Optimal Testing for Crowd Workers

Jonathan Bragg
University of Washington

Mausam
IIT Delhi

Daniel S. Weld
University of Washington

Crowdsourcing

- Fun

The ESP Game

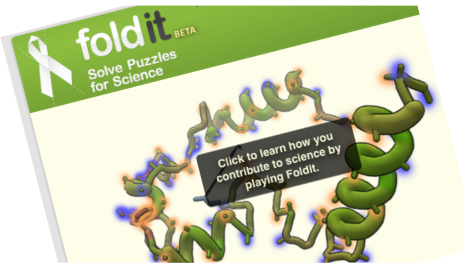

- Altruism

ZOONIVERSE

GALAXY ZOO

- Money

amazon mechanical turk beta
Artificial Artificial Intelligence

upwork™
formerly oDesk

Supervised learning

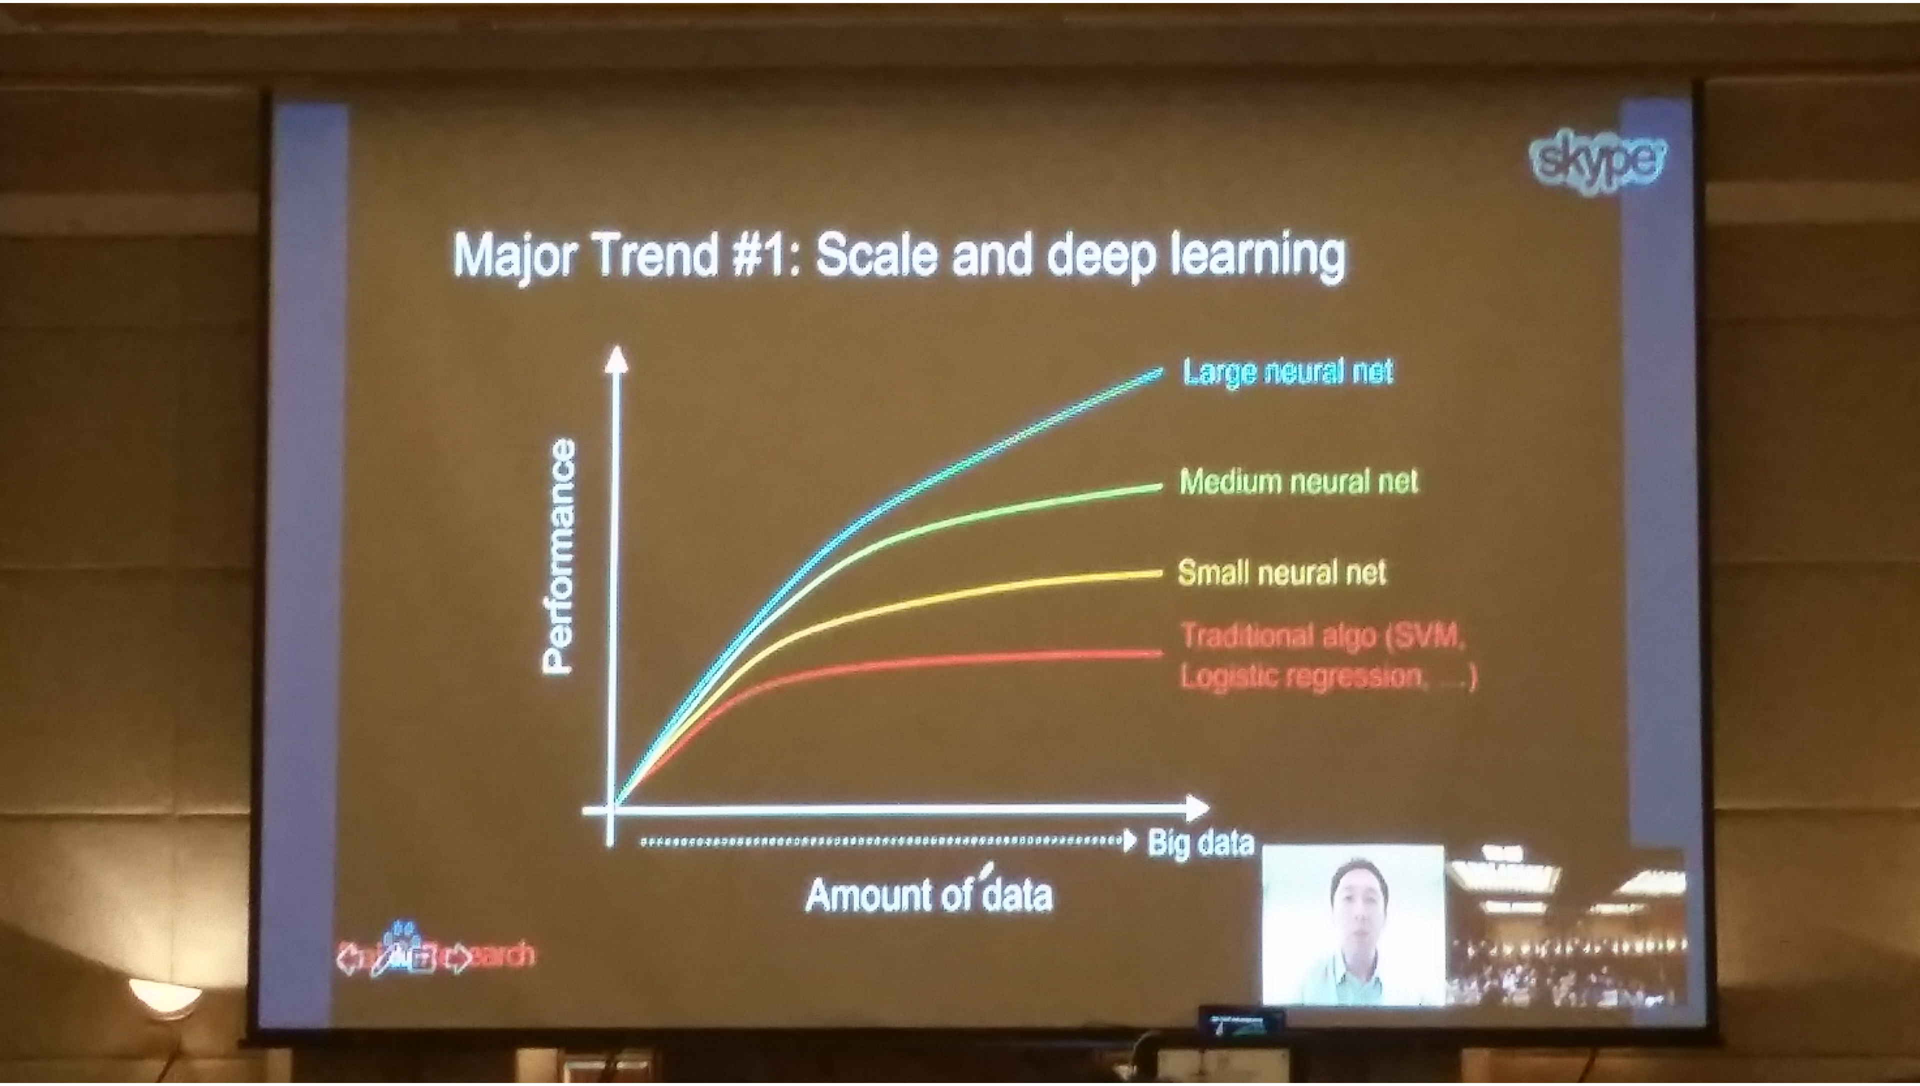

Quality control

High quality

Low quality

Quality control

High quality

?

Low quality

Quality control in practice

- Gold question insertion
- Worker filtering

Quality control in practice

- Gold question insertion — What fraction?
- Worker filtering — What accuracy?

Quality control parameters

- Researchers (NLP relation extraction task)
  - 13% gold, 67% accuracy threshold [1]
  - 30% gold, 85% accuracy threshold [2]
  - 20% gold, 80% accuracy threshold [3]
- Industry
  - CrowdFlower: 20% gold, 80% accuracy threshold

[1] Angeli et al. 2014. In EMNLP.
[2] Gormley et al. 2010. In NAACL.
[3] Zhang et al. 2012. In ACL.

Problems

- How to set parameters?
- Static policy is sub-optimal
  - Heterogeneous workers
  - Non-stationary worker population

Goal

Automatically insert gold questions to maximize **quality** & **quantity** of individual worker answers

Agent architecture

Give me answers
above accuracy **a***

Gold questions
Base policy

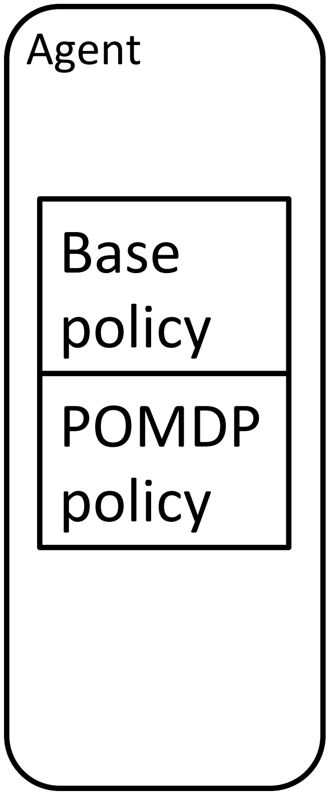

Agent architecture

Give me answers above accuracy **a***

Gold questions
Base policy

Agent

Base policy

POMDP policy

Actions:
Ask **test** question
Ask **work** question
Replace worker

Environment

amazon mechanical turk beta

Gold (test) questions

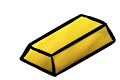

Agent architecture

Give me answers
above accuracy **a***

Gold questions
Base policy

Agent

Base
policy

POMDP
policy

Actions:
Ask **test** question
Ask **work** question
Replace worker

Observations:
Correct test answer
Incorrect test answer
Worker **leaves**

Environment

amazon mechanical turk beta

Gold (test)
questions

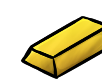

Agent architecture

Unobservable!

Reward: $\begin{cases} 1 & \text{if correct } \textbf{work} \text{ answer} \\ \text{-PENALTY} & \text{if incorrect } \textbf{work} \text{ answer} \end{cases}$

Give me answers above accuracy **a***

Gold questions
Base policy

Agent

Base policy

POMDP policy

Actions:
Ask **test** question
Ask **work** question
Replace worker

Observations:
Correct test answer
Incorrect test answer
Worker **leaves**

Environment

amazon mechanical turk beta

Gold (test) questions

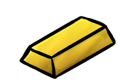

Reinforcement learning

- Challenges
  - Unobservable rewards
  - Exploration should not upset workers
- Our approach
  - Explore with base policy (e.g., “test 20%”)
  - Exploit with learned POMDP policy

POMDP formulation

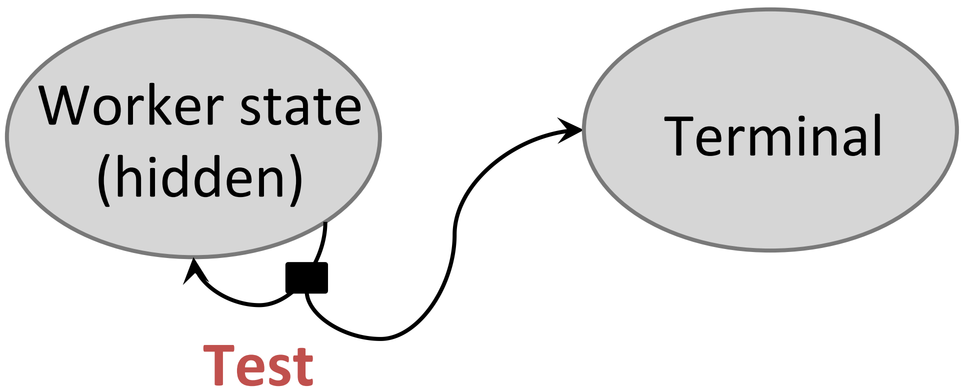

Reward: 0

Observation: Correct / incorrect

POMDP formulation

Set PENALTY s.t. Reward > 0 iff worker accuracy > **a***

Reward: P(correct) * 1 + (1 – P(correct)) * (-PENALTY)
Observation: NULL

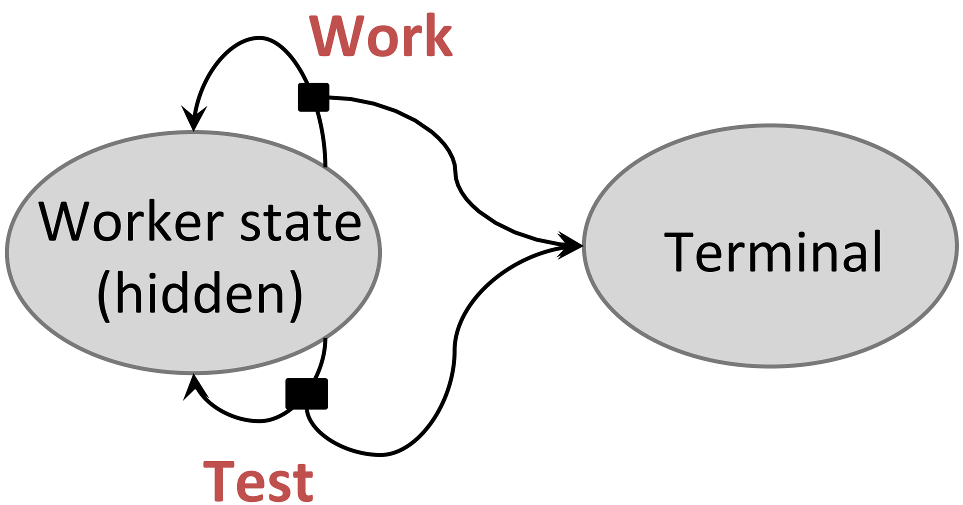

Reward: 0
Observation: Correct / incorrect

POMDP formulation

Set PENALTY s.t. Reward > 0 iff worker accuracy > **a***

Reward: P(correct) * 1 + (1 – P(correct)) * (-PENALTY)
Observation: NULL

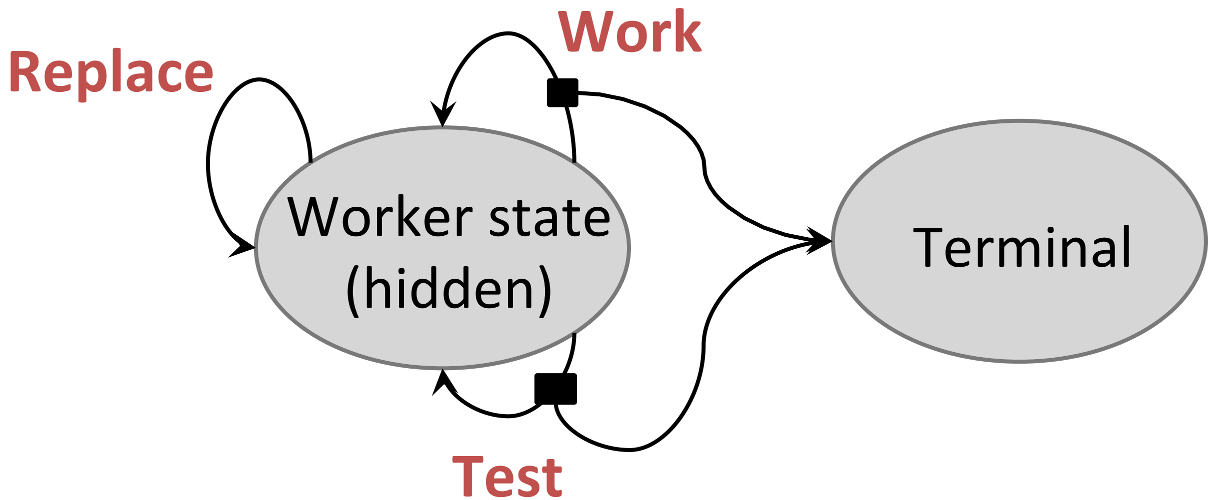

Reward: 0
Observation: Correct / incorrect

Worker state transitions

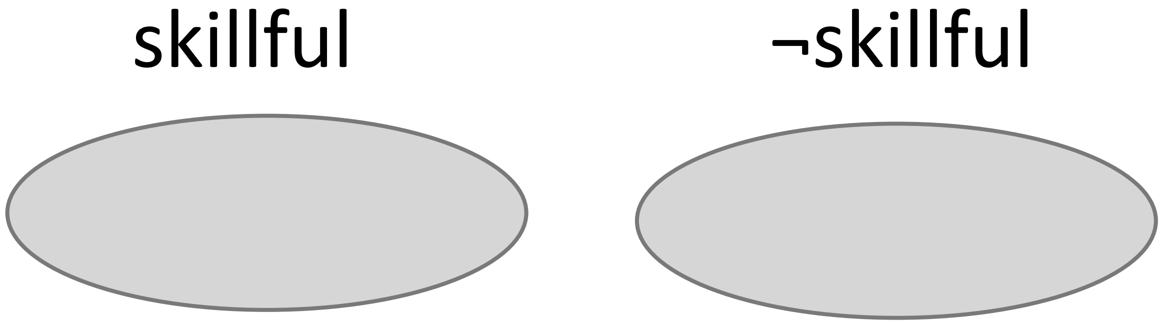

Worker state transitions

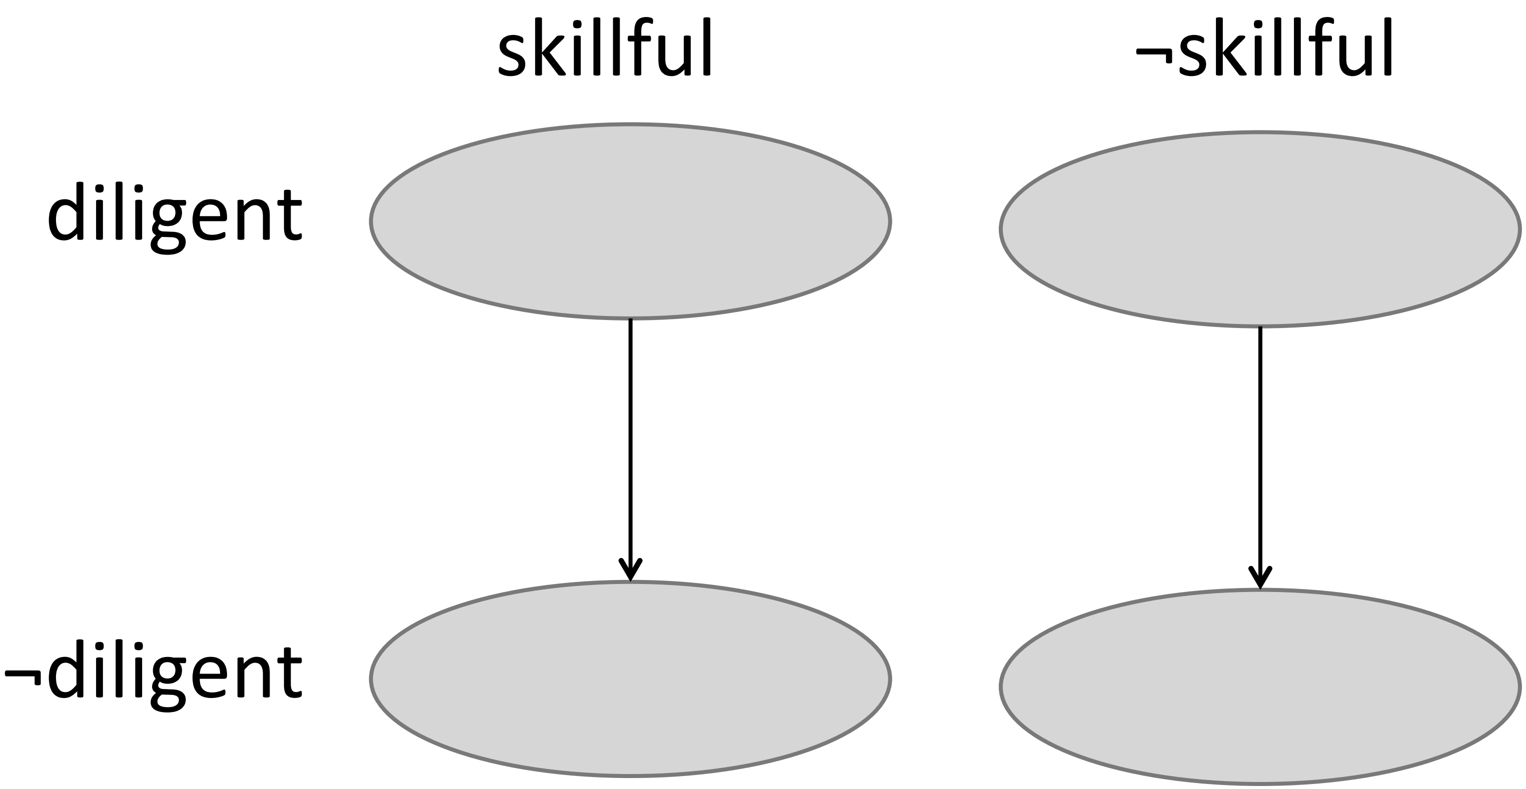

Worker state transitions

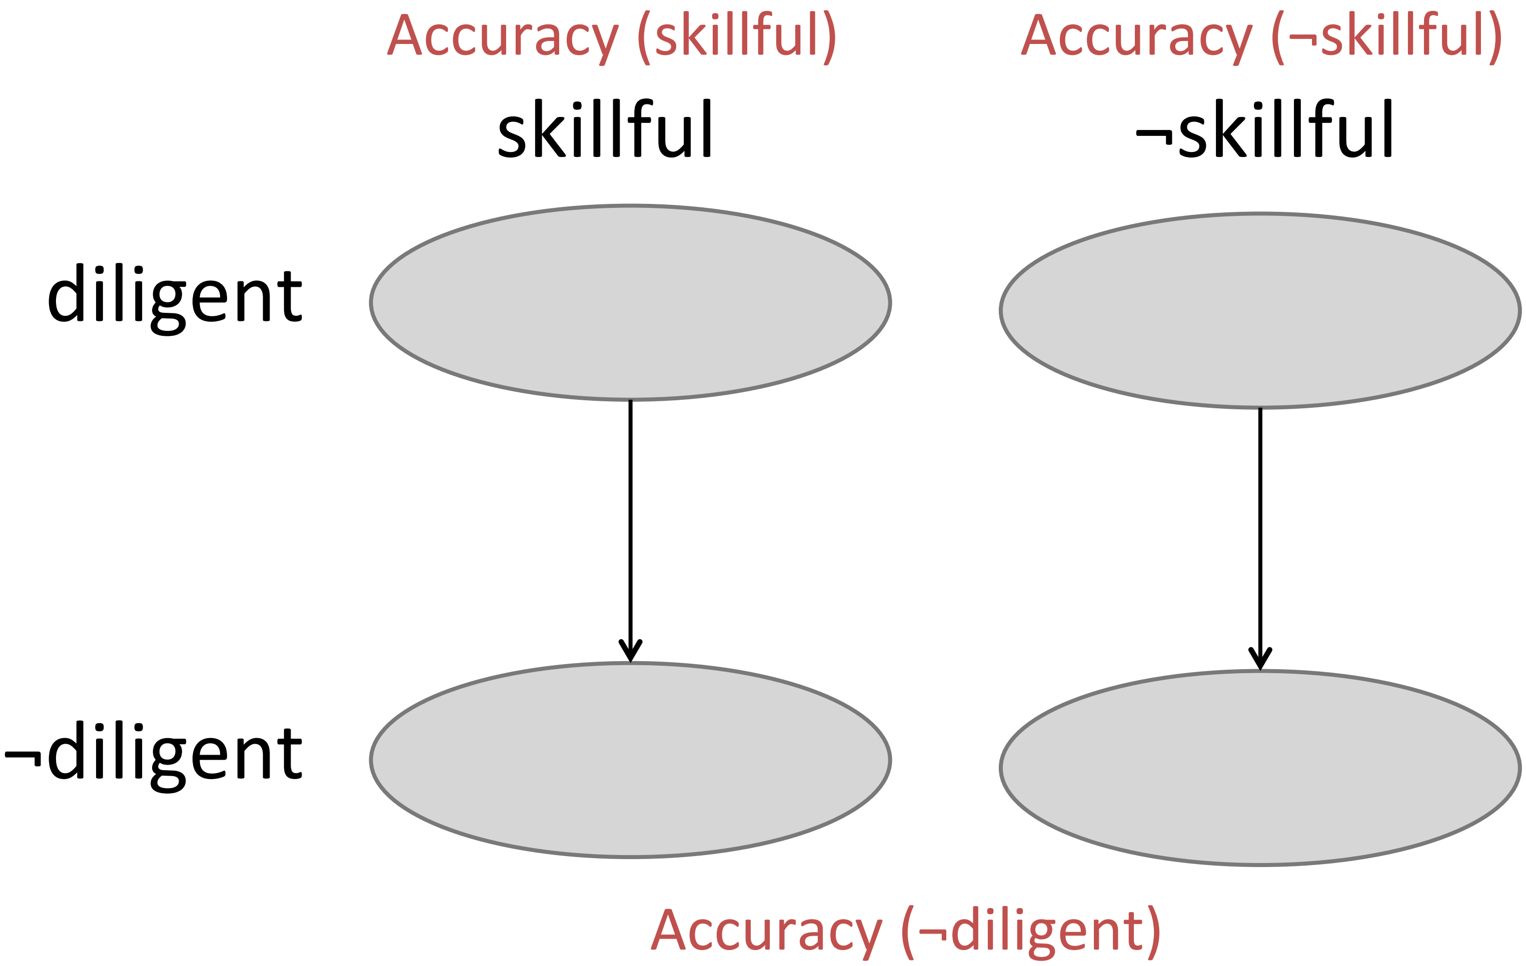

POMDP parameters to estimate

- Probability worker leaves
- Probability worker becomes ¬diligent
- Accuracy (skillful)
- Accuracy (¬skillful)
- Accuracy (¬diligent)
- Ratio of skillful to ¬skillful workers

Experiments

- Existing datasets from Mechanical Turk: LinWiki, LinTag, [1] and Rajpal [2]
- Desired accuracy a* = 0.85

[1] Lin et al. 2012. In AAAI.
[2] Rajpal et al. 2015. In ICML workshop.

Task: Named Entity Linking

"Only two states -- Vermont and **Washington** -- this year joined five others requiring private employers to grant leaves of absence to employees with newborn or adopted infants."

- ○ Washington, D.C.
- ● Washington (state)

Submit

Policies

- Work-only
- Test-and-replace
  - Test 20%
  - Replace if accuracy < 0.85
- POMDP-RL
  - Base policy: Test-and-replace (above)
  - Exploration budget: 20 workers

LinWiki dataset

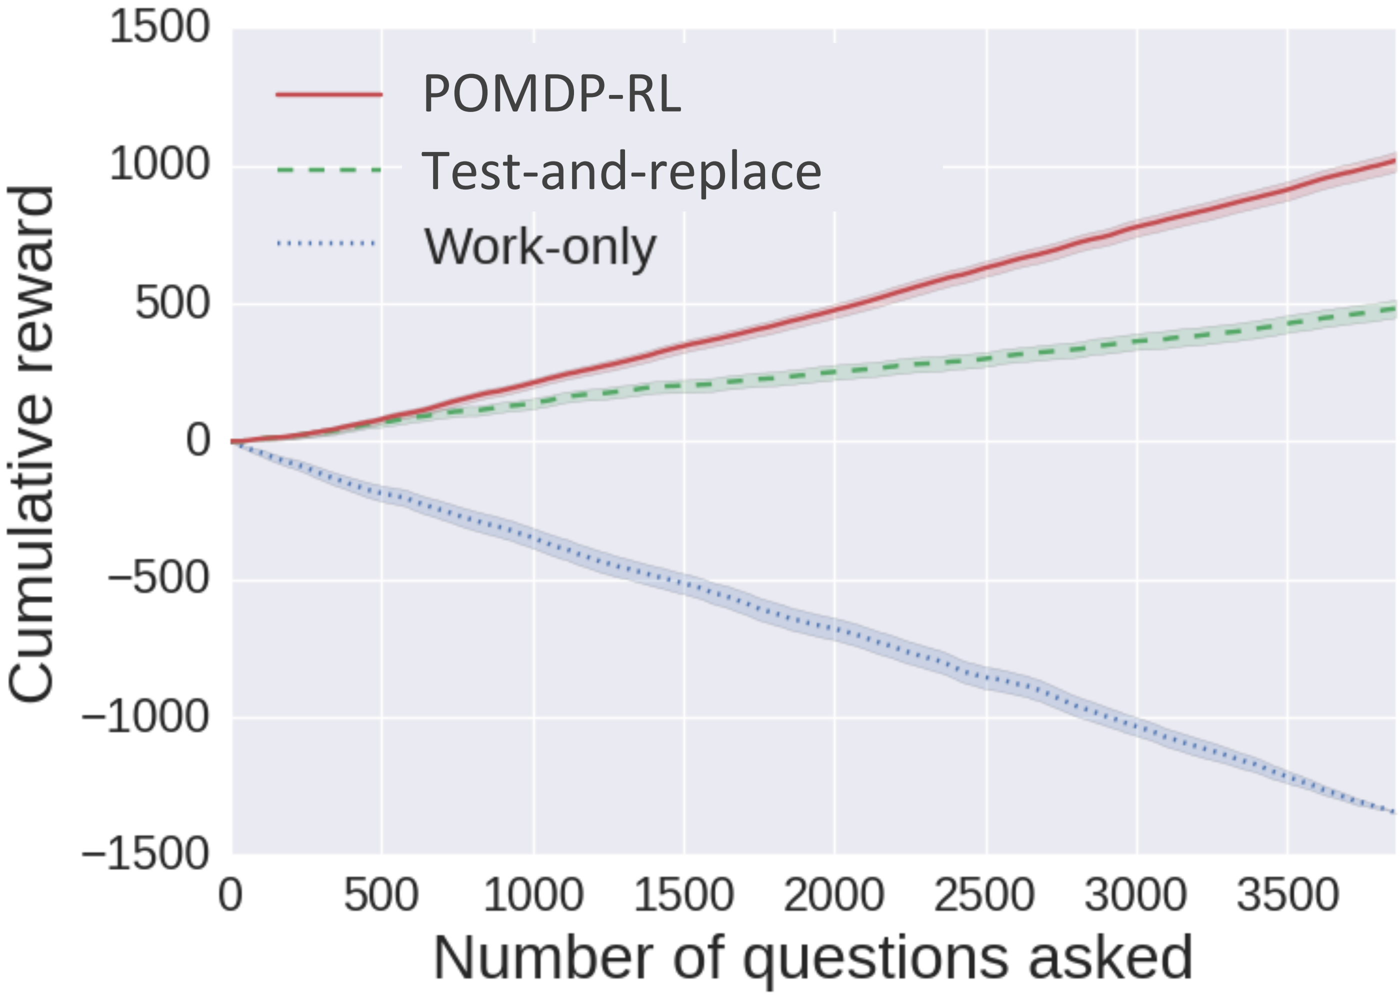

Rajpal dataset

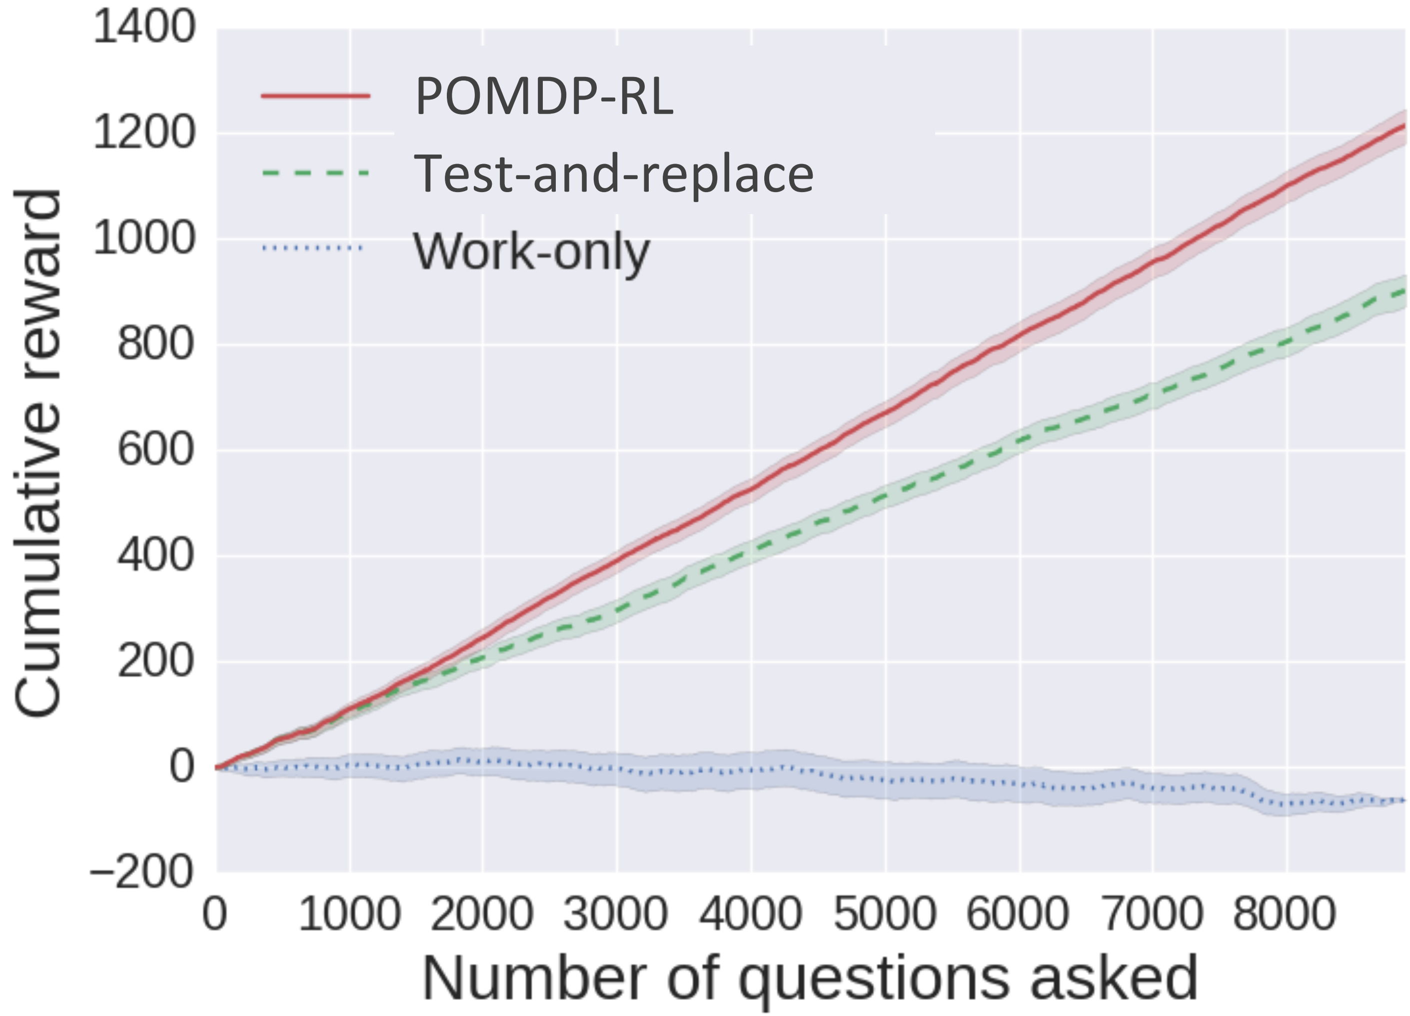

LinTag dataset

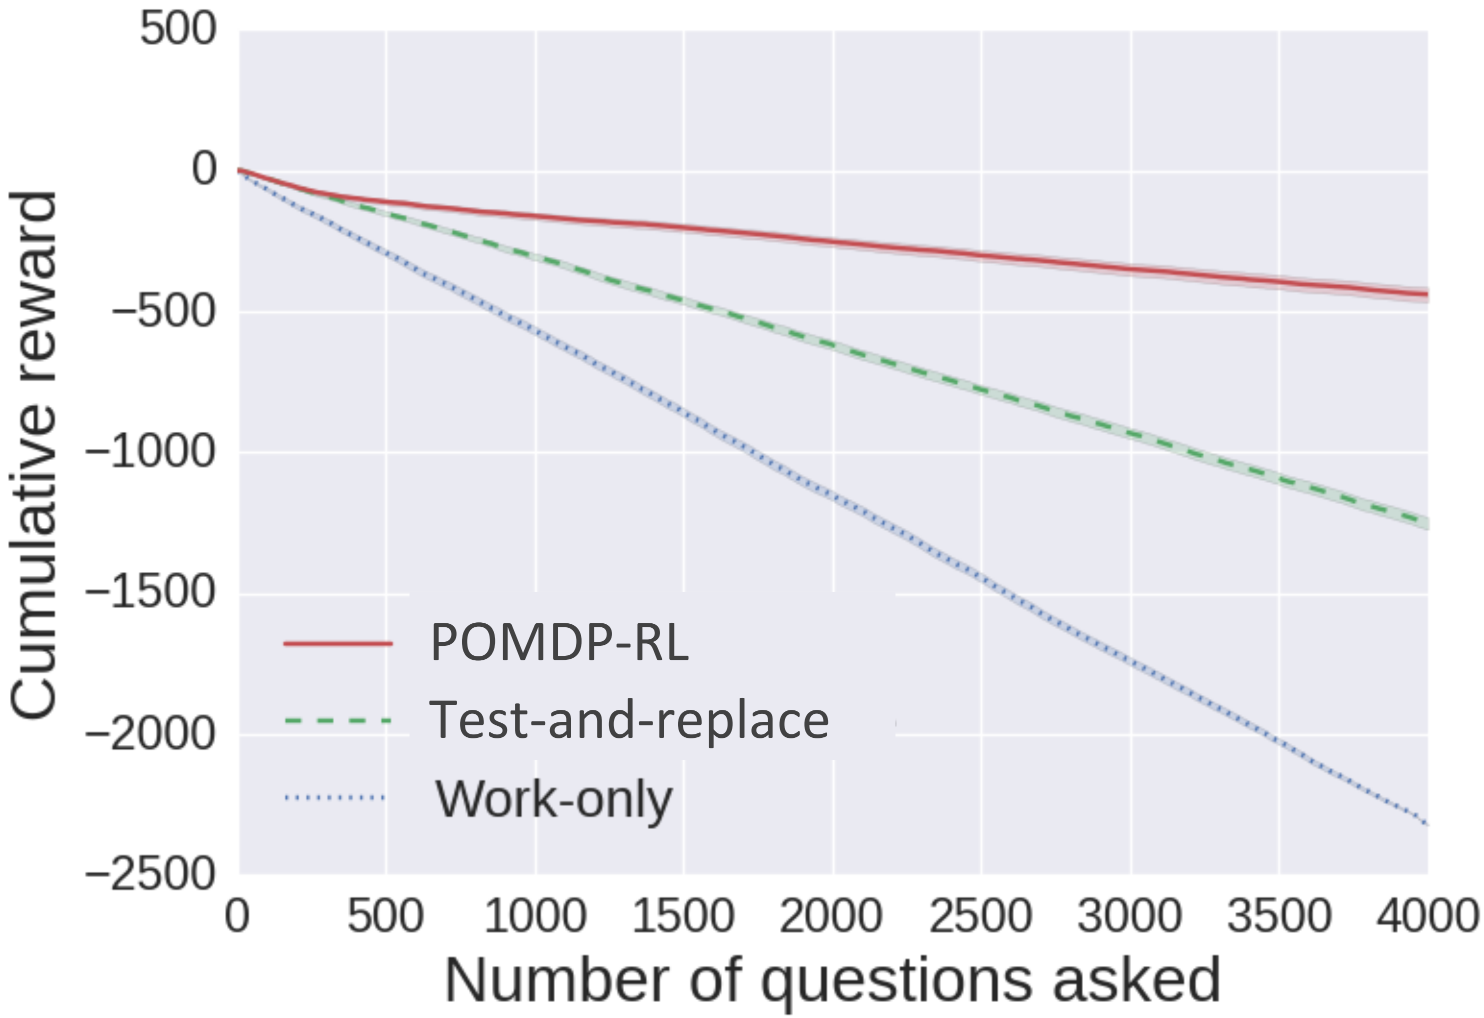

LinTag worker qualities

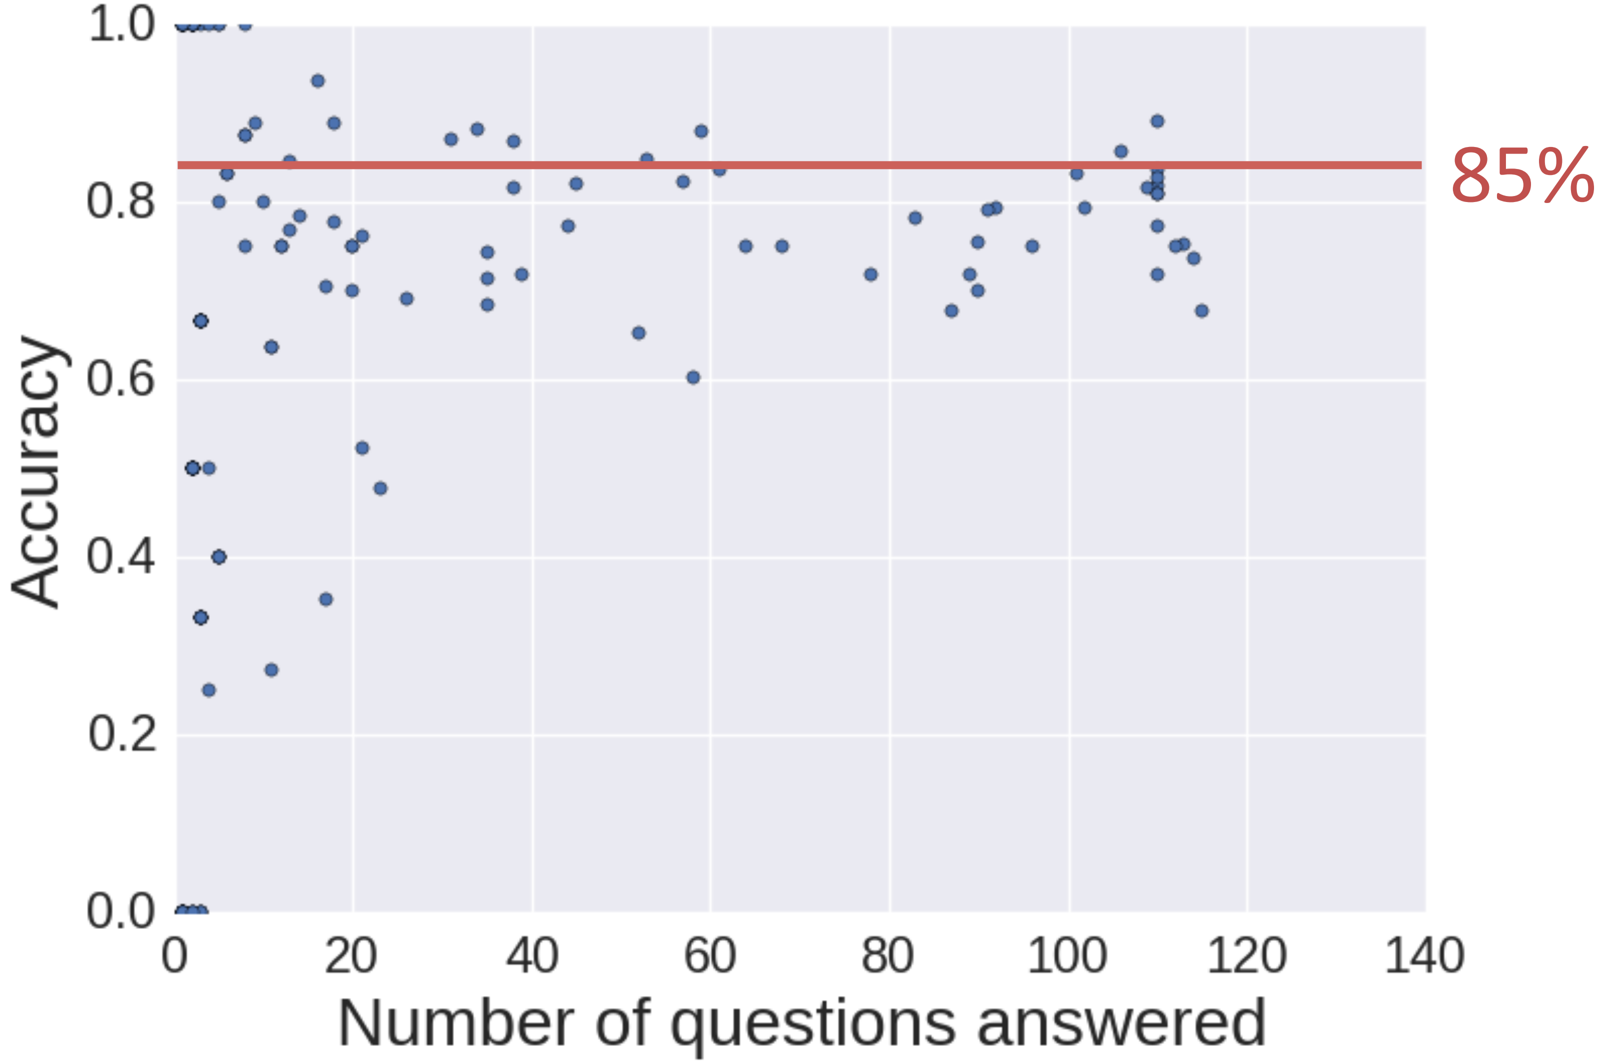

Related work

- Reinforcement learning in educational games [Mandel et al. AAMAS ’14, AAAI ’15, ’16]
- Task-centric control for crowdsourcing quality [Dai et al. AAAI ’10, ‘13; Kamar et al. AAMAS ‘12; Lin et al. AAAI ‘12; Bragg et al. HCOMP ’13]
- Worker-centric control for crowdsourcing engagement [Kobren et al. ‘15, Yin & Chen ’15]

Summary

- Problem: automatically insert gold questions to maximize quality & quantity of crowd work
- POMDP-based reinforcement learning algorithm
- Robust to parameter variations (see paper)
- Up to 111% more reward than common policies
- Code available online

Future directions

- Subjective tasks
- Multiple answers per question
- Worker training [1] & engagement

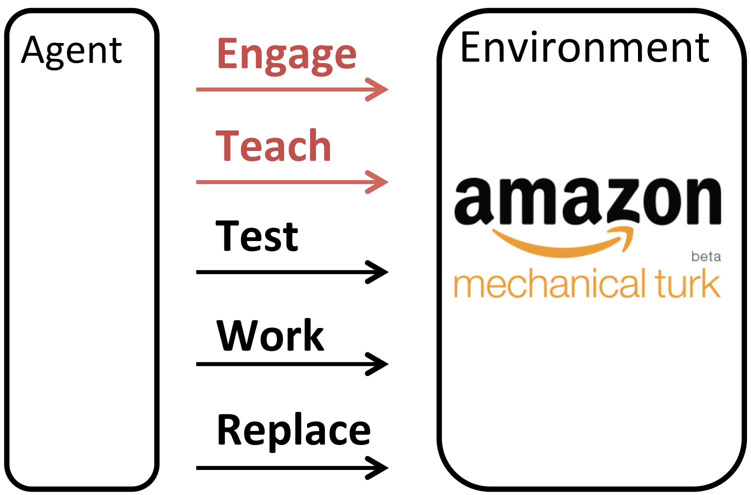

[1] Bragg et al. 2015. In ICML workshop.

Thanks!

- **Support:** NSF, ONR, WRF/Cable Professorship, Google, Bloomberg
- **Data:** Chris Lin, Shreya Rajpal, Karan Goel, workers on Mechanical Turk
- **Code:** https://crowdlab.cs.washington.edu/optimal-training-and-testing-for-crowd-workers