Reward Determination in Crowdsourcing

Kiatten Mittons

Nathan DeMaria | Rees Klintworth | Derek Nordgren

Citation

Amos Azaria, Yonatan Aumann, and Sarit Kraus (2014). Automated agents for reward determination for human work in crowdsourcing applications, autonomous agents and multiagent systems, 28(6):934-955.

Problem Introduction

Rees Klintworth

Research Goal

Automate the assignment of differing individual rewards in a crowdsourcing application

Crowdsourcing

- Goal broken into tiny increments of work
  - Too many to manually determine reward for each
- For simplicity, most problems are divided with equal rewards
  - Doesn’t make economic sense

Problem Setting

- Small, identical tasks must be completed
- Any one of available human workers can complete a task
- There are costs associated with bringing in new worker and making an offer

TCP

- TCP - Task Completion Problem
  - A = set of tasks
  - O = set of possible rewards (payment) for a task
  - T = set of types
    - Decision functions for workers
- TCP(A, pi) - minimize cost to satisfy all tasks in A, subject to type distribution pi

M-TCP

- Each task is broken up into milestones
- Each milestone for a task must be completed by the same worker
- A worker can leave mid-task
    - Compensated for milestones completed
- Upfront or Stepwise

Requester

- Makes an offer to a worker for a task
  - If declined, can either:
    - Move on to a new worker
    - Make another offer to the worker
- Has a history of offers and their acceptance/rejection
- Must assign each task to some worker

Workers

- Infinite amount of workers exist
- Any worker can complete any task
- Each worker has a decision policy governed by history
  - Rejects only
- Any cost for the worker is included in decision policy

Minimizing Cost

- The requester attempts to minimize the total cost associated with completing all tasks
- Doesn’t care about:
  - Best candidate
  - Cheapest candidate for a particular task

Problem is NP-Hard

- Maps to Set-Cover problem
- Set-Cover is a known NP-Hard problem
- Solution set grows exponentially large

Simplifying Assumptions

- Two restrictions considered in order to solve the problem
- RPBA - each worker type has a reservation price
- NBA - only one offer is made to each worker
- Different algorithm for each type

Bargaining Effect

- Basis for NBA
- If given a higher offer after declining a lower offer, you are less likely to accept the higher offer than if only offered the higher offer
- Explored in referenced work

Agents

Nathan DeMaria

Reservation Price-Based Agent (RPBA)

- Type of RPBA agents is defined by their reservation price
- Accept any offer above
- Reject any offer below
- Requester knows:
  - all reserve prices
  - frequencies of each reservation price

Optimal Algorithm for RPBA

- Offers must be exactly reservation prices
- For each reservation price
  - Estimate the cost of an interaction with that agent assuming you will negotiate up to that price, and then move on to the next agent if that maximum price is declined
- Find the reservation price that minimizes that cost, use the table for that price

RPBA Algorithm for Stepwise M-TCP

- Similar to basic RPBA, with an added layer of complexity
  - Offers need to be negotiated at each milestone
  - Requester knows frequencies of each type (a type is a unique combination of reservation prices for each milestone)
  - Based on milestone offers accepted/rejected so far, the requestor adjusts the probabilities that the agent is of each type

No-Bargaining Agent (NBA)

- Assumes workers will consider first offer
- Workers require a much higher price for second offer

Optimal Algorithm for NBA

u(x) = probability of worker accepting $x offer

Expected cost per worker

$$C_c + C_o + x \cdot u(x)$$

Expected cost per assignment

$$cost(x) = \frac{C_c + C_o + x \cdot u(x)}{u(x)}$$

Type Elicitation - RPBA

- Vickrey Auction
  - For crowdsourcing, use the last 3 bids
- Becker-DeGroot-Marschak mechanism
  - User submits a bid
  - Random number is generated
  - User is paid the generated number if it is greater than the bid
- Both incentivize truthful bidding
- Cluster bids into reservation prices

Type Elicitation - NBA

- Sigmoid function is assumed
- Pick test offers
- Observe % accepted at each offer
- Find u(x) by fitting a sigmoid function to those points

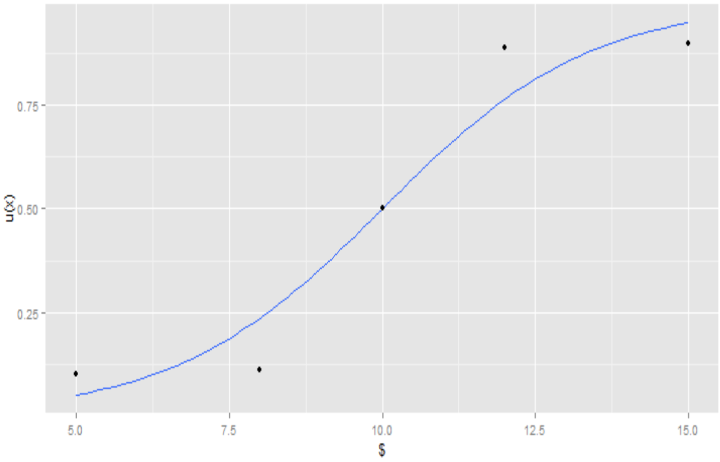

Experimentation

Derek Nordgren

Experimentation Setup

- Performed using Amazon’s Mechanical Turk service
- Both single milestone and 5-milestone tasks were simulated
- New worker cost = $0.20
- New offer cost = $0.04

Experimentation Setup

- Re-performed experiment from Related Work
- Performed new experiment
- RPBA, NBA and human “experts” performance were analyzed

Experiment 1

- Previously published experiment
- Workers are required to identify a unique shape

Experiment 2

- Workers are required to encode a paragraph of text
- TCP - Requester must accomplish 25 tasks
- M-TCP - Requester must accomplish 25, 5-milestone tasks

Results - Aggregated across Experiments

- The NBA performs best, though only slightly better than RPBA
- Both agents significantly outperform human “experts”

Results, cont’d

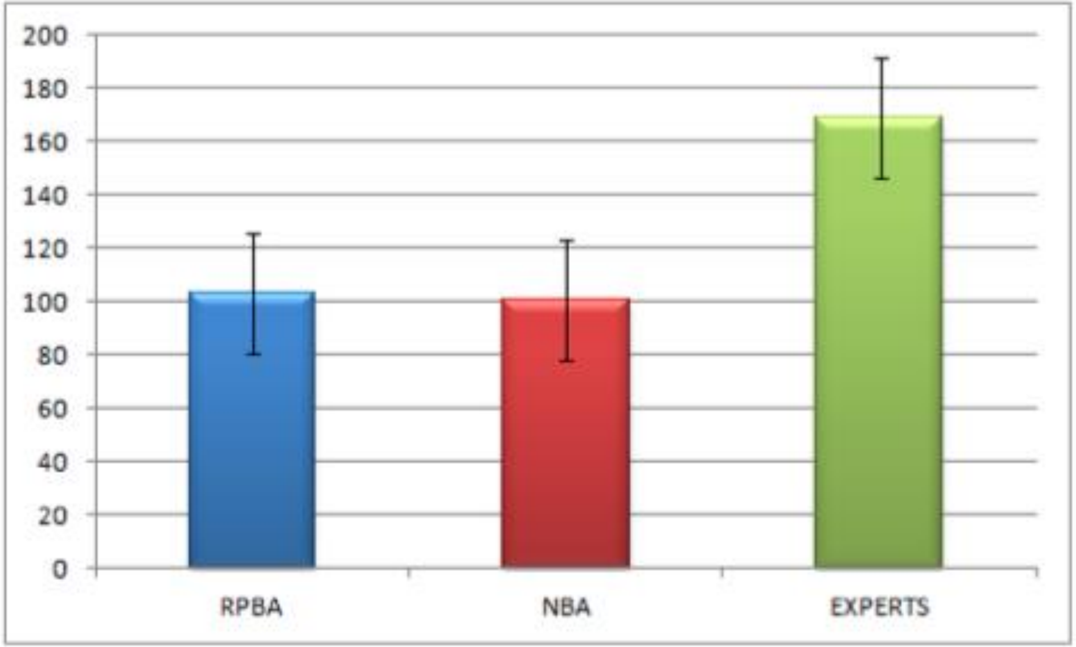

Fig. 3 Average cost per goal over all three settings

Results, cont'd

Table 6 Average cost per goal over all three settings

agent	average cost
RPBA	102.9
NBA	100.7
EXPERT #1	118.2
EXPERT #2	121.1
EXPERT #3	294.1
EXPERT #4	142.5

Agent Recommendation

The authors recommend the NBA agent

- Easier to implement
- All workers are paid identically (more fair)
- Data collection (sigmoid calibration) is simpler than RPBA’s auction process

Agent Recommendation

The authors recommend the NBA agent

- May significantly outperform the RPBA in some conditions
- Every interaction builds a more accurate strategy

Merits of Bargaining

In general, bargaining was found to be “fruitless.”

- Adds to expense
- Does not improve completion rate

NBA avoids bargaining altogether.

Sunk Cost Effect

NBA also avoids “sunk cost effect”

- In M-TCPs, human’s cost was much higher
- People tend to spend more money when money has already been spent
- Requires more research

Reward Determination Schedule

Both agents performed much better with upfront reward determination schedule.

- Reduced costs ~ 15%
- Requester commitment lowers cost required by workers
- Human agents incurred more costs upfront

Cost Modification

- A higher Cc results in higher offers
- A lower Co means more offers from RPBA
- Changes in Co do not affect NBA

Future Work

Incorporate milestone repetitiveness into NBA to account for worker

- Expertise
- Boredom

Questions?