EPSRC Engineering and Physical Sciences Research Council

ORCHID

Efficient Crowdsourcing of Unknown Experts using Multi-Armed Bandits

Long Tran-Thanh, Sebastian Stein, Alex Rogers and Nicholas R. Jennings
Agents, Interaction and Complexity Research Group
Electronics and Computer Science
University of Southampton

Motivation

Crowdsourcing systems:

combine machine intelligence with mass of human intelligence

ZOONIVERSE

amazon mechanical turk
Artificial Artificial Intelligence

Expert crowdsourcing:

New challenges: quality of work matters

- Hiring a worker is costly
- Limited by a financial budget

Existing systems: oDesk, freeLancer, eLance, geniousRocket, blurGroup

Offline version: Bounded knapsack

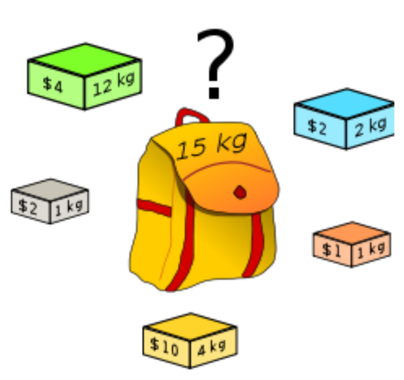

$$\max \sum_{i=1}^{N} x_i v_i \quad \text{s.t.} \quad \sum_{i=1}^{N} x_i w_i \le C,$$

$$\forall i: \quad 0 \le x_i \le L_i \quad x_i \text{ integer}$$

NP-hard problem → Approximation algorithms

Density ordered greedy approach

Performance analysis

Theorem 1 *Let $0 < \varepsilon, \beta < 1$. Suppose that $\varepsilon B \ge \sum_{j=1}^{N} c_j$. With probability $(1-\beta)^N$, the performance regret of the bounded ε–first approach is at most*

$$2 + \varepsilon B d_{\max} + 2N\left(\sqrt{\frac{B\left(-\ln\frac{\beta}{2}\right)\sum_{j=1}^{N} c_j}{\varepsilon}}\right), \qquad (4)$$

where $d_{\max} = \max_{i \ne j}\left|\frac{\mu_i}{c_i} - \frac{\mu_j}{c_j}\right|$ (i.e. the largest distance between different density values).

Theorem 2 *Let ε_{opt} denote the value that minimises Equation 4. By setting the exploration budget to be $B\varepsilon_{\text{opt}}$, the regret of the bounded ε–first algorithm is at most*

$$2 + 3B^{\frac{2}{3}}\left(N^2\left(-\ln\frac{\beta}{2}\right)\sum_{j=1}^{N} c_j d_{\max}\right)^{\frac{1}{3}}.$$

Model Description

Bounded Multi-Armed Bandit model

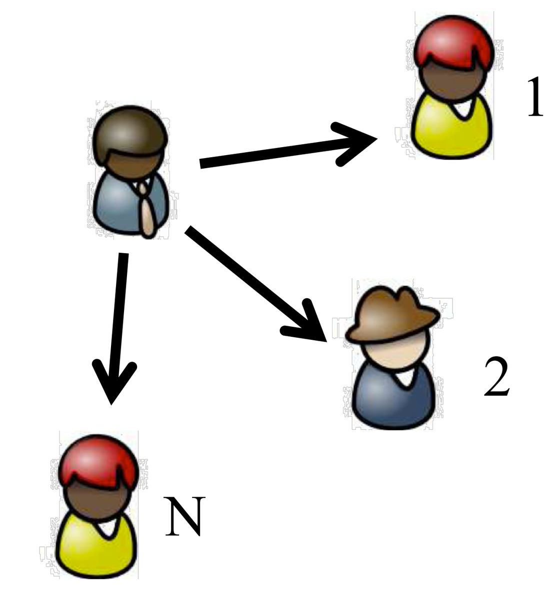

- N workers, financial budget B
- Each worker i:
  - Hiring cost c_i
  - Working task limit L_i
  - Quality of work: drawn from *unknown* distribution with expected value μ_i

Objective: maximise the total quality of assigned work, w.r.t. budget B

$$A^* = \arg\max_A \sum_{i}^{N} \mathbb{E}\left[N_i^B(A)\right]\mu_i. \quad P\left(\sum_{i}^{N} N_i^B(A)\, c_i \le B\right) = 1$$

Online learning

Bounded Epsilon-first: $\varepsilon B \leftrightarrow (1-\varepsilon)B$

- Exploration: uniformly assign tasks to workers

  $\hat{\mu}_i$: estimate of worker i's expected quality

- Exploitation: form a bounded knapsack from estimates + solve with density ordered greedy

$$\max \sum_{i=1}^{N} \hat{\mu}_i x_i^{\text{exploit}} \quad \text{s.t.} \quad \sum_{i=1}^{N} c_i x_i^{\text{exploit}} \le (1-\epsilon)B,$$

$$\forall i: \quad 0 \le x_i^{\text{exploit}} \le L_i - x_i^{\text{explore}}$$

Empirical Evaluation

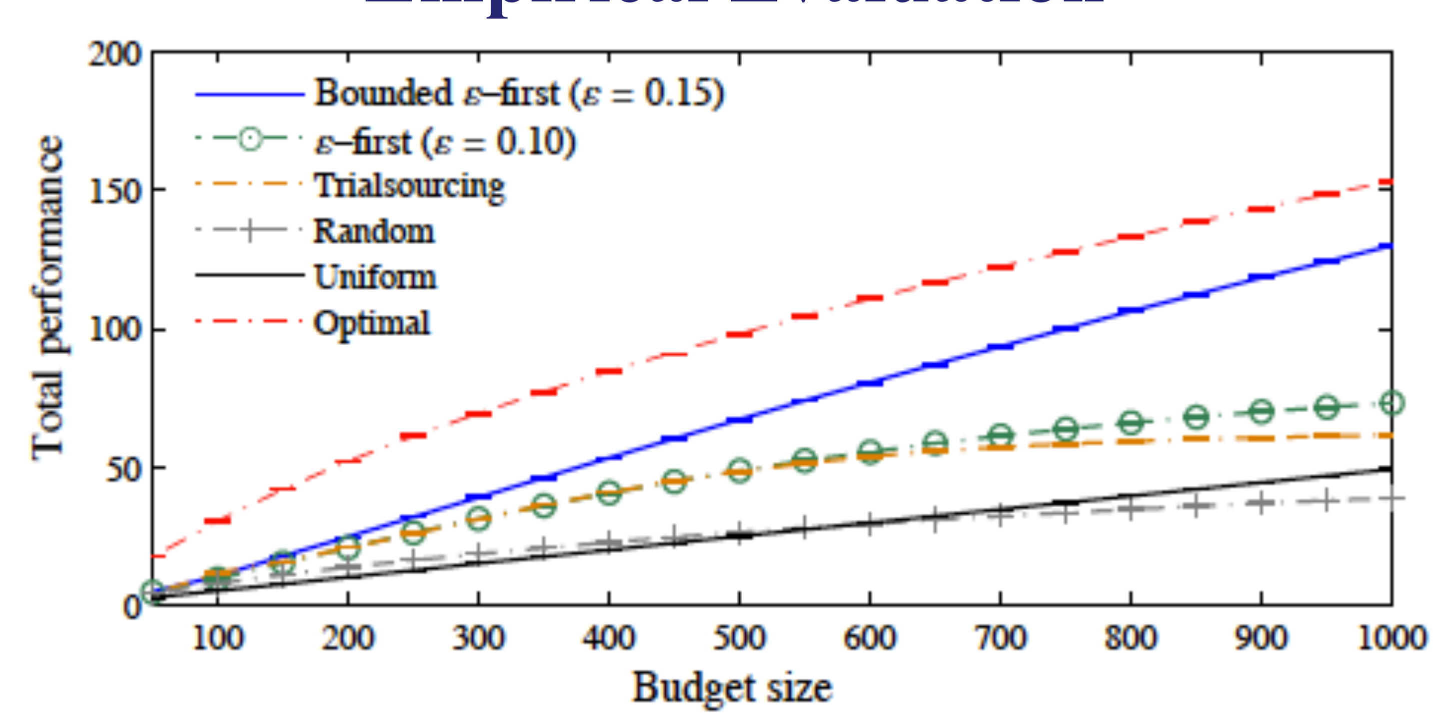

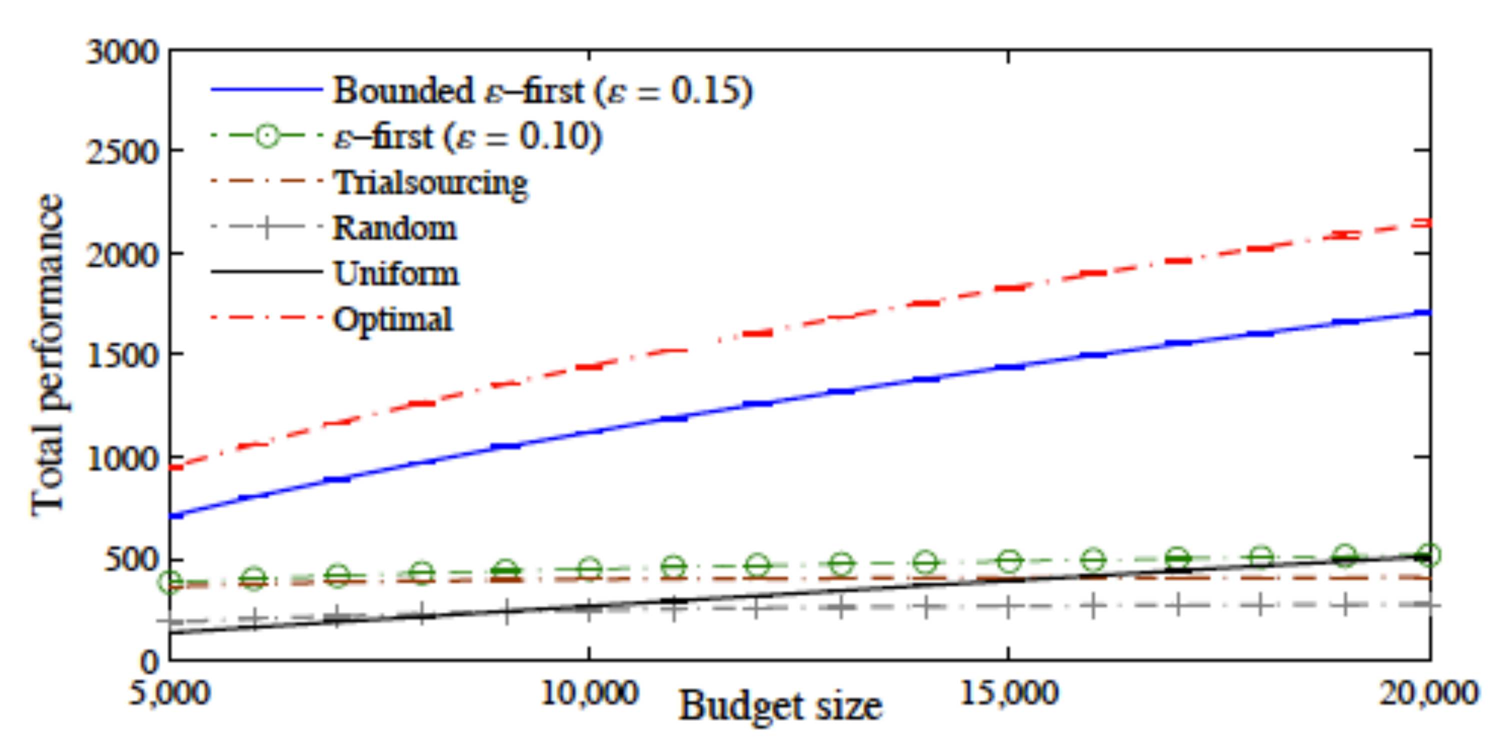

BAE SYSTEMS ACFR AUSTRALIAN CENTRE FOR FIELD ROBOTICS SECURE UNIVERSITY OF Southampton