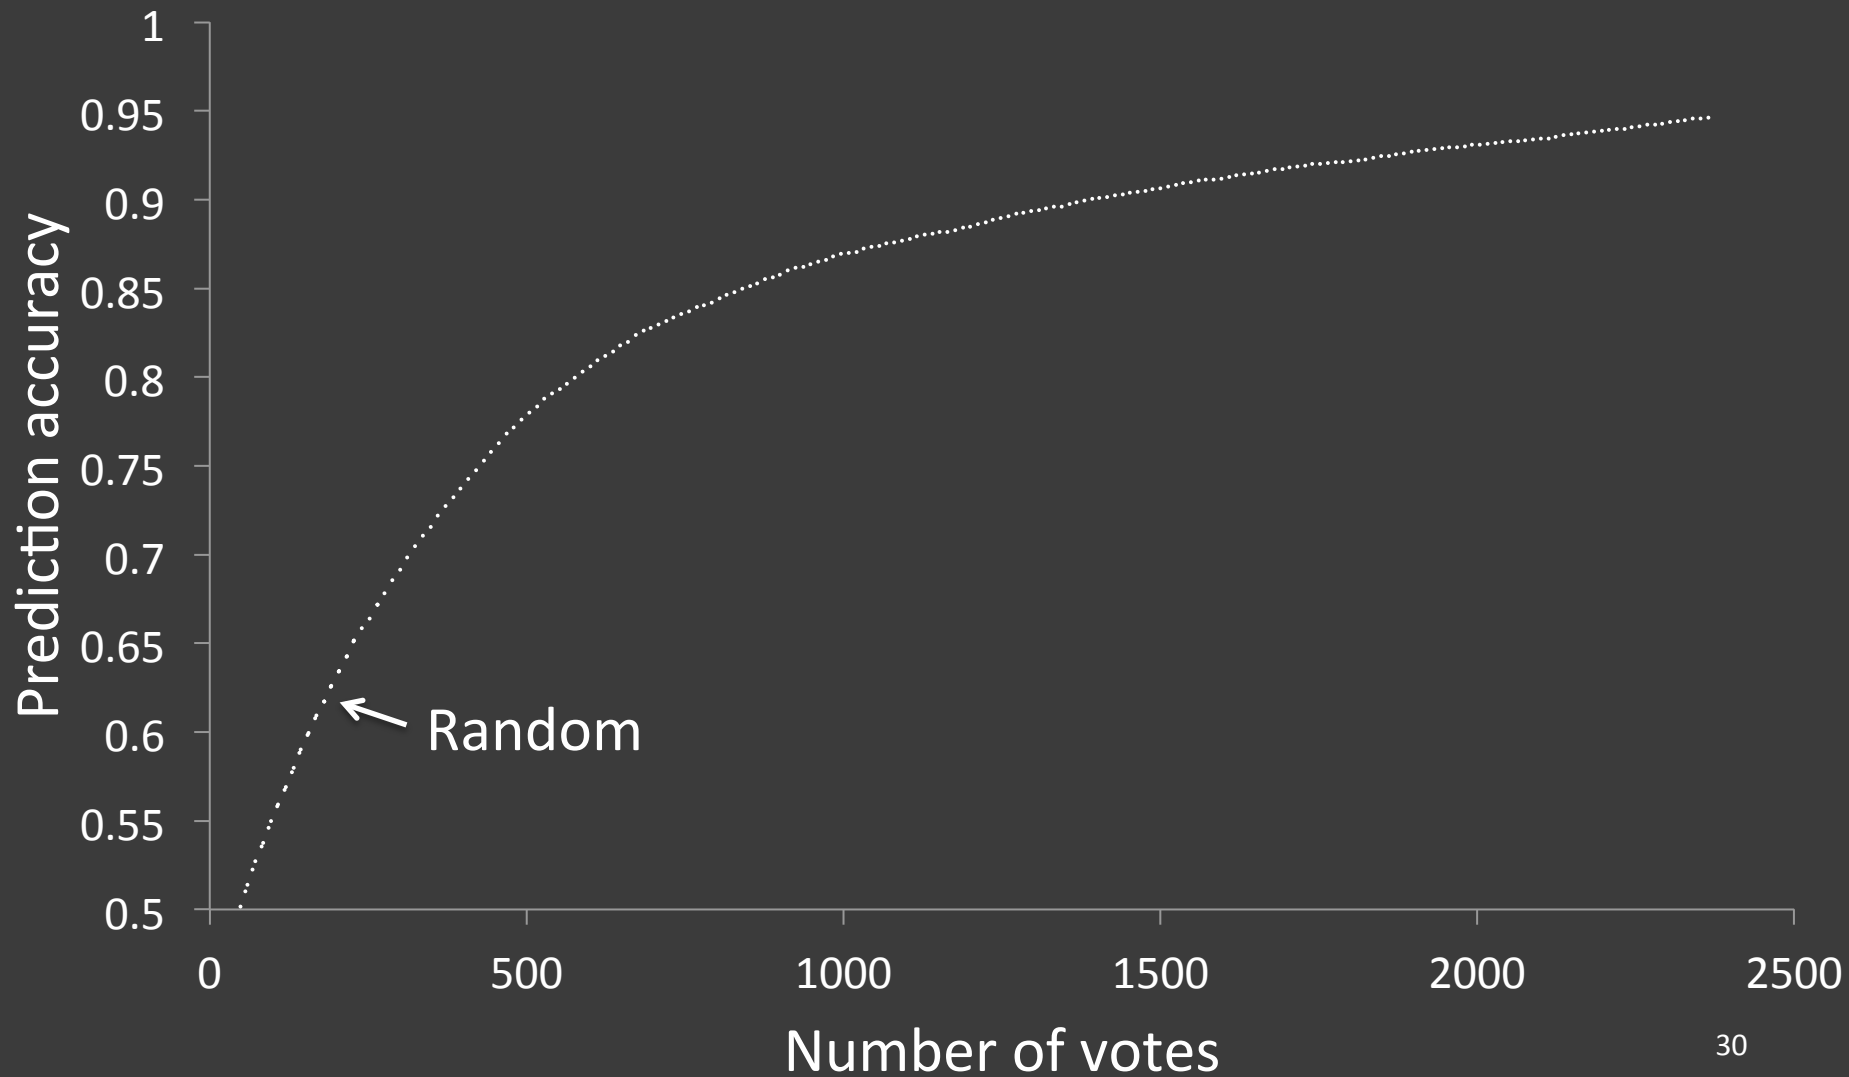
Adaptivity benefit



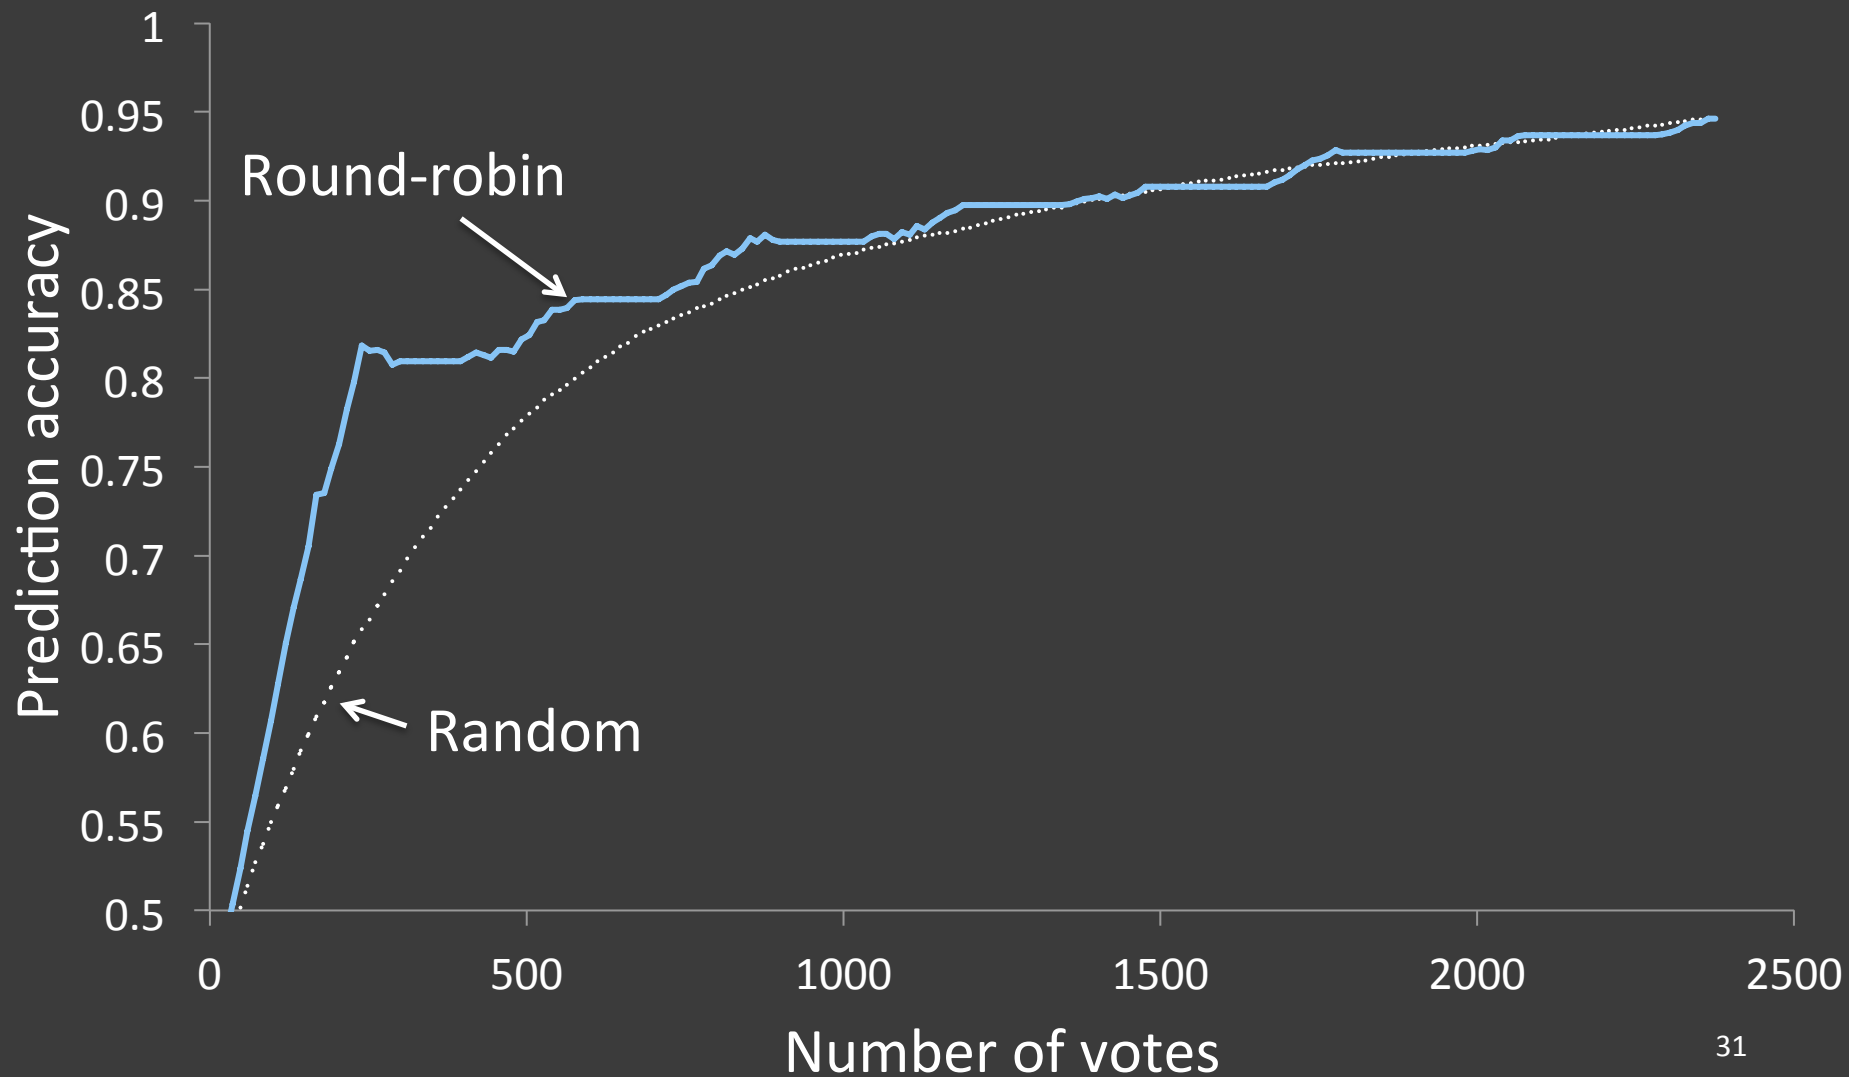
Live experiment



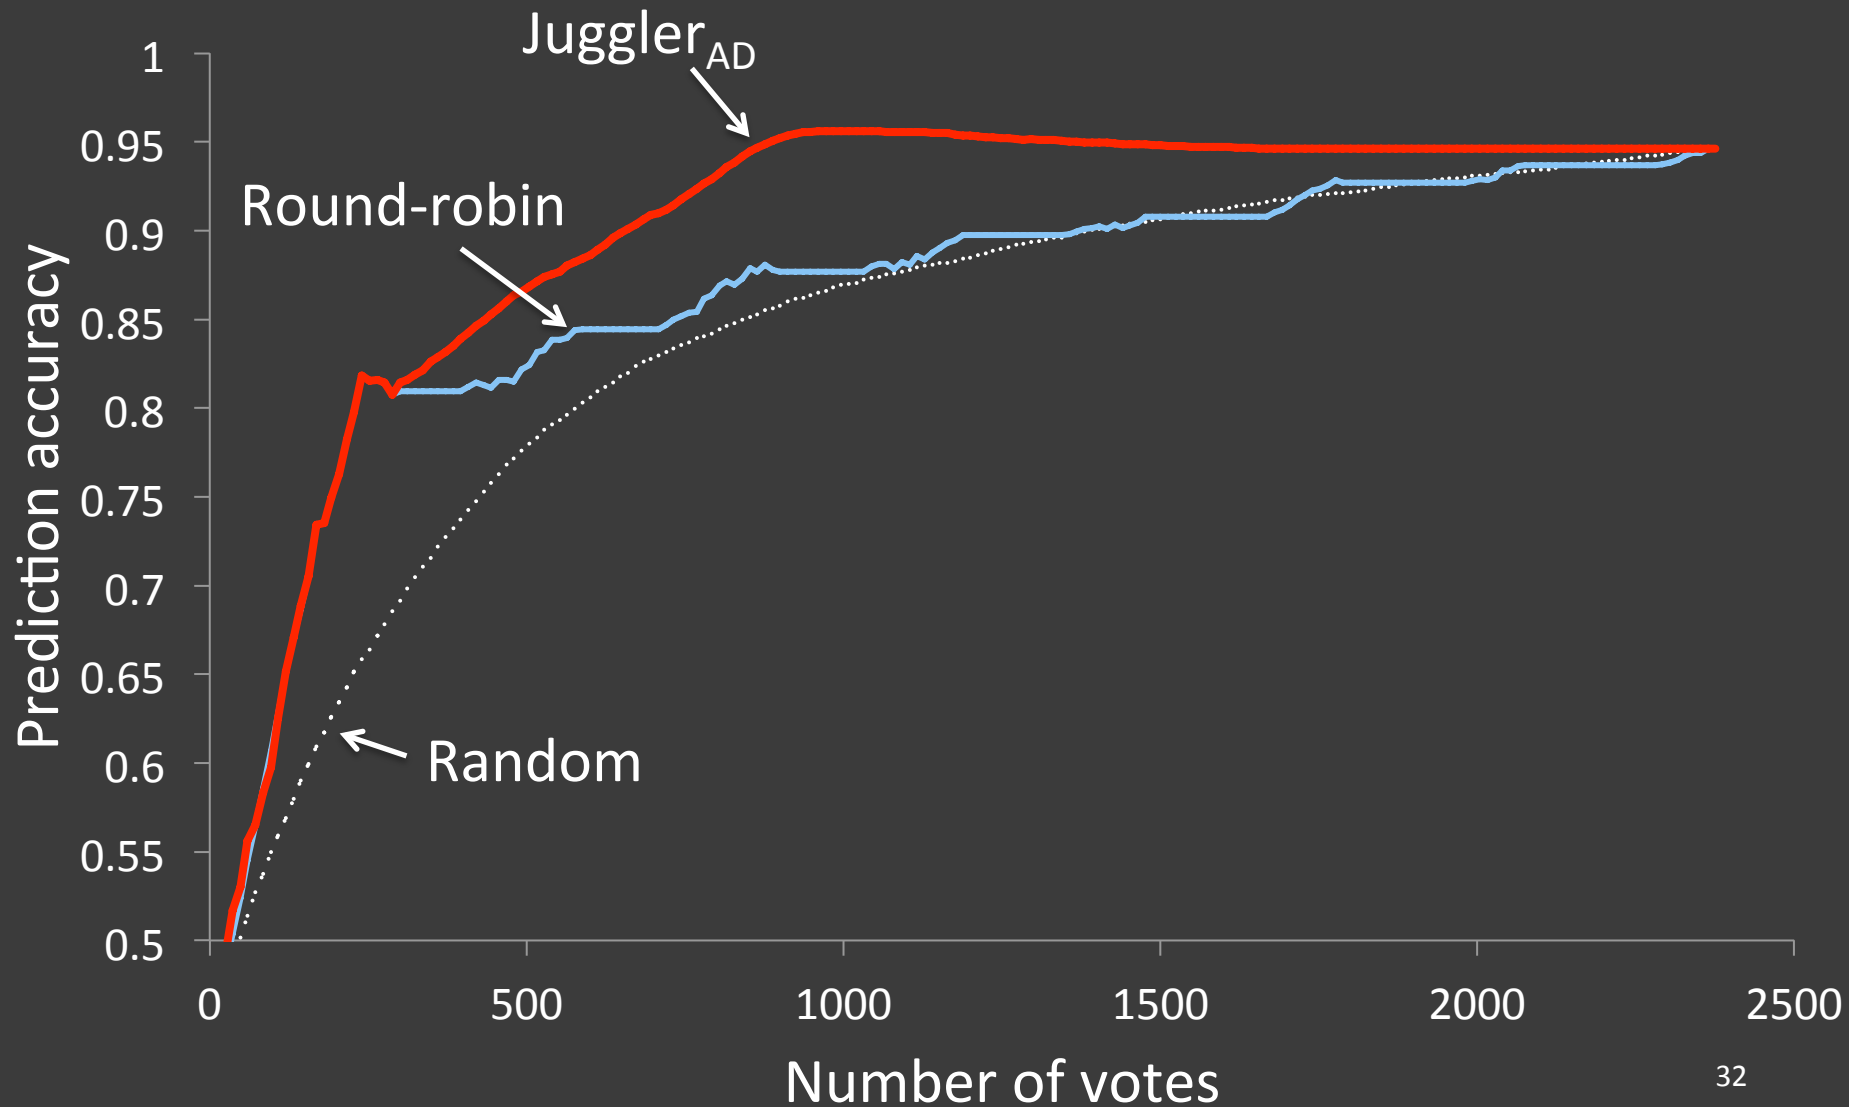
Live experiment



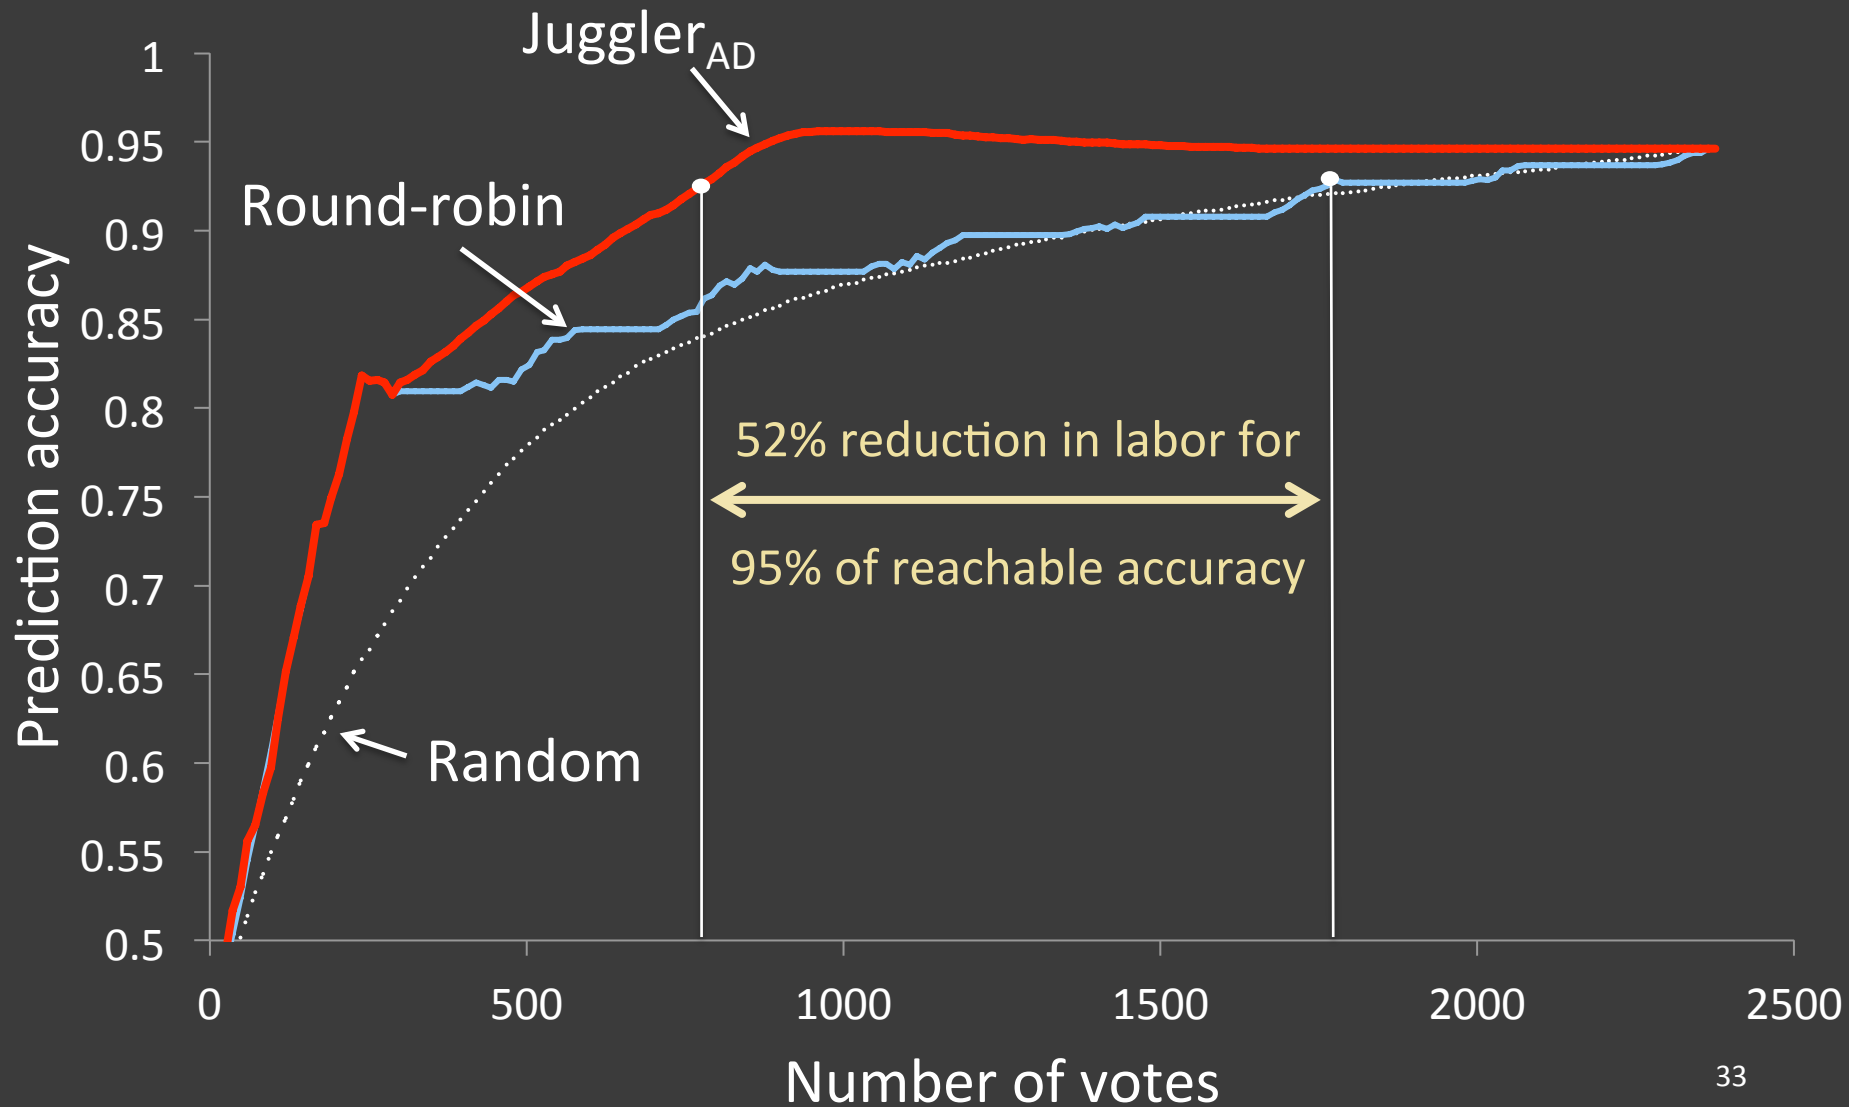
Live experiment




Live experiment



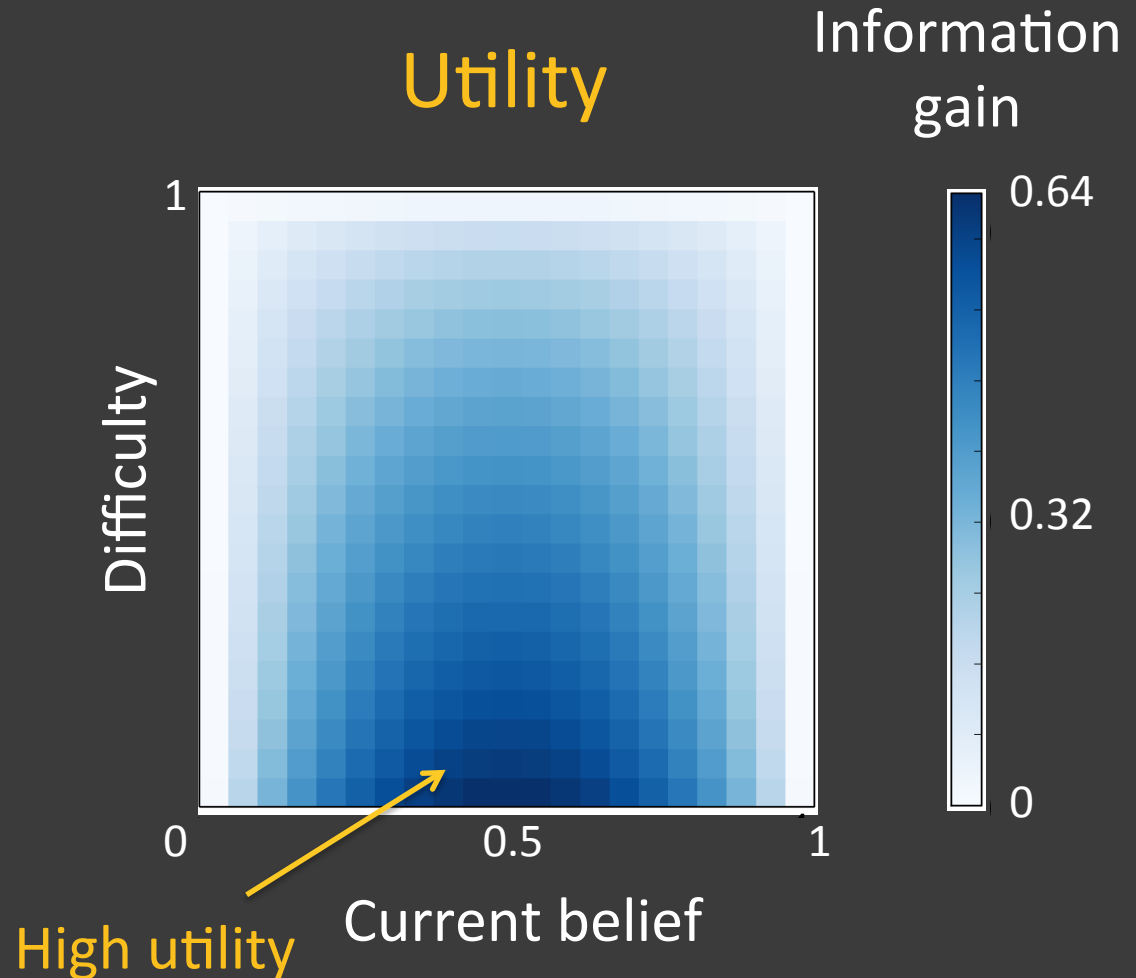
Live experiment




Scalability



Binning questions




Binning questions




Binning questions


Lower skill



Higher skill




Information
gain




Larger-scale experiment

100 workers
2k questions




Larger-scale experiment

100 workers
2k questions




Larger-scale experiment

100 workers
2k questions




Larger-scale experiment

100 workers
2k questions




Larger-scale experiment


100 workers
2k questions




Variable response times



Variable response times




Variable response times




Summary

- Workers not willing to wait
- Offline algorithm Juggler_{OFF}
- Adaptive algorithm Juggler_{AD} Reduce labor by 52%
 - Scalability
 - Variable response times

Future directions



Thanks!



UW
CrowdLab

Microsoft
Research

iit

