Fairness and Efficiency in Online Class Matching

MohammadTaghi Hajiaghayi [1] Shayan Chashm Jahan [1] Mohammad Sharifi [2] Suho Shin [1] Max Springer [1]

[1]University of Maryland, College Park, USA [2]Sharif University of Technology, Tehran, Iran

Online Matching

The online matching problem is a fundamental challenge in algorithm design where vertices from one set (e.g., buyers, jobs) arrive sequentially and must be matched to a fixed set of offline vertices (e.g., items, workers) without future knowledge. Decisions are irrevocable, requiring algorithms to optimize objectives such as maximizing total weight or ensuring fairness under constraints.

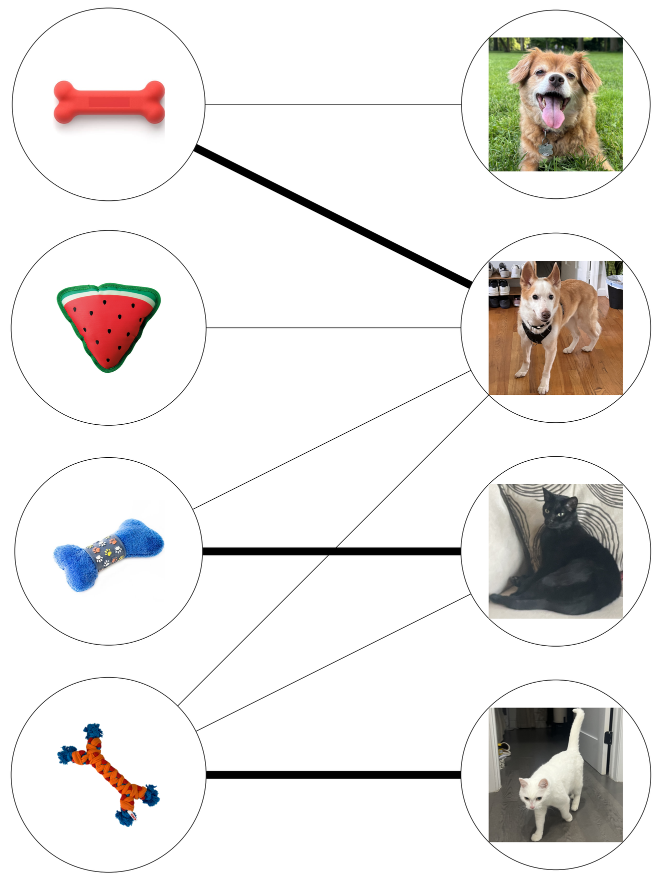

No randomized algorithm can achieve a better competitive ratio of $1 - 1/e$ in expectation. The RANKING algorithm gives this guarantee.

Online Fair Matching

In real-world scenarios, fairness and diversity often require considering ”group” or ”class” constraints.

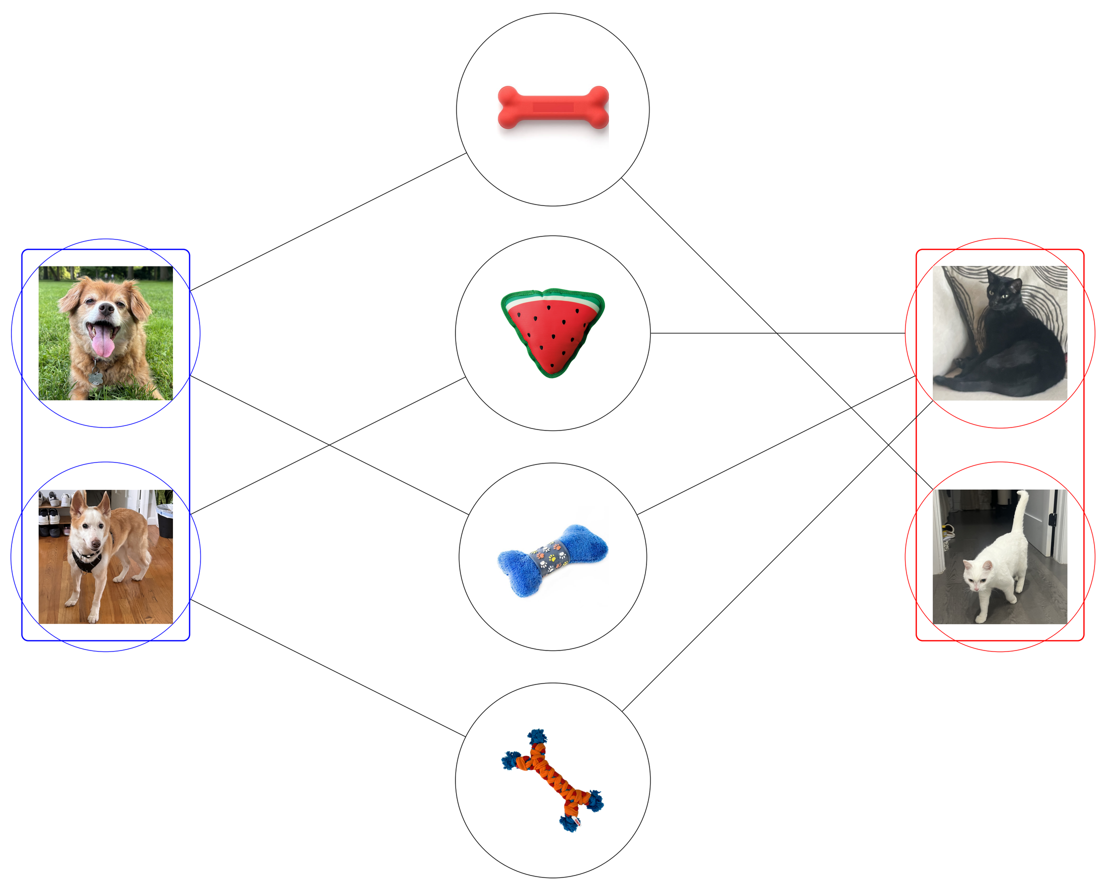

Main Contributions

- **Impossiblity Result 1:** No non-wasteful algorithm can achieve α-CEF in expectation for $\alpha > \frac{e^2-1}{e^2+1}$, or $1 - 1/e$-CMMS.
- **Impossibility Result 2:** For any $\varepsilon > 0$, there exists a problem instance such that no (possibly randomized) non-wasteful online algorithm with an α-CEF guarantee can achieve an approximation to the USW objective greater than $\frac{1}{1+\alpha} + \varepsilon$.
- **Impossibilty Result 3:** No non-wasteful *divisible* algorithm can achieve α-CEF for $\alpha > 0.677$.
- **Algorithmic Results:** Non-wasteful random assignment satisfies non-wastefulness, $1/2$-CEF, $1/2$-CMMS, and $1/2$-USW.

Problem Formulation

We study online matching in a bipartite graph $G = (N, M, E)$, where N (agents) is partitioned into k known classes $N_1, \ldots, N_k$, and M represents items [2]. Agents $a \in N$ and items $o \in M$ are adjacent if $(a, o) \in E$. A matching is represented by $X = (x_{a,o})$, where $x_{a,o} \in \{0, 1\}$ indicates whether o is matched to a. For divisible matchings, $x_{a,o} \in [0, 1]$.

Fairness Definitions:

- A matching is α-CEF if for all classes $i, j \in [k]$, $v_i(A_i) \geq \alpha \cdot v_i^*(A_j)$ where $v_i^*(\cdot)$ denote class i's “optimistic” value.
- Let μ_i be the value for each class when asked to maximize the worst partition it can create when optimally allocating items $\mu_i = \max_{M \in \mathcal{M}} \min_{j \in [k]} v_i^*(X_j)$ where $\mathcal{M}$ is the set of all indivisible matchings. A matching is α-CMMS if for every class $i \in [k], v_i(A_i) \geq \alpha \cdot \mu_i$.

Efficiency Definition: Maximize the number of matched agents (so-called “utilitarian social welfare”)

Impossibility Results

No better than $1 - 1/e$-CEF:

Consider $k = 2$ classes of n and n arriving items:

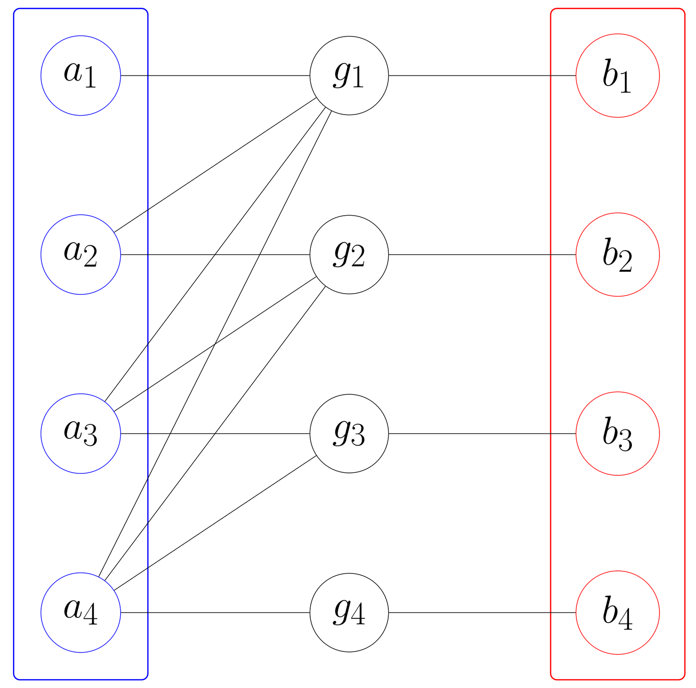

Price of Fairness:

- For some α, choose integers p, q such that $|\alpha - \frac{p}{q}| < \epsilon$.
- k classes: classes 1 to $k - 1$ have q agents, the k-th class has $q(k - 1)$.

Adversarial input sequence:

- Phase One:
  - $p(k - 1) + q$ items arrive, each item is liked by every agent (edges to all agents).
- Phase Two:
  - $k - 1$ groups of items arrive sequentially.
  - **Group** i: Contains q items, each item is liked only by agents in **Class** i.

Algorithm Overview

Algorithm 1 Random

```
for o ∈ M do
    S_o ← ∅
    for i ∈ [k] do
        if ∃a ∈ N_i s.t. (a, o) ∈ E and x_a = 0 then
            S_o ← S_o ∪ {i}
        end if
    end for
    Pick an i ∈ S_o uniformly at random
    Pick an a ∈ N_i with (a, o) ∈ E and x_a = 0 uniformly at random
    Set x_{a,o} = 1
end for
```

Key Takeaways

- We provide the first non-wasteful algorithm that simultaneously obtains approximate class fairness guarantees in expectation, resolving a major open problem posed by [2] in an affirmative manner.
- We demonstrate that the expected guarantees are almost tight, but conjecture that the optimal competitive ratio for randomized algorithms is $1 - 1/e$.
- We provide a strengthened upper bound for the *divisible* matching setting, further resolving an open problem left by [2].
- Defined the “price of fairness,” the necessary trade-off between an optimal and a fair matching.

Proof Idea

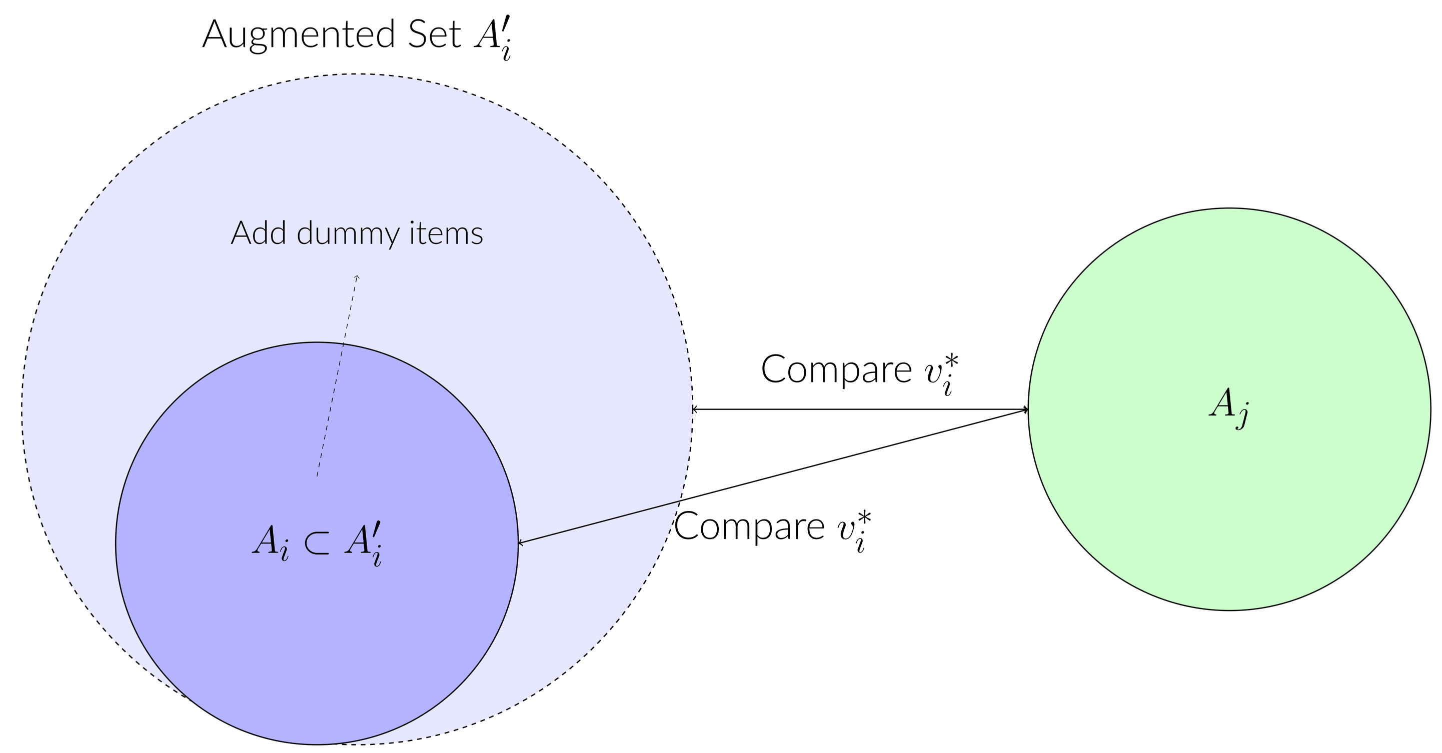

1. $v_i^*(A_i') \leq 2 \cdot v_i(A_i)$
2. $\mathbf{E}[v_i^*(A_i')] \geq \mathbf{E}[v_i^*(A_j)]$
3. $\Rightarrow \mathbf{E}[v_i(A_i)] \geq \frac{1}{2} \cdot \mathbf{E}[v_i^*(A_j)]$

Acknowledgements

This work is partially supported by DARPA QuICC, NSF AF:Small #2218678, NSF AF:Small #2114269, Army-Research Laboratory (ARL) #W911NF2410052, and MURI on Algorithms, Learning and Game Theory.

Max Springer was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 1840340. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

The full version of our paper can be found on arXiv with the QR-code above.

References

[1] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

[2] H. Hosseini, Z. Huang, A. Igarashi, and N. Shah. Class fairness in online matching. *Artificial Intelligence*, 335:104177, 2024.

[3] R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 352–358, 1990.

[4] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 125–131, 2004.