

# Planning to Control Crowd-Sourced Workflows

Daniel S. Weld  
University of Washington



**UW**  
CrowdLab

(Joint Work with Jonathan Bragg, Lydia Chilton, Peng Dai, Shih-Wen Huang, James Landay, Chris Lin, Angli Liu, Andrey Kolobov, Mausam & Stephen Soderland)



Jonathan  
Bragg



Lydia  
Chilton




Peng  
Dai



Shih-Wen  
Huang



James  
Landay



Chris  
Lin

# Thanks



Angli  
Liu



Andrey  
Kolobov



Mausam




Stephen  
Soderland



# Crowdsourcing

- Performing work by **soliciting effort** from many people
- **Combining the efforts** of volunteers/part-time workers (each contributing a small portion) to produce a large or significant result



# Crowdsourcing Successes



190 M reviews of 4.4 M businesses



Answers to 7.1 M prog. questions



Universal reference for anything

# Citizen Science



800,000 volunteers – Hubble images  
Discovered “Hanny’s Voorwerp” black-hole  
“Pea galaxies”

# eBird



Crowdsourced bird count & identification  
Migration shift -> effect of climate change



Game to find 3D structure of proteins.  
Solved 15 year outstanding AIDS puzzle


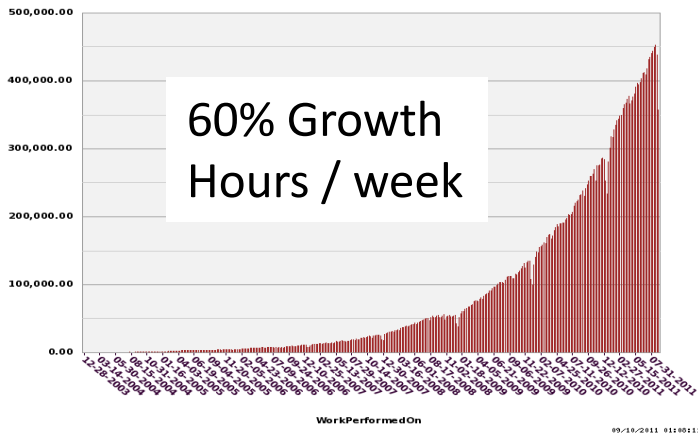
# Labor Marketplaces

Will Grow to \$5B by 2018 [Staffing Industry Analysts]



- 2.7 million workers
- 540,000 requestors
- 35M hours worked in 2012

spare5



# AI in Crowdsourcing

## Collective assessment

belief propagation  
variational inference

mean field approximation  
independent Bayes classifier

Mallows model

hierarchical clustering

Bayesian bias mitigation

two coin model

Bayes net

Bayesian aggregation

majority vote

Chinese restaurant process

HybridConfusion

EM

gold questions

multidimensional wisdom of crowds

ordinal-discrete mixture model

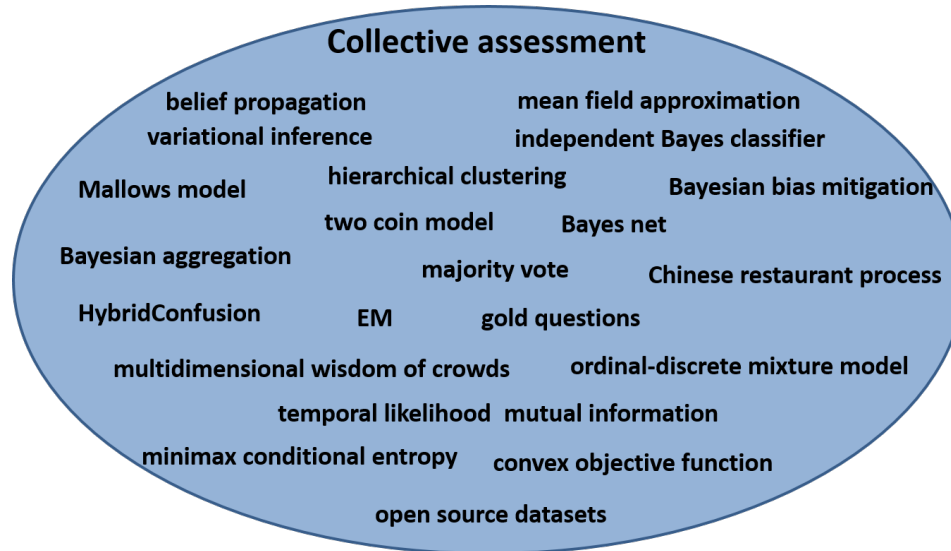
temporal likelihood mutual information

minimax conditional entropy

convex objective function

open source datasets

# AI in Crowdsourcing



- Collective assessment
  - State estimation/tracking
  - passive




# The Rest of Crowdsourcing


- control of simple tasks
  - optimize redundancy for best quality-cost tradeoff
- complex tasks
  - optimize workflows; pick the BEST
- task routing
  - finding the right workers
- make workers skilled
  - training; when? how much?

Sequential Optimization  
(Decision-Theoretic Control)

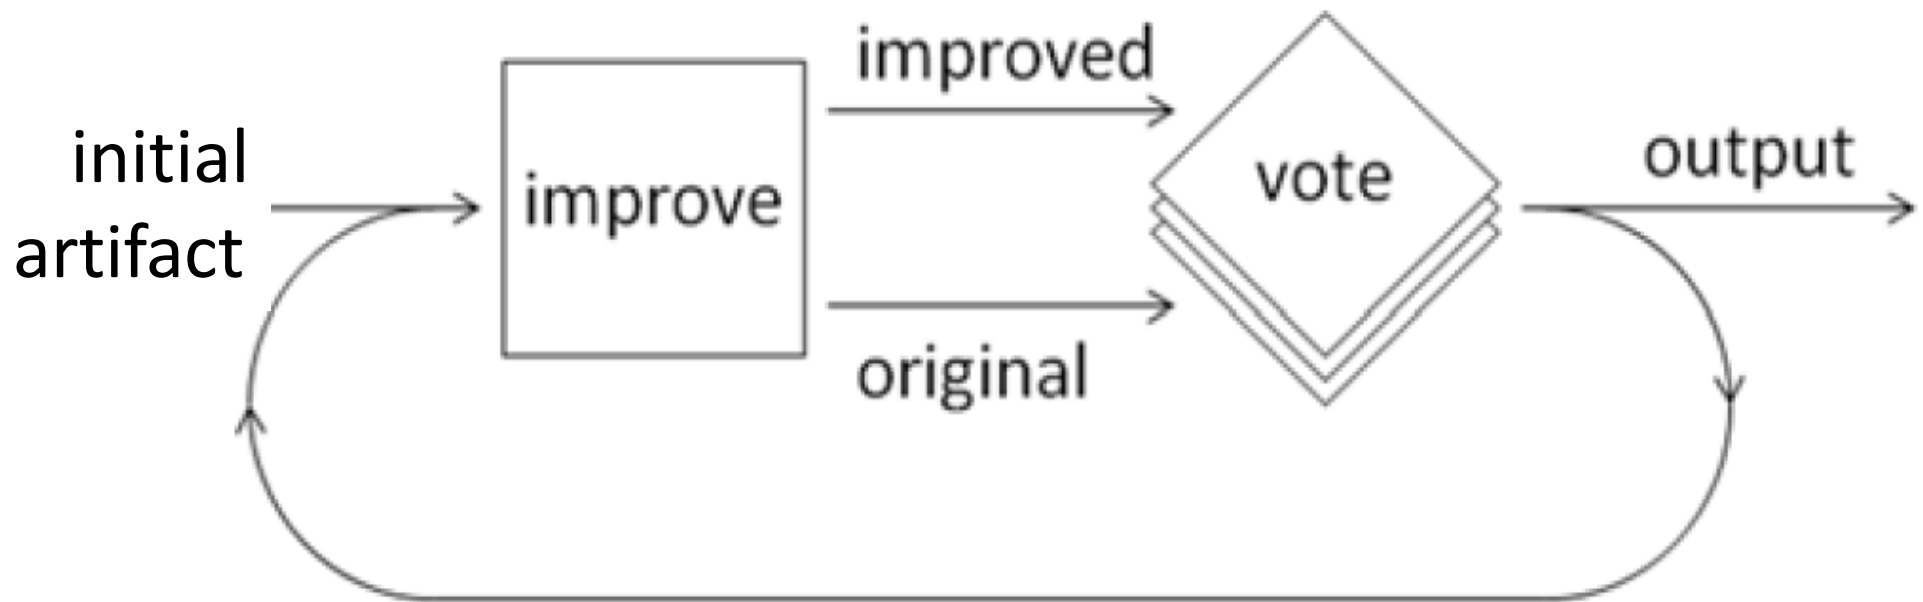
# Iterative Improvement Workflow



# Iterative Improvement Workflow




# Iterative Improvement Workflow



# Iterative Improvement Workflow

[Little et al, 2010]



## First version

A parial view of a pocket calculator together with some coins and a pen.



## Version after 8 iterations

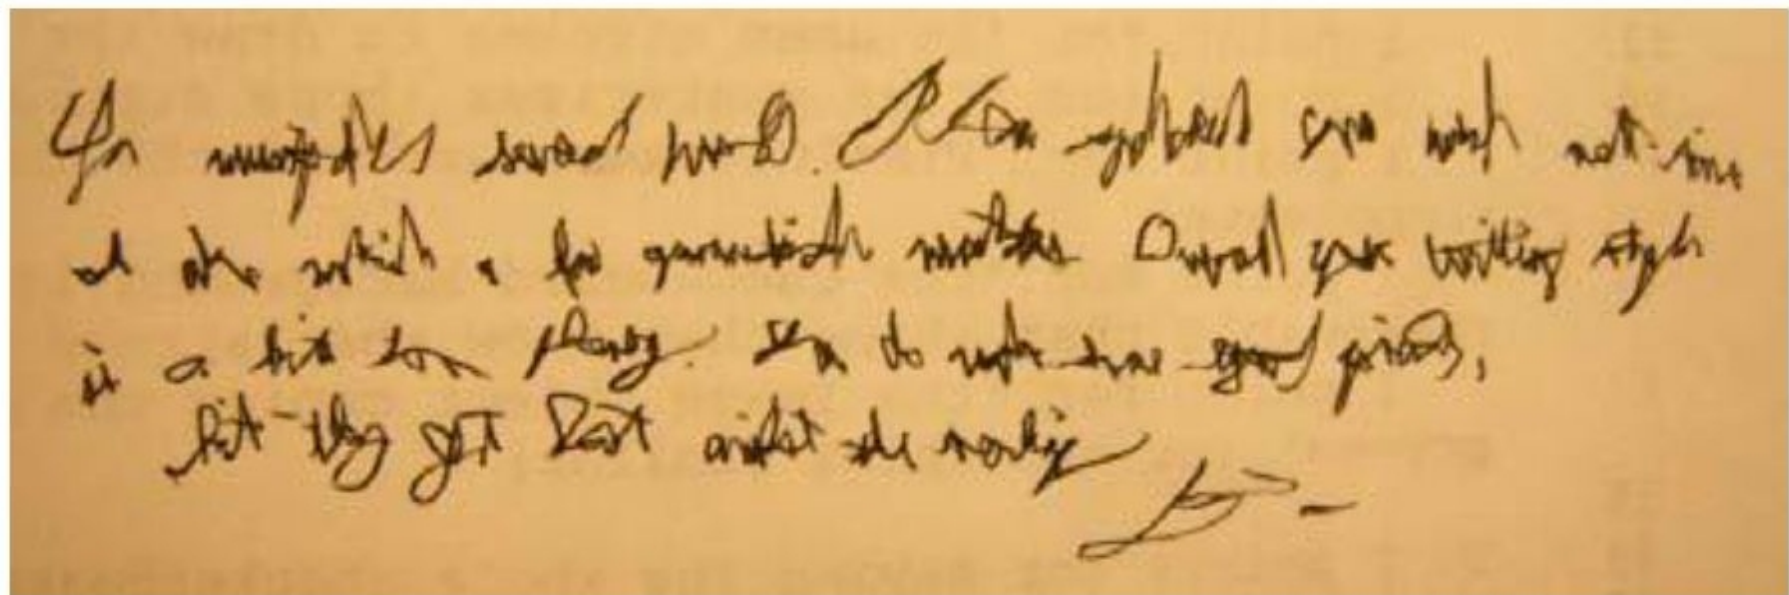
A CASIO multi-function, solar powered scientific calculator.

A blue ball point pen with a blue rubber grip and the tip extended.

Six British coins; two of £1 value, three of 20p value and one of 1p value.

Seems to be a theme illustration for a brochure or document cover treating finance - probably personal finance.

For ~~unmistakable~~ sweet words. How ~~pleasant~~ can such not-ions  
of the which a few ~~garment-like~~ ~~months~~ ~~Overall~~ ~~your~~ ~~writing~~ ~~style~~  
is a bit too ~~fluffy~~. So do not ~~lose~~ ~~your~~ ~~friend~~,  
let-thy get lost amidst the noise  
L.

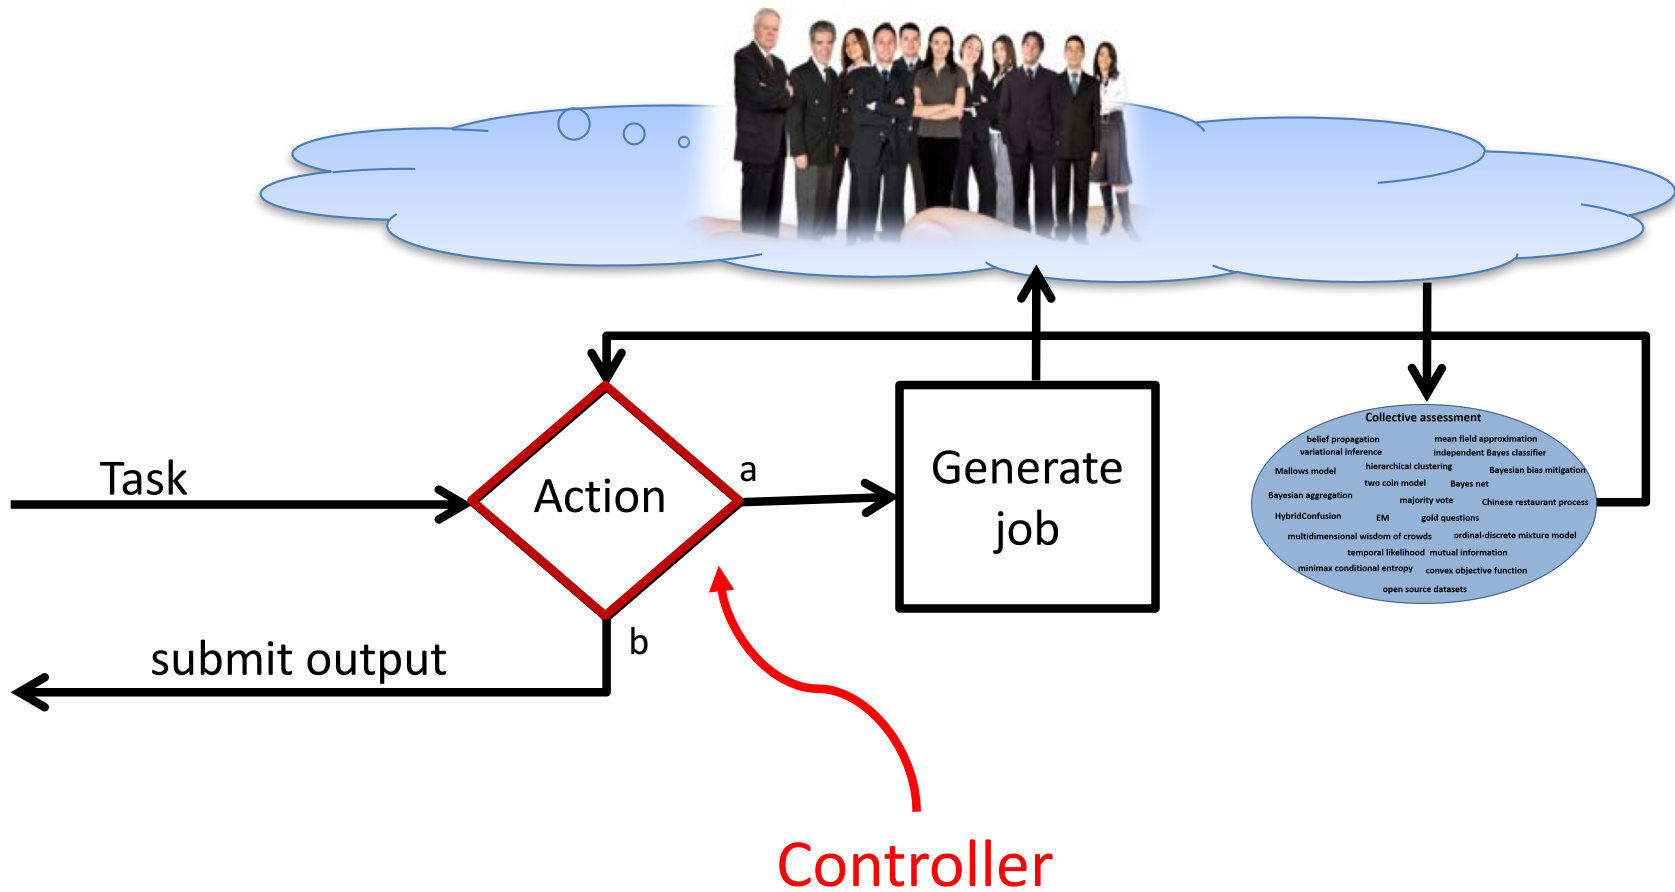


You misspelled several words. Please spellcheck  
your work next time. I also notice a few grammatical mistakes. Overall your writing style  
is a bit too phoney. You do make some good points,  
but they get lost amidst the verbiage  
B-

“You (misspelled) (several) (words). Please spellcheck your work next time. I also notice a few grammatical mistakes. Overall your writing style is a bit too **phoney**. You do make some good (points), but they **got** lost amidst the **(writing)**. **(signature)**”

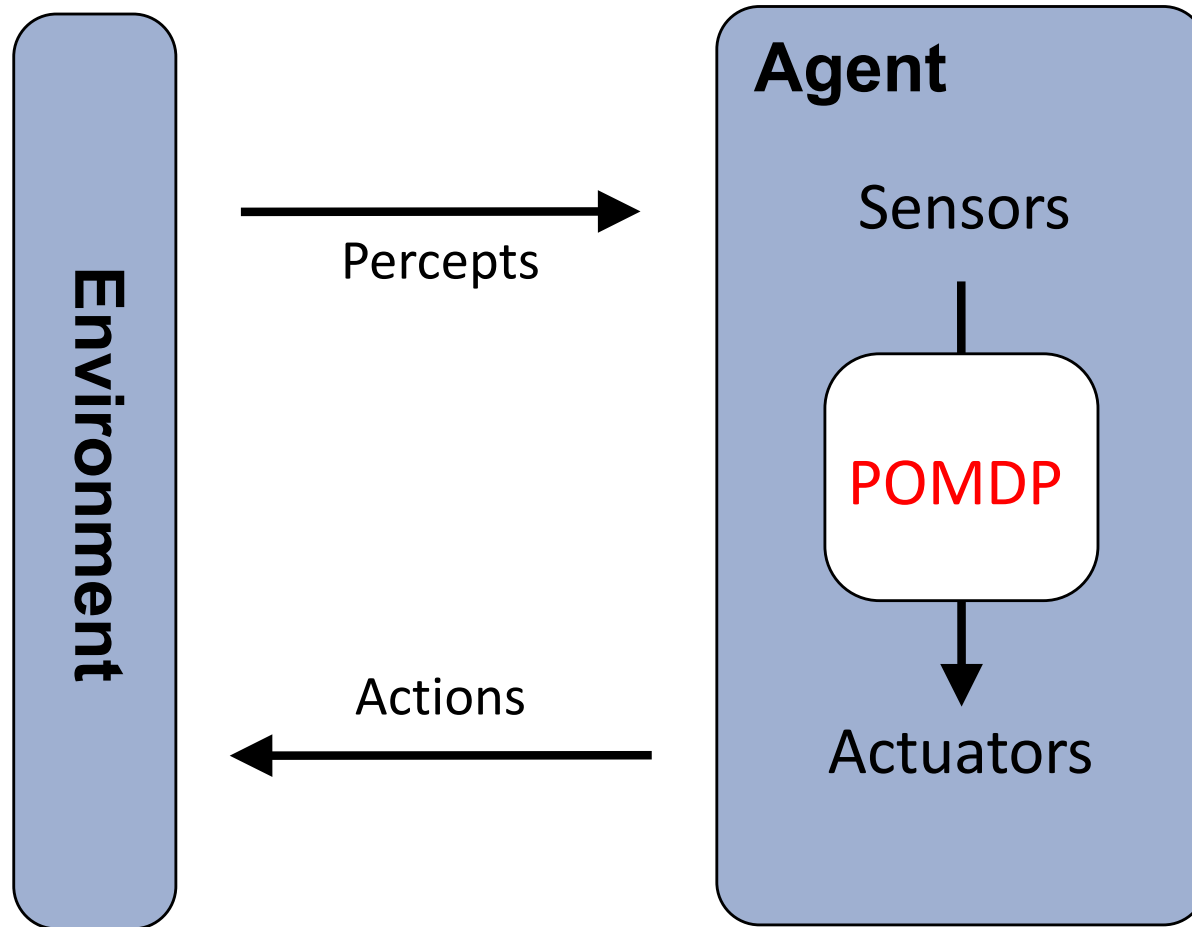
According to our ground truth, the highlighted words should be “flowery”, “get”, “verbiage” and “B-” respectively.

# Controller for a Task







# Artificial Intelligence 101



# Markov Decision Process




Input:



**World State**


$$s = \langle x, y \rangle$$



**Actions**

$$P(s' | s, a)$$

Cost  $c$



Observe: Next State  $s' = \langle x', y' \rangle$

$$\text{Reward} = f(s, a, s')$$


Output:

Construct **policy**,  $\pi : S \rightarrow A$ , that chooses best action for each state  
I.e., actions that **maximize expected reward – costs** over time


While **learning** action & reward probabilities  
(Reinforcement learning)

# Partially-Observable Markov Decision Process


Input:




**Belief State**  
 $P(s)$



**Actions**  
 $P(s' | s, a)$   
Cost  $c$



Observe: Noisy Sensor =  $f(s')$   
Reward



Output:

Construct *policy*,  $\pi : S \rightarrow A$ , that chooses best action for each state  
I.e., actions that *maximize expected reward – costs* over time

While *learning* action & reward probabilities  
(Reinforcement learning)

# Solving the POMDP


**Constructing the policy,  $\pi$ , to choose the best action**

- Many algorithms
  - Point-based methods
  - UCT on discretized space
  - Lookahead search with beta distribution belief states


$$Q^*(s, a) = \sum_{s'} P(s' | s, a) [ R(s, a, s') + \gamma \text{Max}_a Q^*(s, a) ]$$

- Exploration / exploitation problem
  - $\epsilon$ -greedy
  - UCB / Multi-armed bandit

From



To



*(Hidden)*  
**World State**

$\langle x, y \rangle$  coords


**Actions**

Move  
Grasp


**Costs**

Power used

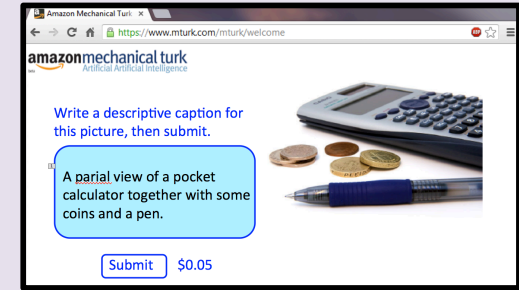
**Reward**



From



To



*(Hidden)*  
**World State**

$\langle x, y \rangle$  coords

**Actions**

Move  
Grasp

**Costs**

Power used

**Reward**

Quality  $Q_1, Q_2 \in (0, 1)$

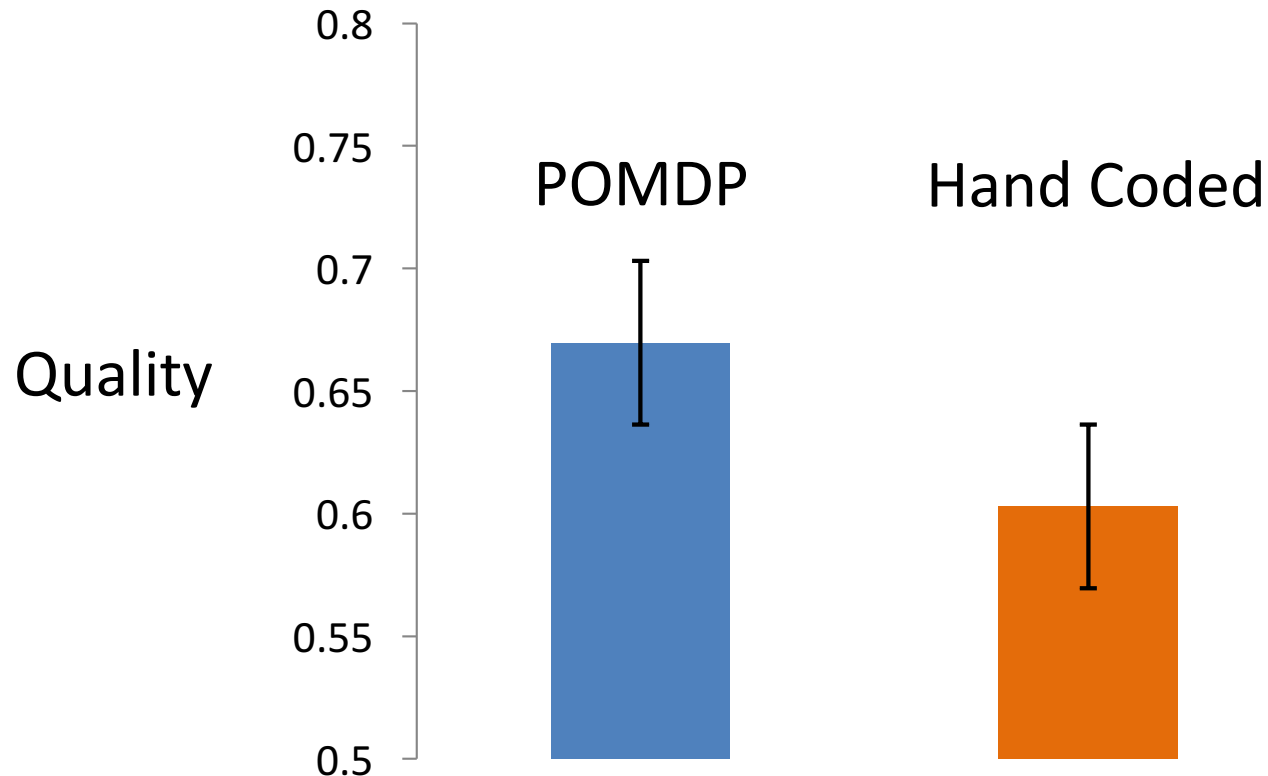
Improve caption task  
Vote best caption

\$\$ paid to workers

$F(\text{quality returned})$

# Comparison

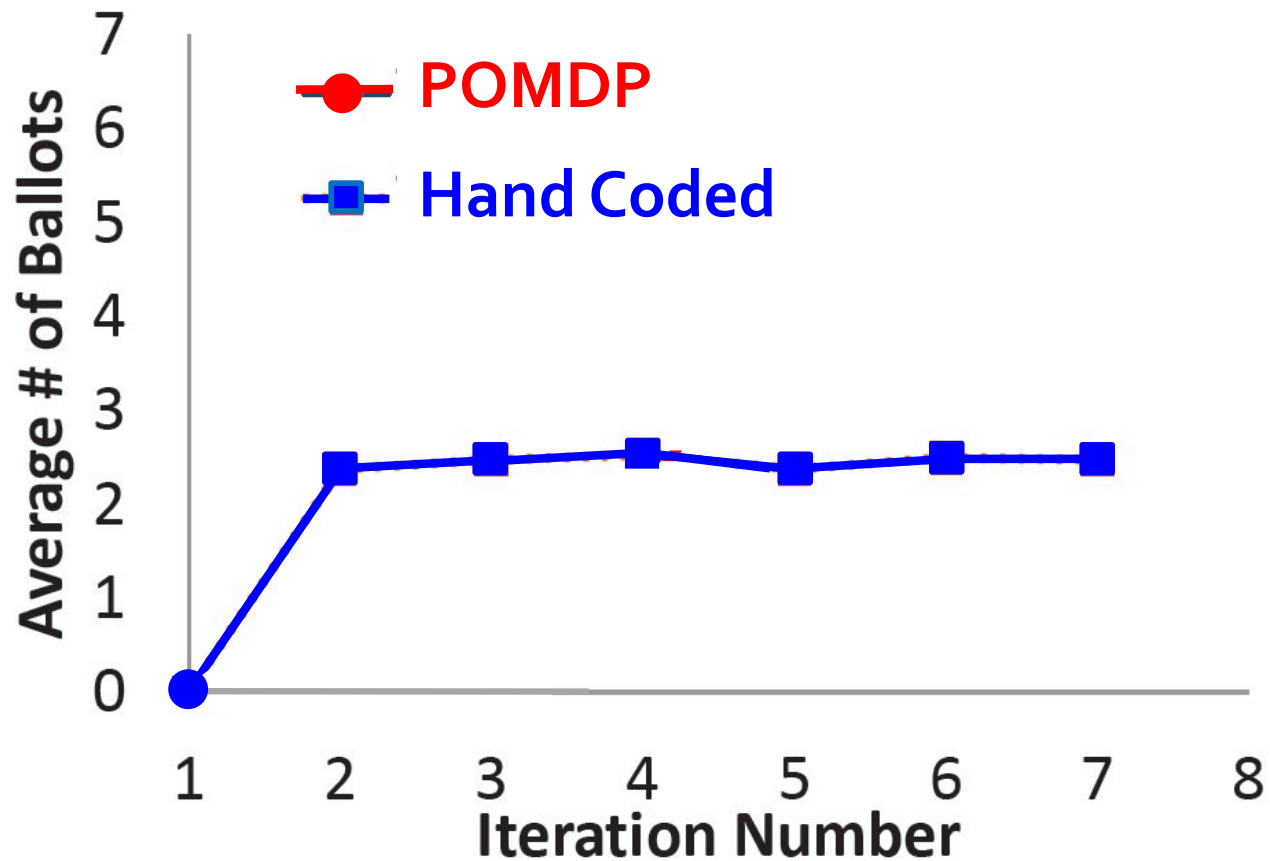
[Dai, Lin, Mausam, Weld AIJ'13]



40 images, same average cost

Controlling quality: ***POMDP 30% cheaper***

# Allocation of Human Labor





# Human Labor Redirected

