Ariana Tang, Naveen Raman, Fei Fang, Ryan Shi

Contextual Budget Allocation for Food Rescue Volunteer Engagement

1. Background: Food Rescue as Contextual-RMAB

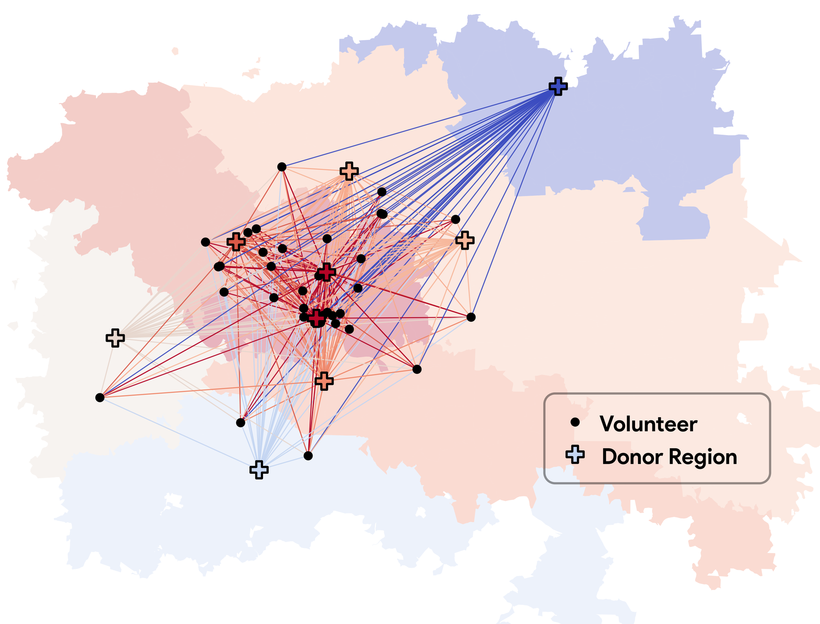

🍲 Existing algorithms improve volunteer engagement but worsen geographical disparities.

🥗 Food Rescue Platforms (FRPs) match surplus food to low-resource communities.

🍔 Volunteers are essential, but engagement
is uncertain and uneven across regions.

🍩 Our goal: engage volunteers efficiently and fairly across all regions.

Key challenge: A dynamic, resource-limited allocation problem
→ modeled as a contextual restless multi-armed bandit (RMAB).

2. Model & Preliminaries: Contextual-RMAB

Standard RMAB:

N arms (volunteers), binary state (active/inactive).

Each step: choose each arm's action {0,1} under total budget B.

Objective: maximize long-term reward across arms.

Contextual-RMAB:

Adds K contexts (e.g. regions).

$$\langle N, \mathcal{S}, \mathcal{A}, K, \{r_i^k\}, \{P_i^k\}, \mathcal{F} \rangle$$

Context k is chosen i.i.d. with a known prior distribution at each time step.

Context-specific budgets:

$\sum a_i^t \le B_k$, with $\mathbb{E}[B_k] \le B$ for all k

🍱 Food Rescue Mapping of **Contextual RMAB**:
Volunteers = arms; Regions = contexts; State = active/inactive volunteer.
Action = notify volunteer about trip.
Reward = successful food pickup.
Transition dynamics depend on distance, history, and volunteer types

Food rescue platforms rely on volunteers, but engagement is unfair.
We model volunteer allocation as a contextual RMAB with region-specific budgets.
Our Mitosis algorithm ensures fairness while maximizing food rescued.

3.5 The Mitosis Algorithm

Mitosis solves a **zero-order combinatorial optimization problem** over budget allocations in contextual RMABs, combining no-regret exploration with a branch-and-bound tree guided by LP upper bounds.

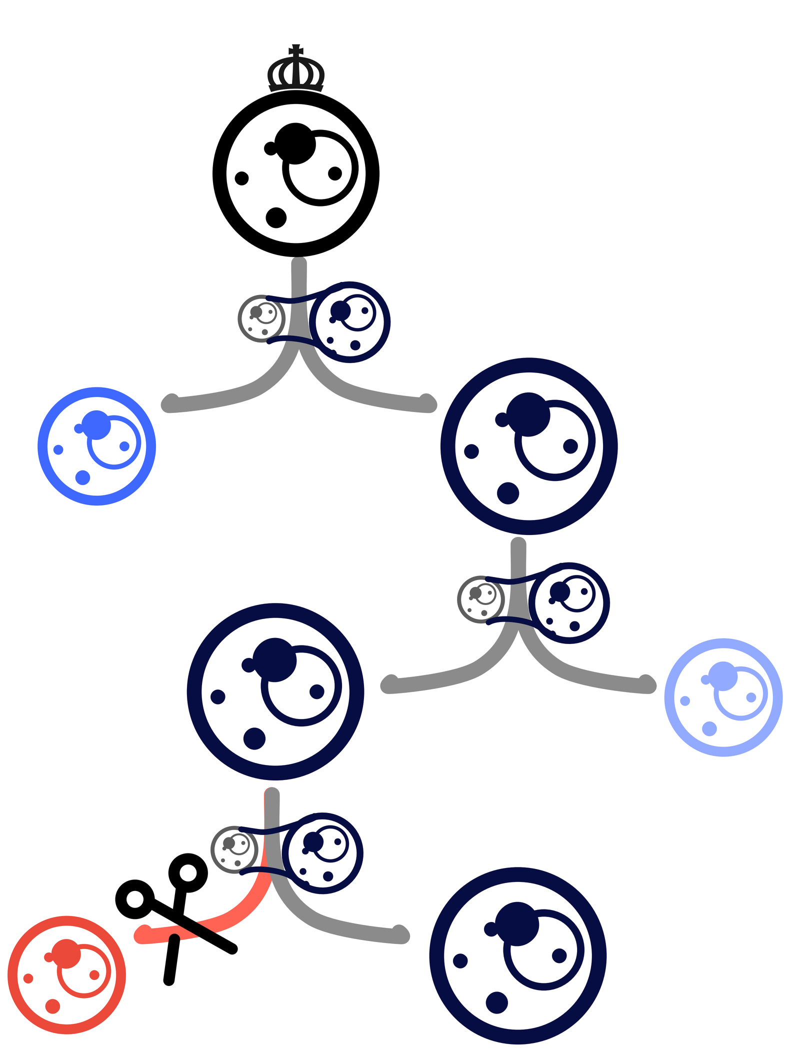

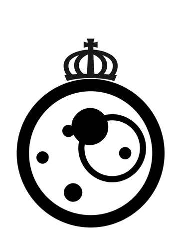

Root node: the full feasible region

Arms are candidate budget allocations explored via **UCB**.

UCB-guided exploration:
Pull the most promising arms first.

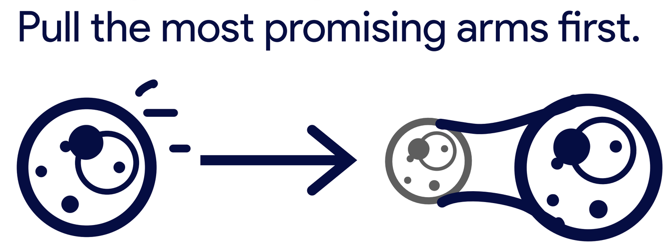

Root node **split into a new arm** when it's pulled.

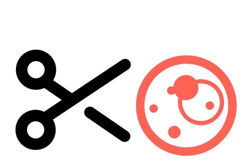

Arms are cut if too small compared to its LP-upperbound

Mitosis achieves optimal allocation with no-regret guarantee for contextual-RMAB, despite the solution space's combinatorial structure (Theorem 3)

Mini-comparison with Kleinberg (2008)'s Zooming Algorithm for continuous MAB

	Mitosis	Zooming Algorithm
Problem type	Zero-order RMAB optimization	Lipschitz MAB in metric space
Search Space	Discrete, exponential arms	Continuous / metric balls
Tree Growth	Guided by LP upper bounds of	Guided by metric smoothness
Guarantees	No-regret (UCB1-tight)	No-regret

3. Theory Results: Cocc & Mitosis

Definition Whittle Index (intuitive): For each volunteer i in state s, the Whittle index w_i(s) is the marginal subsidy at which being active or idle yields equal long-term value, i.e. it quantifies the priority of notifying that volunteer by balancing immediate reward and future engagement.

Algorithms for Contextual-RMAB:

Standard Whittle Index Algorithm (context-agnostic RMAB policy).
→ Can be arbitrarily bad in contextual-RMAB (Theorem 1).

Cocc (Contextual Whittle Index) – LP-based heuristic with context-dependent budgets.
→ Approximation ratio at most 5/6 of optimal (Theorem 2).

Mitosis – bandit-based search for budget allocation.
→ Achieves optimal allocation with no-regret guarantee (Theorem 3).
→ Computationally efficient: matches Branch-and-Bound's optimality at a fraction of the cost.

4. Applications & Experiments

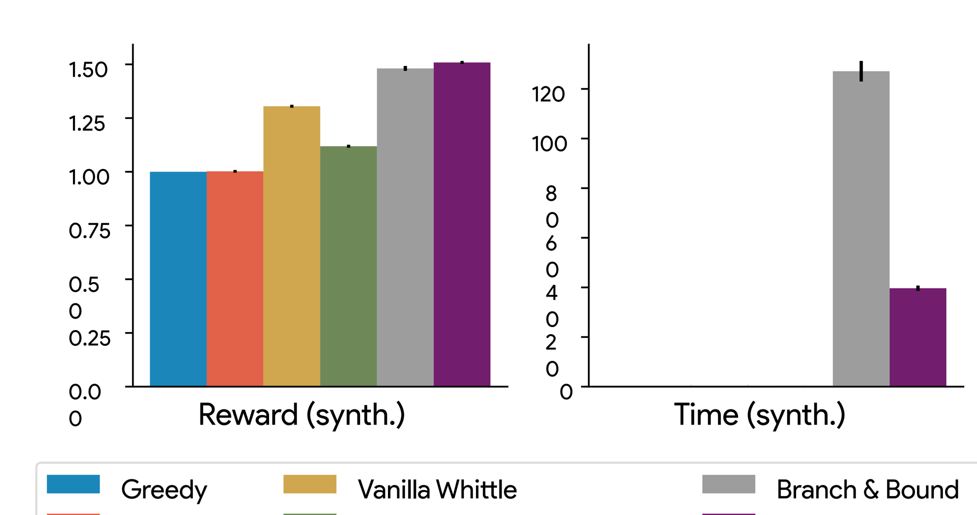

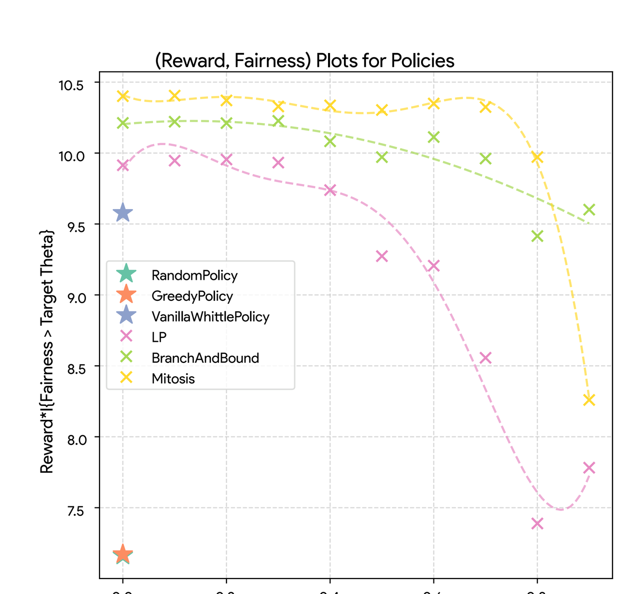

Pareto frontiers for (proportional) fairness. Mitosis achieves higher fairness levels

Real-World Food Rescue Data:
Sampled from >500,000 volunteers across Penn.
Context-aware algorithms (Cocc, Mitosis) outperform random, greedy, vanilla Whittle.
Mitosis ≈ optimal reward, with computation far lower than Branch-and-Bound.