Optimizing Relevance and Diversity in Online Matching Markets: A Time-Adaptive Attenuation Approach

EVAN YIFAN XU*, Southeast University, China
PAN XU, New Jersey Institute of Technology, USA

Real-world online matching markets (OMMs) often involve multiple objectives, such as maximizing relevance and diversity in online recommendation and crowdsourcing systems. In this paper, we propose a generic bi-objective maximization model for OMMs with the following features: (1) there are two types of agents—offline and online—with online agents arriving dynamically and stochastically; (2) upon each online agent's arrival, an immediate and irrevocable decision must be made regarding which subset of relevant offline agents to assign; and (3) each offline and online agent has a specific matching capacity, i.e., an upper bound on the number of allowable matchings. Our model supports two general linear objective functions defined over all possible assignments to online agents. We formulate a bi-objective linear program (LP) and design an LP-based parameterized algorithm. Departing from prevalent non-adaptive attenuation methods, we introduce a time-adaptive attenuation framework that achieves an almost tight competitive ratio for each objective. To complement our theoretical analysis, we implement the proposed algorithm and evaluate it against several heuristics using two real-world datasets. Extensive experimental results demonstrate the flexibility and effectiveness of our approach, validating our theoretical predictions.

JAIR Associate Editor: Yair Zick

JAIR Reference Format:
Evan Yifan Xu and Pan Xu. 2025. Optimizing Relevance and Diversity in Online Matching Markets: A Time-Adaptive Attenuation Approach. *Journal of Artificial Intelligence Research* 83, Article 37 (August 2025), 15 pages. DOI: 10.1613/jair.1.16635

1 Introduction

Matching markets involve heterogeneous agents—typically from two distinct parties—who are paired to achieve mutual benefits. Over the past decade, these markets have rapidly evolved with the rise of the Internet, giving birth to a new paradigm known as Online Matching Markets (OMMs). Prominent examples include rideshare platforms (e.g., Uber and Lyft), crowdsourcing markets (e.g., Amazon Mechanical Turk, Cad Crowd, and crowdSPRING), and online advertising platforms (e.g., Google and Meta).

OMMs differ from traditional matching markets, where all agent information is disclosed in advance, in two key aspects: (**F1**) Agents from at least one side join the market dynamically and are referred to as online agents (e.g., riders, workers, and impressions); (**F2**) Upon the arrival of an online agent, a prompt and irrevocable matching decision must be made due to the limited "patience" of these agents.

OMMs have attracted considerable attention in both the computer science and operations research communities. A substantial body of work has been devoted to designing matching policies for ridesharing platforms [5, 25, 7, 42, 13] and crowdsourcing markets [6, 21, 12, 41, 36].

*Corresponding Author.

Authors' Contact Information: Evan Yifan Xu, ORCID: 0000-0001-9965-0086, xyf@seu.edu.cn, Southeast University, Nanjing, Jiangsu, China; Pan Xu, ORCID: 0000-0002-6362-6727, pxu@njit.edu, New Jersey Institute of Technology, Newark, NJ, USA.

DOI: 10.1613/jair.1.16635

Real-world OMMs often involve optimizing multiple objectives. A common goal is to generate a set of *relevant and diverse* offline agents for each arriving online agent. For example, in online recommendation systems (e.g., Netflix and Amazon), a well-studied problem is how to provide each user with a set of *relevant and diversified* recommendations; see, *e.g.,* [35, 34, 16, 31, 26]. In crowdsourcing markets, prior work [32, 3] has explored how to generate a set of *relevant and diversified* tasks for each online worker to enhance worker engagement. Similarly, in online team formation, studies [37, 29] have examined how to form a team of *relevant and diversified* workers for each arriving task to boost the team performance.

The applications mentioned above share several common features. First, the overarching goal is to optimize both relevance and diversity, even though these two objectives often conflict with each other. Second, capacity constraints are typically present on both sides of the matching process. For example, in online recommendation systems, both ads and users have associated capacities: each ad can be shown only a limited number of times due to budget constraints, and each user should be presented with a limited number of ads to avoid fatigue. Similar capacity constraints apply to tasks and workers in crowdsourcing markets.

Several common approaches have been proposed to optimize multiple objectives in Online Matching Markets (OMMs). One approach is to divide the online process into short, discrete time windows, transforming the problem into a fully offline setting where all online agent arrivals are known in advance. This allows us to bypass the challenges introduced by the two key features of OMMs—dynamic arrivals and irrevocable decisions (**F1** and **F2**); see, *e.g.,* [11]. Another approach involves defining a single meta-objective that balances multiple goals, thereby reducing the problem to a standard single-objective optimization. Examples of such meta-objectives include submodular functions[43, 33, 18] and linear or nonlinear combinations of multiple criteria [3]. In this setting, a common strategy is to design a greedy-based algorithm to optimize the meta-objective; see, *e.g.,* [2].

1.1 Main Model

In this paper, we propose a generic model of bi-objective maximization that can overcome shortcomings in previous approaches. Suppose we have a bipartite graph $G = (I, J, E)$, where I and J denote the sets of types of offline and online agents.[1] An edge $e = (i, j)$ indicates the feasibility of pairing agents i and j.[2] Let E_i (E_j) denote the set of edges incident to i (j). Each offline agent i has a capacity $c_i \in \mathbb{Z}^+$ in the way that i can be matched up to c_i times, where c_i reflects the limited budget on offline agent i (*e.g.,* tasks or ads). Each online agent j is associated with a collection of subsets of E_j, denoted by $\mathcal{S}_j \subseteq 2^{E_j}$, which consists of all feasible assignments with respect to j due to its capacity/patience. Let $\mathcal{S} = \cup_{j \in J} \mathcal{S}_j$, the collection of all feasible assignments. We are given two arbitrary metrics $\mathbf{d}$ and $\mathbf{w}$ on $\mathcal{S}$ such that for each $S \in \mathcal{S}_j$, $\mathbf{d}(S)$ and $\mathbf{w}(S)$ capture the respective utilities under the two metrics in the assignment of S to agent j.[3] Under our assumption, upon each arrival of agent j, the system should select an assignment $S \in \mathcal{S}_j$ for her, and if successful (*i.e.,* no violation of any budgets from offline agents), the system will gain utilities of $\mathbf{w}(S)$ and $\mathbf{d}(S)$ under the two metrics, respectively. Note that each assignment $S \in \mathcal{S}_j$ is a subset of *edges* incident to agent j, which encodes the information of both the online agent type j and a set of its neighboring offline agents. In this way, we allow for online-agent-dependent metrics such that $\mathbf{w}(S)$ can differ from $\mathbf{w}(S')$ even when S and S' include the same set of offline agents (but involve different online agents). All information of $G = (I, J, E)$, $\{c_i, \mathcal{S}_j\}$, and $\{\mathbf{w}(S), \mathbf{d}(S)\}$ is known as part of the input. Here are a few other features in our model. Throughout this paper, we denote $[n] = \{1, 2, \ldots, n\}$ for a generic positive integer n.

[1]We categorize a collection of analogous agents into clusters referred to as 'online(offline) agent types.' For instance, in spatial crowdsourcing scenarios, workers situated in the same geographic area are classified under one worker type.

[2]We refer to an agent of type j (i) directly as agent j (i).

[3]For example, in the context of online recommendations, $\mathbf{w}(S)$ and $\mathbf{d}(S)$ with $S \in \mathcal{S}_j$ can be interpreted as the overall relevance and diversity of ads in S with respect to online user (of type) j.

Arrivals of Online Agents. We have a given finite time horizon T, and during each time $t \in [T] \doteq \{1, 2, \ldots, T\}$, one single agent (of type) $\hat{j}$ will be sampled (*i.e.,* $\hat{j}$ arrives) with replacement such that $\Pr[\hat{j} = j] = r_j/T$ for all $j \in J$ with $\sum_j r_j/T = 1$. Here r_j is called the arrival rate of j. Note that the arrival distribution $\{r_j/T\}$ is known as input, and *independent* and *invariant* throughout the online T rounds. Our arrival setting is called Known Independent and Identical Distribution (KIID), which is mainly inspired from the fact that we can often learn the arrival distribution from historical logs [40]. *Throughout this paper, we assume $T \gg 1$. Part of our theoretical analyses and results are obtained under this assumption.*[4]

Bi-Objective Maximization. Consider a given online matching policy ALG and a given arrival sequence $\mathcal{A}$ of online agents. Let $\text{ALG}(\mathcal{A}) = \{S_j | j \in \mathcal{A}, S_j \in \mathcal{S}_j\}$ be a (random) allocation, where S_j is the assignment selected by ALG for j. Note that $\mathcal{A}$ can be a multiset of J due to potentially multiple arrivals of online agents. As a result, $\text{ALG}(\mathcal{A})$ can also be a multiset. The overall utilities achieved by ALG on metrics of $\mathbf{w}$ and $\mathbf{d}$ are defined as follows:

$$\textbf{Max-W} : \max \mathbb{E}\Big[\sum_{S \in \text{ALG}(\mathcal{A})} \mathbf{w}(S) \Big]; \tag{1}$$

$$\textbf{Max-D} : \max \mathbb{E}\Big[\sum_{S \in \text{ALG}(\mathcal{A})} \mathbf{d}(S) \Big], \tag{2}$$

where the expectation is taken over the randomness in $\mathcal{A}$ and that used by ALG.

1.2 Main Contriubtions

In this paper, we propose a generic online-matching-based model of bi-objective maximization existing in a wide range of OMMs. Our model features in the following ways. First, it can directly address challenges exclusively existing in OMMs, *i.e.,* the dynamic arrivals of online agents and the real-time decision-making requirements. Second, the model admits two arbitrary metrics defined on collecting all possible valid assignments to online agents, and it aims to maximize the two corresponding linear objectives simultaneously. We design a parameterized online algorithm, $\text{ATT}(\alpha, \beta)$, with α and β being two parameters, which can *smoothly* tradeoff the objectives of relevance and diversity.

Let $\Delta = \max_{S \in \mathcal{S}} |S|$, the largest cardinality of feasible assignments among all online agents. Note that Δ reflects the limited patience or budget/capacity among all individual online agents, and it typically takes a small value in practice; see, *e.g.,* the low tolerance of online users for (offline) ads and the limited capacity of arriving workers for (offline) tasks. In contrast, the budget/capacity on offline agents can be huge in reality. *Fortunately, the performance of the proposed Algorithm 1 does not depend on the capacity of offline agents.*

Theorem 1. *[Section 3.1] For any given $\alpha, \beta \geq 0$ with $\alpha + \beta \leq 1$, $\text{ATT}(\alpha, \beta)$ can achieve competitive ratios at least $\frac{(1-e^{-\Delta})}{\Delta} \cdot (\alpha, \beta) := \left(\frac{(1-e^{-\Delta})}{\Delta} \cdot \alpha, \frac{(1-e^{-\Delta})}{\Delta} \cdot \beta \right)$ simultaneously on any two pre-specified linear objectives, as defined in (1) and in (2).*

Theorem 2. *[Section 4] No algorithm can achieve a competitive ratio of $\frac{1-e^{-(\Delta-1+1/\Delta)}}{\Delta-1+1/\Delta} \cdot (\alpha, \beta)$ simultaneously on any two pre-specified linear objectives, as defined in (1) and in (2), with $\alpha + \beta > 1 + e^{-(\Delta-1+1/\Delta)}$ or $\alpha > 1$ or $\beta > 1$.*

Asymptotic Tightness of Lower Bound in Theorem 1. Our analysis establishes a strong lower bound on the achievable competitive ratio for the adaptive algorithm $\text{ATT}(\alpha, \beta)$. Theorem 1 shows that for any $\alpha, \beta \geq 0$ with $\alpha + \beta \leq 1$, $\text{ATT}(\alpha, \beta)$ achieves a competitive ratio of at least $\frac{(1-e^{-\Delta})}{\Delta} \cdot (\alpha, \beta)$ simultaneously on two predetermined linear objectives, as defined in (1) and in (2). This constructive guarantee is complemented by Theorem 2, which

[4] This is a common assumption among studies of online-matching under KIID; see, *e.g.,* [22, 19, 28, 9].

establishes fundamental limits: no algorithm can achieve a competitive ratio of $\frac{1-e^{-(\Delta-1+1/\Delta)}}{\Delta-1+1/\Delta} \cdot (\alpha, \beta)$ if the condition $\alpha + \beta > 1 + e^{-(\Delta-1+1/\Delta)}$ (or $\alpha > 1$ or $\beta > 1$) holds. This impossibility result implies that any algorithm achieving this competitive ratio must satisfy $\alpha + \beta \leq 1 + e^{-(\Delta-1+1/\Delta)}$. Notably, the critical competitive ratio factor from Theorem 2 is $\frac{1-e^{-(\Delta-1+1/\Delta)}}{\Delta-1+1/\Delta}$. As Δ becomes large, the term $1/\Delta$ diminishes, causing $\Delta - 1 + 1/\Delta$ to approach $\Delta - 1$. Consequently, this factor asymptotically approaches $\frac{1-e^{-(\Delta-1)}}{\Delta-1}$. Comparing this with our algorithm's guaranteed performance factor of $\frac{1-e^{-\Delta}}{\Delta}$ from Theorem 1, we observe that these two expressions are very close for large Δ. This proximity suggests that our lower-bound result in Theorem 1 is **asymptotically tight**. The conditions on $\alpha + \beta$ further support this: our algorithm $\text{ATT}(\alpha, \beta)$ achieves its bound for $\alpha + \beta \leq 1$, while the limit implied by Theorem 2 is $\alpha + \beta \leq 1 + e^{-(\Delta-1+1/\Delta)}$, indicating the gap is characterized by the small term $e^{-(\Delta-1+1/\Delta)}$. This indicates that the $\text{ATT}(\alpha, \beta)$ algorithm is nearly optimal for large Δ within the permissible range of α and β.

1.3 Related Work

Since the influential work of [23], online matching and related models have attracted substantial interest over the past thirty years, partly driven by the rise of Online Matching Markets (OMMs). Our model falls within the paradigm of multi-objective online optimization, which has recently been employed to address trade-offs between conflicting objectives in OMMs. For example, [15] studied a variant of the online submodular welfare maximization problem where each offline agent is associated with two distinct submodular utility functions. The works of [24, 30, 38] proposed bi-objective models aimed at balancing fairness and profit in the context of ridesharing. The model of [39] considers a specific application scenario in online matching markets with two contextual objectives. These models differ significantly from the one and setting described in this paper. Another approach to handling trade-offs between conflicting objectives is to formulate the problem as the maximization or minimization of a single submodular function; see, e.g., [14, 2].

Table 1. Glossary of Notation

I (J)	Set of types of offline (online) agents
c_i	Capacity on offline agent i
E_i (E_j)	Set of edges incident to i (j)
$\mathcal{S}_j$	Collection of all feasible assignments $S \subseteq E_j$ to j
$\mathcal{S}$	$\cup_{j \in J} \mathcal{S}_j$, collection of all feasible assignments
Δ	$\max_{S \in \mathcal{S}} \lvert S \rvert$, largest cardinality of all feasible assignments
$\mathbf{w}(S)$	Utility associated with S under metric $\mathbf{w}$
$\mathbf{d}(S)$	Utility associated with S under metric $\mathbf{d}$
T	Total number of online rounds
r_j	Arrival rate (or expected number of arrivals) of j

2 Preliminaires and Main Techniques

Competitive Ratio. The competitive ratio is a commonly-used metric to evaluate the performance of online algorithms. Consider an online maximization problem with one single objective. Let $\text{ALG}(\mathcal{I}) = \mathbb{E}_{\mathcal{A} \sim \mathcal{I}}[\text{ALG}(\mathcal{A})]$ denote the expected performance of ALG on an instance $\mathcal{I}$, where the expectation is taken over the randomness in both the arrival sequence $\mathcal{A}$ and that in ALG. Let $\text{OPT}(\mathcal{I}) = \mathbb{E}_{\mathcal{A} \sim \mathcal{I}}[\text{OPT}(\mathcal{A})]$ denote the expected performance of a *clairvoyant optimal* on the instance $\mathcal{I}$, where $\text{OPT}(\mathcal{A})$ refers to the optimal performance *after* observing the full arrival sequence $\mathcal{A}$. We say ALG achieves an online competitive ratio of $\rho \in [0, 1]$ if $\text{ALG}(\mathcal{I}) \geq \rho \, \text{OPT}(\mathcal{I})$ for all possible inputs $\mathcal{I}$. The example below shows that the natural Greedy can have an arbitrarily bad performance on some metrics. Generally, Greedy will assign every arriving agent j with a then available assignment $S \in \mathcal{S}_j$ that has the largest utility under a certain metric.

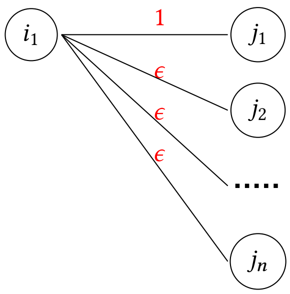

$T = n$
$c_{i_1} = 1$
$c_j = 1, \forall j \in J$
$\mathcal{S}_j = \{\{(i_1, j)\}\}, \forall j \in J$
$\mathbf{w}(\{(i_1, j_1)\}) = 1$
$\mathbf{w}(\{(i_1, j)\}) = \epsilon, \forall j \neq j_1$
Greedy $\leq \epsilon + 1/n$
OPT-W $\geq 1 - 1/\mathrm{e}$

Fig. 1. A toy example on which the natural Greedy achieves a competitive ratio of zero.

Example 1 (Greedy Achieves a Competitive Ratio of Zero)**.** Consider a toy example as shown in Figure 1. We have a star graph with $I = \{i_1\}$ and $J = \{j_1, \ldots, j_n\}$, where $T = n$ and $r_j = 1$ for every $j \in J$. In other words, during each round $t \in [T]$, a single online agent $\hat{j}$ will be sampled uniformly at random with replacement such that $\Pr[\hat{j} = j_\ell] = 1/n$ for every $\ell \in [n]$. Let $c_{i_1} = 1$ and $\mathcal{S}_j = \{\{(i_1, j)\}\}$ for all $j \in J$. That is, the offline agent i_1 can be matched only once, and each online agent $j \in J$ accepts an assignment consisting of at most one offline agent upon its arrival. Consider a metric $\mathbf{w}$ where $\mathbf{w}((i_1, j_1)) = 1$ and $\mathbf{w}((i_1, j)) = \epsilon > 0$ for all $j \in J \setminus \{j_1\}$.

Under the metric $\mathbf{w}$, we can verify that:

(1) Greedy achieves an expected utility of $1/n \cdot 1 + (1 - 1/n) \cdot \epsilon \leq 1/n + \epsilon$.
(2) A clairvoyant optimal OPT-W achieves an expected utility of $(1 - 1/\mathrm{e}) \cdot 1 + 1/\mathrm{e} \cdot \epsilon$. This is because OPT-W will assign i_1 to j_1 if j_1 arrives at least once, which happens with probability $1 - 1/\mathrm{e}$.[5]

By definition, Greedy achieves a competitive ratio no greater than $(1/n + \epsilon)/(1 - 1/\mathrm{e}) \approx 0$ under metric $\mathbf{w}$. ∎

Time-adaptive Simulation-based Attenuations. For online optimization problems, the arrival setting of known distributions like here makes simulation-based attenuation plausible since we can simulate both the internal randomness in the algorithm itself and the external randomness from online agents' arrivals. A typical framework is as follows. Suppose we aim for a target competitive ratio of $\tau \in (0, 1)$ and a collection of assignment $\mathcal{S}$. For each $S \in \mathcal{S}$, let x^*_S is the expected number of times that S is made in a clairvoyant optimal (that can be obtained through solving a benchmark LP). Consider an agent of type j that arrives at time t and let $\mathcal{S}_j$ be the set of all assignments relevant to j. By applying Monte-Carlo simulations, we can get a sharp estimate of $\gamma_{S,t} := \Pr[\mathsf{SF}_{S,t}]$ for every $S \in \mathcal{S}_j$, which denotes the probability that assignment S is safe to make with respect to external constraints (*e.g.,* budgets constraints on all offline agents involved in S). Then the sampling distribution over $\mathcal{S}_j$ is carefully designed to include a factor of $\tau/\gamma_{S,t}$ such that every assignment $S \in \mathcal{S}_j$ will be selected with probability *equal to* $\tau \cdot x^*_S/T$ during every round $t \in [T]$, and thus, the total expected number of selections of S is equal to $\tau \cdot x^*_S$. In this way, every assignment S is selected with a fraction of τ times of the clairvoyant optimal, which immediately leads to a competitive ratio of τ. The framework above is called non-adaptive in the way that every assignment S will be selected with a uniform probability $(\tau \cdot x^*_S/T)$, regardless of its arrival time. Almost all

[5]The probability that j_1 never arrives during the $T = n$ rounds is $(1 - r_{j_1}/T)^T = (1 - 1/T)^T$, which approximates $1/\mathrm{e}$ as $T \to \infty$. This means it arrives at least once with a probability of $1 - 1/\mathrm{e}$.

current existing attenuations are non-adaptive; see, *e.g.,* stochastic knapsack [27], stochastic matching [1, 8], and (online) matching policy design in rideshare [39, 17, 13].

In this paper, we propose a framework of time-adaptive attenuations. Specifically, we carefully craft a time-dependent sampling distribution for every arriving online agent j at time t such that every assignment $S \in \mathcal{S}_j$ will be selected with probability equal to $\tau \cdot x^*_S \cdot (1 - \Delta/T)^{t-1}/T$ at time t; see more details in Algorithm 1.

3 Main Algorithm

Throughout this paper, we use OPT-W and OPT-D to denote the clairvoyant optimal algorithms (and the corresponding performance) under metrics $\mathbf{w}$ and $\mathbf{d}$, respectively. Recall that $\mathcal{S} = \cup_{j \in J} \mathcal{S}_j$, the collection of all assignments for online agents. For each agent i, let $\mathcal{S}_i = \{S | \exists j \in J : S \in \mathcal{S}_j \text{ and } (i, j) \in S\}$, the collection of assignments involving i. For each $S \in \mathcal{S}$, let x_S be the expected number of times that S is selected in the clairvoyant optimal. Our benchmark LPs can be stated as follows.

$$\max \sum_{S \in \mathcal{S}} \mathbf{w}(S) \cdot x_S \tag{3}$$

$$\max \sum_{S \in \mathcal{S}} \mathbf{d}(S) \cdot x_S \tag{4}$$

$$\sum_{S \in \mathcal{S}_j} x_S \leq r_j \quad \forall j \in J \tag{5}$$

$$\sum_{S \in \mathcal{S}_i} x_S \leq c_i \quad \forall i \in I \tag{6}$$

$$0 \leq x_S \qquad \forall S \in \mathcal{S}. \tag{7}$$

Throughout this paper, we use LP (3) and LP (4) to denote the two LPs of Objective (3) and (4) each with Constraints (5) to (7).

Lemma 1. *The optimal values of* LP (3) *and* (4) *are valid upper bounds for the expected performance of* OPT-W *and* OPT-D*, respectively.*

Proof. By linearity of expectation, we can verify that the two objectives (3) and (4) correctly code **Max-W** and **Max-D**, respectively. It will suffice to show the validities of all constraints. Consider a given arrival sequence $\mathcal{A} = (A_j | j \in J)$, where A_j is the number of arrivals of agent j. For each $S \in \mathcal{S}$, let $X^{\mathcal{A}}_S$ be the number of times that S is selected in an offline optimal after observing $\mathcal{A}$. First, we have $\sum_{S \in \mathcal{S}_j} X^{\mathcal{A}}_S \leq A_j$ since we can select at most one assignment from $\mathcal{S}_j$ upon each arrival of j. Taking expectation over $\mathcal{A}$ on both side, we get Constraint (5). Second, $\sum_{S \in \mathcal{S}_i} X^{\mathcal{A}}_S \leq c_i$ since i can be included in at most c_i different assignments for every arrival sequence. Taking expectation over $\mathcal{A}$, we get Constraint (6). The last one is trivial. Thus, we are done. □

3.1 An Algorithm with Time-adaptive Attenuations

Let $\alpha \in [0, 1]$ and $\beta \in [0, 1]$ be two parameters with $\alpha + \beta \leq 1$. Our first algorithm, ATT(α, β), is formally stated in Algorithm 1.

We now justify Step (9) in ATT(α, β). Observe that

$$\sum_{S \in \mathcal{S}_j} (\alpha x^*_S + \beta y^*_S)/r_j = \alpha \sum_{S \in \mathcal{S}_j} x^*_S/r_j + \beta \sum_{S \in \mathcal{S}_j} x^*_S/r_j \leq \alpha + \beta \leq 1.$$

The first inequality above is due to Constraint (5) in benchmark LPs and the second inequality follows from our assumption. To validate Step (9) in ATT(α, β), it suffices to show that $(1 - \Delta/T)^{t-1} \leq \gamma_{S,t}$ for all S and t, as presented in the lemma below.

Algorithm 1: An Algorithm with Time-adaptive Attenuations ATT(α, β), with $0 \leq \alpha, \beta \leq 1$ and $\alpha+\beta \leq 1$.

1 **Offline Phase**:

/* The offline phase will output $\{\gamma_{S,t} | S \in \mathcal{S}, t \in [T]\}$, where $\gamma_{S,t}$ denotes the probability that every offline agent involved in S has at least one capacity at (the beginning of) t (**we call S is safe at t**). */

2 Solve LP (3) and LP (4), and let $\{x_S^*\}$ and $\{y_S^*\}$ be the respective optimal solutions.

3 *Initialization*: When $t = 1$, set $\gamma_{S,t} = 1$ for all $S \in \mathcal{S}$.

4 **for** $t = 2, 3, \ldots, T$ **do**

5 Applying the Monte-Carlo method to simulate Step 8 to Step 10 for all the rounds $t' = 1, 2, \ldots, t-1$ of Online Phase, we can get a sharp estimate of the probability that each assignment $S \in \mathcal{S}$ is safe at t, denoted by $\gamma_{S,t}$.

6 **Online Phase**:

7 **for** $t = 1, 2, \ldots, T$ **do**

8 Let an online agent (of type) j arrive at time t.

9 Sample an assignment $S \in \mathcal{S}_j$ with probability $\big(\alpha x_S^* + \beta y_S^*\big)/r_j \cdot \big((1-\Delta/T)^{t-1}/\gamma_{S,t}\big)$.

10 Select S if S is safe and reject j otherwise.

Lemma 2. $(1-\Delta/T)^{t-1} \leq \gamma_{S,t}$ *for all* $S \in \mathcal{S}$ *and* $t \in [T]$.

Proof. We prove by induction over $t = 1, \ldots, T$. When $t = 1$, we have $\gamma_{S,t} = 1$. Thus, we are done. Consider a given time $t' \geq 2$ and $S' \in \mathcal{S}$. Let $I(S')$ be the set of offline agents involved in S'. Recall that $\gamma_{S',t'}$ is the probability that every $i \in I(S')$ has at least one capacity at t'. Consider a given $i \in I(S')$. Let $\mathcal{N}_i$ be the set of neighbors of i in the graph G. For each $j \in \mathcal{N}_i$, let $\mathcal{S}_{ij} = \{S \in \mathcal{S}_j | (i,j) \in S\}$, the collection of assignments including edge (ij). Let $U_{i,t'}$ be the (random) number of times that i gets matched during the previous $t'-1$ rounds. For each $t < t'$, $j \in \mathcal{N}_i$, and $S \in \mathcal{S}_{ij}$, let $Y_{j,t} = 1$ indicate that j arrives at t, $Z_{S,t} = 1$ indicate if S is sampled at t, and $\mathsf{SF}_{S,t} = 1$ indicate that S is safe at t. By definition, we have $\mathbb{E}[Y_{j,t}] = r_j/T$, $\mathbb{E}[Z_{S,t}] = \big(\alpha x_S^* + \beta y_S^*\big)/r_j \cdot \big((1-\Delta/T)^{t-1}/\gamma_{S,t}\big) := \eta_{S,t}$, and $\mathbb{E}[\mathsf{SF}_{S,t}] = \gamma_{S,t}$. Thus, we have

$$\mathbb{E}[U_{i,t}] = \mathbb{E}\Big[\sum_{t<t'}\sum_{j\in\mathcal{N}_i}\sum_{S\in\mathcal{S}_{ij}} Y_{j,t}\cdot Z_{S,t}\cdot \mathsf{SF}_{S,t}\Big] \tag{8}$$

$$= \sum_{t<t'}\sum_{j\in\mathcal{N}_i}\sum_{S\in\mathcal{S}_{ij}} (r_j/T)\cdot \eta_{S,t}\cdot\gamma_{S,t} \tag{9}$$

$$= \sum_{t<t'}\sum_{j\in\mathcal{N}_i}\sum_{S\in\mathcal{S}_{ij}} \big(\alpha x_S^* + \beta y_S^*\big)\cdot(1-\Delta/T)^{t-1}/T \tag{10}$$

$$= \sum_{t<t'}\sum_{S\in\mathcal{S}_i} \big(\alpha x_S^* + \beta y_S^*\big)\cdot(1-\Delta/T)^{t-1}/T \tag{11}$$

$$\leq c_i \cdot \sum_{t<t'}(1-\Delta/T)^{t-1}/T \tag{12}$$

$$= c_i \cdot \Big(1-\big(1-\Delta/T\big)^{t'-1}\Big)/\Delta. \tag{13}$$

Equality (11) is due to the fact $\{\mathcal{S}_{ij} | j \in \mathcal{N}_i\}$ are disjoint and $\cup_{j\in\mathcal{N}_i}\mathcal{S}_{ij} = \mathcal{S}_i$. Inequality (12) follows from Constraint (6) suggesting $\sum_{S\in\mathcal{S}_i} x_S^* \leq c_i$ and $\sum_{S\in\mathcal{S}_i} y_S^* \leq c_i$, and the fact $\alpha+\beta \leq 1$. Now we analyze the probability that S' is safe at t'.

$$\gamma_{S',t'} = \Pr\Big[\bigwedge_{i \in I(S')} \big(U_{i,t'} \leq c_i - 1\big)\Big] \tag{14}$$

$$\geq 1 - \sum_{i \in I(S')} \Pr[U_{i,t} \geq c_i] \tag{15}$$

$$\geq 1 - \sum_{i \in I(S')} \frac{E[U_{i,t}]}{c_i} \geq \Big(1 - \frac{\Delta}{T}\Big)^{t'-1}. \tag{16}$$

Note that Inequality (15) is due to the union bound; the first inequality on (16) is due to Markov's inequality; the second one on (16) is due to Inequality (13) and the fact $|I(S')| \leq |S'| \leq \Delta$. □

Here we present the proof of Theorem 1.

Proof of Theorem 1. Consider a given j and an assignment $S \in \mathcal{S}_j$. Let X_S be the random number of times that S is selected by $\mathrm{ATT}(\alpha, \beta)$. For each $t \in [T]$, let $Y_t = 1$ indicate that j arrives at t and $Z_{S,t} = 1$ indicate that S is sampled at t. Thus,

$$\begin{aligned}
\mathbb{E}[X_S] &= \sum_{t=1}^{T} \mathbb{E}[Y_t \cdot Z_{S,t} \cdot \mathsf{SF}_{S,t}] \\
&= \sum_{t=1}^{T} (r_j/T) \cdot \big(\alpha x_S^* + \beta y_S^*\big)/r_j \cdot \big((1 - \Delta/T)^{t-1}/\gamma_{S,t}\big) \cdot \gamma_{S,t} \\
&= \big(\alpha x_S^* + \beta y_S^*\big) \sum_{t=1}^{T} (1 - \Delta/T)^{t-1}/T \\
&= \big(\alpha x_S^* + \beta y_S^*\big) \cdot (1 - \mathrm{e}^{-\Delta})/\Delta. \text{ (By taking } T \to \infty)
\end{aligned}$$

Thus, we have $\mathbb{E}[X_S] = \big(\alpha x_S^* + \beta y_S^*\big)(1 - \mathrm{e}^{-\Delta})/\Delta$. By linearity of expectation, we claim that the expected performance of ATT under the metric $\mathbf{w}$, denoted by $\mathrm{ATT}_{\mathbf{w}}$, satisfies that

$$\begin{aligned}
\mathrm{ATT}_{\mathbf{w}} &= \sum_{S \in \mathcal{S}} \mathbf{w}(S) \cdot \mathbb{E}[X_S] \geq \sum_{S \in \mathcal{S}} \mathbf{w}(S) \cdot (\alpha x_S^*) \cdot \frac{1 - \mathrm{e}^{-\Delta}}{\Delta} \\
&= \alpha \cdot \mathrm{LP}\,(3) \cdot (1 - \mathrm{e}^{-\Delta})/\Delta \geq \text{OPT-W} \cdot \alpha \cdot (1 - \mathrm{e}^{-\Delta})/\Delta.
\end{aligned}$$

The last inequality above is due to Lemma 1. The analysis above suggests that $\mathrm{ATT}(\alpha, \beta)$ achieves a competitive ratio at least $\alpha \cdot (1 - \mathrm{e}^{-\Delta})/\Delta$ for the metric $\mathbf{w}$. We can argue in the same way that $\mathrm{ATT}(\alpha, \beta)$ achieves a competitive ratio at least $\beta \cdot (1 - \mathrm{e}^{-\Delta})/\Delta$ for the metric $\mathbf{d}$. □

4 Proof of Theorem 2

Example 2. Consider the projective plane $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ of order $d - 1$ [10], where $d - 1$ is a prime number. $\mathcal{H}$ is a hypergraph such that: (1) $\mathcal{H}$ is d-uniform, d-regular, and intersecting; (2) $|\mathcal{V}| = |\mathcal{E}| = d^2 - d + 1$; and (3) the natural canonical LP on $\mathcal{H}$ has an optimal value of $d - 1 + 1/d$. Now, based on $\mathcal{H}$, we construct an instance as follows.

For ease of notation, we use $v \in \mathcal{V}$ and $f \in \mathcal{E}$ to denote a hyper-vertex and a hyper-edge, respectively. For each $v \in \mathcal{V}$, we create an offline agent type i_v. Similarly, for each $f \in \mathcal{E}$, we create two online agent types, j_f and

j'_f, such that j_f and j'_f each admit a single valid assignment. Specifically, $\mathcal{S}_{j_f} = \{S_{j_f}\}$ with $S_{j_f} = \{(i_v, j_f) : v \in f\}$, and $\mathcal{S}_{j'_f} = \{S_{j'_f}\}$ with $S_{j'_f} = \{(i_v, j'_f) : v \in f\}$.[6]

Let $I := \{i_v : v \in \mathcal{V}\}$ and $J = \{j_f : f \in \mathcal{E}\} \cup \{j'_f : f \in \mathcal{E}\} := J_1 \sqcup J_2$, such that $|I| = |J_1| = |J_2| = d^2 - d + 1$. For each online agent $j \in J$, set $r_j = 1/d$. Thus, the total sum of arrival rates among all online agents in J is $R := \sum_{j \in J} r_j = 2(d - 1 + 1/d)$. We assume $R \ll T$.[7] Observe that $\Delta = \max_{j \in J} \max_{S \in \mathcal{S}_j} |S| = d$. Set $c_i = 1$ for all $i \in I$. Consider a uniform metric for $\mathbf{w}$ and $\mathbf{d}$ such that for each pair of online agents, (j_f, j'_f), and their associated unique valid assignments, $(S_{j_f}, S_{j'_f})$, we have $\mathbf{w}(S_{j_f}) = 1$ and $\mathbf{d}(S_{j_f}) = 0$, and $\mathbf{w}(S_{j'_f}) = 0$ and $\mathbf{d}(S_{j'_f}) = 1$. ■

Lemma 3. *LP-(3) and LP-(4) each have an optimal value of at least $d - 1 + 1/d$ for Example 2.*

Proof. First, focus on LP-(3) with the metric $\mathbf{w}$. We can ignore online agents in J_2 and their associated assignments, as their values under the metric $\mathbf{w}$ are zero by definition (see Example 2).

Let $\mathcal{S}_1 = \bigcup_{j \in J_1} \mathcal{S}_j$ be the collection of all valid assignments for online agents in J_1. Consider the solution $x_S = 1/d$ for all $S \in \mathcal{S}_1$. This solution $\{x_S\}_{S \in \mathcal{S}_1}$ can be verified to be feasible for LP-(3) with an objective value equal to $d - 1 + 1/d$. This implies that the optimal value of LP-(3) is at least $d - 1 + 1/d$.

Similarly, for LP-(4) with the metric $\mathbf{d}$, we focus on agents in J_2. Let $\mathcal{S}_2 = \bigcup_{j \in J_2} \mathcal{S}_j$. The solution $x_S = 1/d$ for all $S \in \mathcal{S}_2$ can be verified as feasible for LP-(4) and yields an objective value of $d - 1 + 1/d$. Thus, the optimal value of LP-(4) is also at least $d - 1 + 1/d$. □

Proof of Theorem 2. Consider Example 2 and any given online algorithm ALG. Since the hypergraph $\mathcal{H}$ is intersecting, each hyper-vertex $v \in \mathcal{V}$ appears in every hyper-edge $f \in \mathcal{E}$. Since each offline agent i_v has a unit matching capacity, this means ALG can make at most one assignment given any arrival of an online agent. Let $R := \sum_{j \in J} r_j = 2(d - 1 + 1/d)$ be the sum of arrival rates among all online agents in J. Note that the probability that there is at least one arrival of an online agent is equal to

$$1 - \left(1 - \frac{R}{T}\right)^T \approx 1 - \mathrm{e}^{-R} = 1 - \mathrm{e}^{-2(d-1+1/d)},$$

as $T \to \infty$. Note that for any assignment S, we have $\mathbf{w}(S) + \mathbf{d}(S) = 1$ by definition. This implies that the sum of the expected utilities achieved by ALG under metrics $\mathbf{w}$ and $\mathbf{d}$ is no more than $1 - \mathrm{e}^{-2(d-1+1/d)}$.

Suppose ALG achieves a competitiveness of $\frac{1-\mathrm{e}^{-(\Delta-1+1/\Delta)}}{\Delta-1+1/\Delta} \cdot (\alpha, \beta) = \frac{1-\mathrm{e}^{-(d-1+1/d)}}{d-1+1/d} \cdot (\alpha, \beta)$, for some $\alpha, \beta \geq 0$ simultaneously on the two pre-specified linear objectives **Max-W** and **Max-D** (corresponding to LP (3) and LP (4)). By Lemma 3, each benchmark LP has an optimal value of at least $d - 1 + 1/d$. Thus, the expected utilities achieved by ALG under metrics $\mathbf{w}$ and $\mathbf{d}$ are at least $\left(1 - \mathrm{e}^{-(d-1+1/d)}\right) \cdot \alpha$ and $\left(1 - \mathrm{e}^{-(d-1+1/d)}\right) \cdot \beta$, respectively. Since the sum of the expected utilities achieved by ALG under metrics $\mathbf{w}$ and $\mathbf{d}$ is no more than $1 - \mathrm{e}^{-2(d-1+1/d)}$, we have:

$$\alpha + \beta \leq \frac{1 - \mathrm{e}^{-2(d-1+1/d)}}{1 - \mathrm{e}^{-(d-1+1/d)}} = 1 + \mathrm{e}^{-(d-1+1/d)} = 1 + \mathrm{e}^{-(\Delta-1+1/\Delta)}.$$

Thus, we have established the first part of our claim (referring to the condition on $\alpha + \beta$ in Theorem 2). This implies that for any algorithm to achieve the stated competitive ratio factor with performance parameters (α, β), the condition $\alpha + \beta \leq 1 + \mathrm{e}^{-(\Delta-1+1/\Delta)}$ must hold.

Now, we address the second part of the claim in Theorem 2, concerning the conditions on individual α and β. Consider the objective **Max-W** defined in (1). For this objective, we focus on online agents in J_1 and their

[6] Observe that for the two online agents j_f and j'_f, their unique admissible assignments S_{j_f} and $S_{j'_f}$ involve the same set of offline agents, which is $\{i_v : v \in f\}$, but involve different online agents (j_f and j'_f, respectively). Thus, we still view the two assignments as distinct.

[7] This implies that in each round, the probability that some online agent will arrive is $R/T \ll 1$. We can create a dummy online agent with no valid assignments that arrives with probability $1 - R/T$.

Table 2. Parameter settings, where the default settings are marked as bold.

Parameters	Description	Setting	
		AMT Data	MovieLens 1M
$\lvert I \rvert$	Numbers of offline agent types	20	
$\lvert J \rvert$	Numbers of offline agent types	20	
Ψ	Uniform budget for all offline agent types	25	10
Δ	Uniform capacity for all online agent types	{**2**,4}	
Distance Function	Similarity functions to calculate the relevance and diversity	{**Jaccard**, Euclidean}	
α	Weight for relevance maximization	{0,0.1,0.2,. . .,1}	
β	Weight for diversity maximization	{1,0.9,0.8,. . .,0}	

admissible assignments, since the utilities for **Max-W** from assignments involving agents in J_2 are zero by definition. For any online algorithm ALG, its expected utility for objective **Max-W** is no more than the probability of at least one arrival from J_1 (since each such arrival can yield at most unit utility for **Max-W** and ALG can make no more than one assignment). This probability is:

$$1-\left(1-\frac{R_1}{T}\right)^T \approx 1-\mathrm{e}^{-R_1}=1-\mathrm{e}^{-(d-1+1/d)},$$

where $R_1 = R/2$ is the sum of arrival rates for agents in J_1. The optimal value for LP (3) (representing **Max-W**) is at least $d-1+1/d = \Delta-1+1/\Delta$, as per Lemma 3. This implies that any algorithm can achieve a competitiveness no more than

$$\frac{1-\mathrm{e}^{-(d-1+1/d)}}{d-1+1/d}=\frac{1-\mathrm{e}^{-(\Delta-1+1/\Delta)}}{\Delta-1+1/\Delta},$$

if using LP (3) as the benchmark. This establishes any feasible α must satisfy $\alpha \leq 1$. Similarly, by considering objective **Max-D** and agents in J_2, we can show that $\beta \leq 1$. This completes the proof of Theorem 2. □

5 Experiments

We test our algorithm ATT on two real datasets, namely, an Amazon Mechanical Turk (AMT) dataset collected from crowdsourcing markets (assigning static tasks to dynamic workers) and a MovieLens dataset collected from online recommendations (recommending movies to dynamic users).

AMT Dataset Preprocessing. We first use a crawled dataset [32] from AMT, which includes a total of 152, 221 task groups. For each task group, we have task ID, title, reward, description, requester name, and keywords. According to the predefined 44 keywords (*e.g.,* classification, sentiment analysis and extract information), all task groups can be divided into 26 task types. Similar to [32], we generate online workers synthetically. For each worker, we randomly sample 3 to 7 keywords to capture his/her interest topics.

We construct our graph $G = (I, J, E)$ as follows. Each $i \in I$ denotes a crowdsourcing task type that is represented by a 44-dimensional binary vector capturing all keywords it covers. Each $j \in J$ describes an online worker type with a 44-dimensional binary vector representing the worker's interest topics. An edge $e = (i, j)$ indicates that task type i and worker type j share at least one similar keyword. We downsample from all task types such that $|I| = 20$, and generate 20 worker types such that $|J| = 20$. For each i, we set a uniform budget $\Psi = 25$; for each j, we set a uniform capacity Δ, which is selected from $\{2, 4\}$. We generate a uniform random value $r_j \in [0, 1]$ for each j such that $\sum_{j\in J} r_j = T$. In our experiments, we set $T = 150$.

MovieLens Dataset Preprocessing. We further conduct experiments on a movie rating dataset: MovieLens 1M dataset [20] which includes $1,000,209$ ratings of about $3,900$ movies made by $6,040$ MovieLens users. For each movie we have a unique movie ID, a title including the year of release and genres selected from a set of 18 predefined genres (*e.g.,* comedy, drama, and documentary). For each user, we have an unique user ID and demographic information (*e.g.,* gender, age, and occupation). Here, we randomly sample 3 to 7 genres for each user to represent his/her favorite movie genres. Additionally, for each rating record we have the corresponding user ID and movie ID, the rating score from 1 to 5 and the time when the rating is made.

To set up the proposed online-matching model, we assume that movies are tasks to be rated, and users are online workers who join the system dynamically. Our goal is to assign to every user a set of relevant and diverse movies. For each movie type i and user type j, we create a 18-dimensional binary vector such that it captures the genres covered by i and j. We downsample from all movie types and user types such that $|I| = 20$ and $|J| = 20$. For each i, we set a uniform budget $\Psi = 10$; for each j, we set a uniform capacity $\Delta \in \{2, 4\}$. To get the arrival distribution of online users, we first summarize the rating history for each user and set the number of ratings as their original arrival rate r'_j. We adjust the original rate r'_j by multiplying $T/\sum_{j\in J} r'_j$ such that $\sum_{j\in J} r_j = T$, where we set $T = 150$.

Relevance and Diversity. Recall that $\mathcal{S}_j$ denote the collection of all valid assignments for j and for each $S \in \mathcal{S}_j$, $\mathbf{w}(S)$ and $\mathbf{d}(S)$ represent the respective relevance and diversity of S with respect to j. We use pairwise distance [43, 4] to quantify $\mathbf{w}$ and $\mathbf{d}$: for a given set of S for worker j, we have $\mathbf{w}(S) = \sum_{i\in S} \big(1 - d(i, j)\big)$ and $\mathbf{d}(S) = \sum_{i_m, i_n \in S, m>n} d(i_m, i_n)$, where $d(\cdot)$ refers to the distance between task and worker or two tasks. In the experiment, we consider two distance functions, namely, jaccard [32] and euclidean [35].

Algorithms. In addition to ATT, we have implemented two heuristics, namely, ATT-Boosting (ATT-B) and Greedy, which are stated as follows. Let $\{x^*_S\}$ and $\{y^*_S\}$ be optimal solutions to LP (3) and LP (4), respectively. Consider a worker (of type) j arrives at t in the online phase, and let $\mathcal{S}_{j,t} \subseteq \mathcal{S}_j$ be the set of all *safe* assignments with respect to j at time t. Recall that an assignment $S \in \mathcal{S}_j$ is called safe at t iff every offline agent involved in S has at least one budget at (the beginning of) t. (i) ATT-B(α, β): it can be viewed as a boosted version of ATT(α, β), which will sample an assignment $S' \in \mathcal{S}_{j,t}$ with probability $\big(\alpha x^*_{S'}/\sum_{S\in\mathcal{S}_{j,t}} x^*_S + \beta y^*_{S'}/\sum_{S\in\mathcal{S}_{j,t}} y^*_S\big)$ if $\mathcal{S}_{j,t}$ is not empty. (ii) Greedy: it will with probability α select the assignment maximizing the relevance among all assignments in $\mathcal{S}_{j,t}$ (under the metric of $\mathbf{w}$), and with probability β select the assignment maximizing diversity among all assignments in $\mathcal{S}_{j,t}$ (under the metric of $\mathbf{d}$). Break ties in an arbitrary way if any.

We run the above three algorithms for 500 times and take the average as the final performance for each instance. For each algorithm, we compute these two ratios: the ratio of relevance (**Max-W**) to the optimal value of LP (3) and that of diversity (**Max-D**) to the optimal value of LP (4). We take these two ratios as the final competitive ratios (CR) achieved by the algorithm on relevance and diversity, respectively.

Results and Discussions. Figure 2 and Figure 4 show the results of ATT on the two respective datasets and the related lower bounds in Theorem 1. We observe that competitive ratios (CR) of ATT on both relevance and diversity always stay above the theoretical lower bounds, as stated in Theorem 1.

Furthermore, in the AMT scenario, the competitive ratios achieved by ATT are higher than the MovieLens ones. This is due to that we set a higher Ψ in the AMT scenario. Note that a choice of a higher value of α implies that ATT will prioritize relevance, which will help ATT achieve a higher CR on relevance. Similarly, a choice of a higher value of β suggests that ATT will give more preference to diversity; thus, it will lead to a higher CR on diversity. This trend is more prominent on the MovieLens dataset, as shown in Figure 4.

Figure 3 and Figure 5 show the relative effectiveness and flexibility of ATT compared against ATT-B and Greedy over different choices of distance functions, α and β. We see that ATT-B is the clear winner in almost all cases. This is expected since ATT-B takes advantage of the guidance of LP solutions similar to ATT, and shares

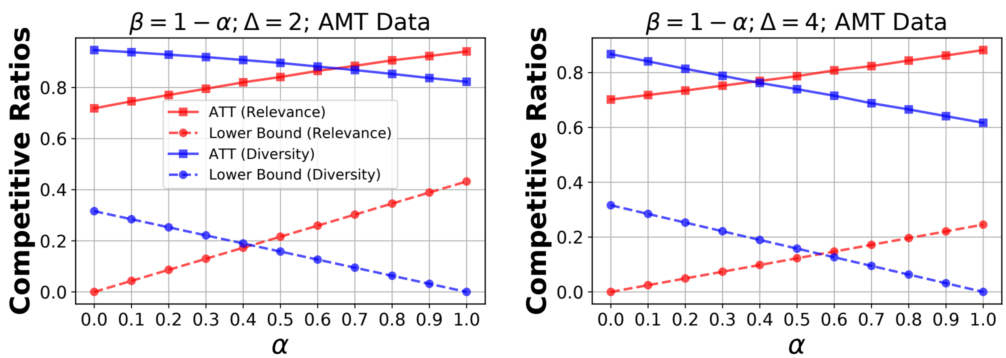

Fig. 2. AMT Dataset: competitive ratios (CR) of relevance and diversity with $\alpha \in \{0, 0.1, 0.2, \ldots, 1\}$ and $\beta = 1 - \alpha$. $|I| = 20, |J| = 20, \Psi = 25, T = 150$. Distance function: Jaccard. $\Delta = 2$ (Left), $\Delta = 4$ (Right).

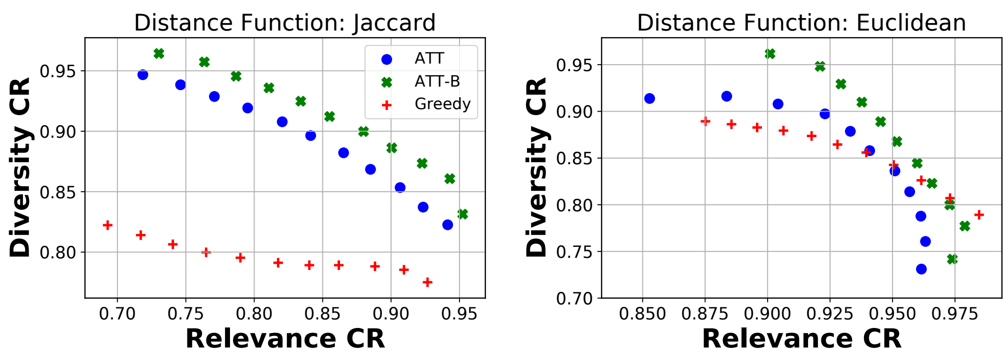

Fig. 3. AMT Dataset: comparison of competitive ratios (CR) achieved by ATT, ATT-B and Greedy with $\alpha \in \{0, 0.1, 0.2, \ldots, 1\}$ and $\beta = 1 - \alpha$. $|I| = 20, |J| = 20, \Psi = 25, \Delta = 2, T = 150$. Distance function: Jaccard (Left), Euclidean (Right).

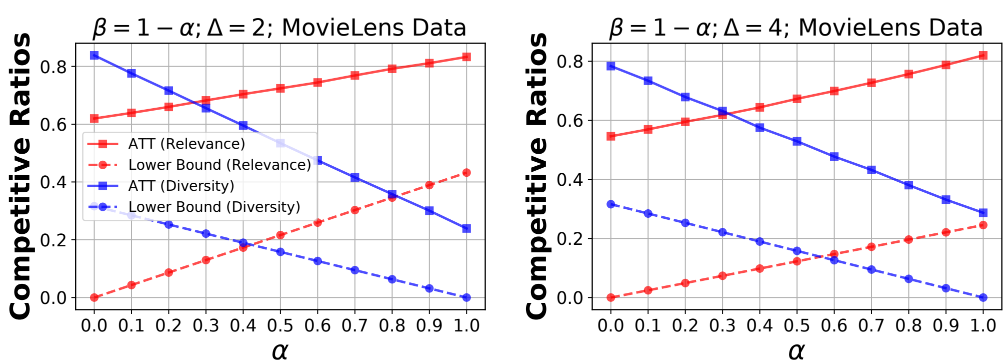

Fig. 4. MovieLens Dataset: competitive ratios (CR) for relevance and diversity with $\alpha \in \{0, 0.1, 0.2, \ldots, 1\}$ and $\beta = 1 - \alpha$. $|I| = 20, |J| = 20, \Psi = 10, T = 150$. Distance function: Jaccard. $\Delta = 2$ (Left), $\Delta = 4$ (Right).

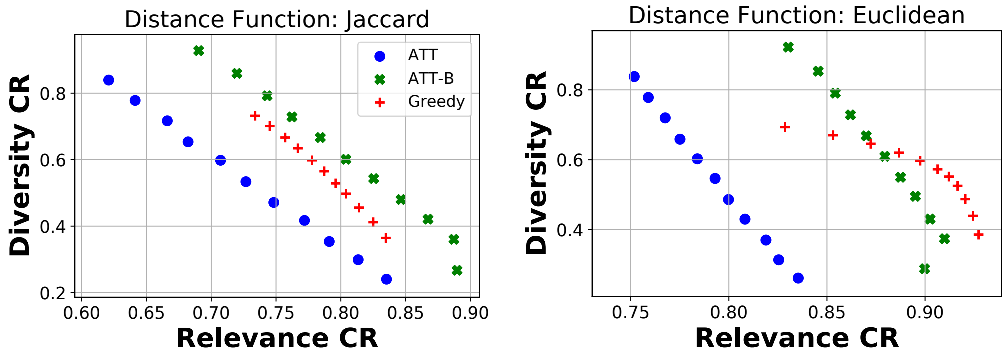

Fig. 5. MovieLens Dataset: comparison of competitive ratios (CR) of ATT, ATT-B and Greedy with $\alpha \in \{0, 0.1, 0.2, \ldots, 1\}$ and $\beta = 1 - \alpha$. $|I| = 20, |J| = 20, \Psi = 10, \Delta = 2, T = 150$. Distance function: Jaccard (Left), Euclidean (Right).

some essence with Greedy. There are cases where Greedy does better than ATT-B, though Greedy tends to be more unstable and vulnerable to the parameter setting. Specifically, on the AMT dataset with the Jaccard distance function, Greedy is strictly dominated by ATT. Also, LP-based algorithms such as ATT and ATT-B can provide a noticeable tradeoff between relevance and diversity under different choices of distance functions.

6 Conclusion and Future Directions

In this paper, we propose a generic online-matching-based model to study bi-objective optimization in online matching markets. Our proposed model captures many real-world applications such as matching tasks to workers in crowdsourcing markets, recommending movies to users in online recommendations, and building a team of workers for arriving tasks in the context of online team formation. We first construct a bi-objective linear program (LP) to serve as a valid upper bound; then design an LP-based parameterized online algorithm to effectively and smoothly balance the two goals. Via two separate parameters, our approaches allow the policy designer to specify overall targets they want to achieve for the two respective objectives. Our work opens a few directions for future research. A natural one is whether it is possible to beat one for the sum of the two competitive ratios. Perhaps we need more assumptions on the inherent relation between the two objectives.

Acknowledgement

Evan Yifan Xu was partially supported by the China Postdoctoral Science Foundation Funded Project under Grant 2025M771567. Pan Xu was partially supported by the NSF CRII Award IIS-1948157 and the BGU-NJIT seed grant.

References

[1] M. Adamczyk, F. Grandoni, and J. Mukherjee. 2015. Improved approximation algorithms for stochastic matching. In *Proceedings of the 23rd Annual European Symposium*. Springer, 1–12.

[2] F. Ahmed, J. P. Dickerson, and M. Fuge. 2017. Diverse weighted bipartite b-matching. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 35–41.

[3] M. Alsayasneh, S. Amer-Yahia, E. Gaussier, V. Leroy, J. Pilourdault, R. M. Borromeo, M. Toyama, and J.-M. Renders. 2017. Personalized and diverse task composition in crowdsourcing. *IEEE Transactions on Knowledge and Data Engineering*, 30, 1, 128–141.

[4] K. El-Arini, G. Veda, D. Shahaf, and C. Guestrin. 2009. Turning down the noise in the blogosphere. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 289–298.

[5] I. Ashlagi, M. Burq, C. Dutta, P. Jaillet, A. Saberi, and C. Sholley. 2019. Edge weighted online windowed matching. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 729–742.

[6] S. Assadi, J. Hsu, and S. Jabbari. 2015. Online assignment of heterogeneous tasks in crowdsourcing markets. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*. Vol. 3, 12–21.

[7] X. Bei and S. Zhang. 2018. Algorithms for trip-vehicle assignment in ride-sharing. In *Proceedings of the AAAI Conference on Artificial Intelligence* number 1. Vol. 32, 3–9.

[8] B. Brubach, K. A. Sankararaman, A. Srinivasan, and P. Xu. 2020. Attenuate locally, win globally: an attenuation-based framework for online stochastic matching with timeouts. *Algorithmica*, 82, 1, 64–87.

[9] B. Brubach, K. A. Sankararaman, A. Srinivasan, and P. Xu. 2016. New algorithms, better bounds, and a novel model for online stochastic matching. In *Proceedings of the 24th Annual European Symposium on Algorithms, ESA*. Vol. 57, 24:1–24:16.

[10] Y. H. Chan and L. C. Lau. 2012. On linear and semidefinite programming relaxations for hypergraph matching. *Mathematical programming*, 135, 1-2, 123–148.

[11] C. Chen, L. Zheng, V. Srinivasan, A. Thomo, K. Wu, and A. Sukow. 2016. Conflict-aware weighted bipartite b-matching and its application to e-commerce. *IEEE Transactions on Knowledge and Data Engineering*, 28, 6, 1475–1488.

[12] J. P. Dickerson, K. Sankararaman, A. Srinivasan, P. Xu, and Y. Xu. 2024. Matching tasks and workers under known arrival distributions: online task assignment with two-sided arrivals. *ACM Transactions on Economics and Computation*, 12, 2, 1–28.

[13] J. P. Dickerson, K. A. Sankararaman, A. Srinivasan, and P. Xu. 2021. Allocation problems in ride-sharing platforms: online matching with offline reusable resources. *ACM Transactions on Economics and Computation (TEAC)*, 9, 3, 1–17.

[14] J. P. Dickerson, K. A. Sankararaman, A. Srinivasan, and P. Xu. 2019. Balancing relevance and diversity in online bipartite matching via submodularity. In *Proceedings of the AAAI Conference on Artificial Intelligence* number 01. Vol. 33, 1877–1884.

[15] H. Esfandiari, N. Korula, and V. Mirrokni. 2016. Bi-objective online matching and submodular allocations. *Advances in Neural Information Processing Systems*, 29, 2747–2755.

[16] M. Esfandiari, R. M. Borromeo, S. Nikookar, P. Sakharkar, S. Amer-Yahia, and S. B. Roy. 2021. Multi-session diversity to improve user satisfaction in web applications. In *Proceedings of the Web Conference*, 1928–1936.

[17] Y. Feng, R. Niazadeh, and A. Saberi. 2019. Linear programming based online policies for real-time assortment of reusable resources. *Chicago Booth Research Paper*, 1, 20-25.

[18] M. Ge, C. Delgado-Battenfeld, and D. Jannach. 2010. Beyond accuracy: evaluating recommender systems by coverage and serendipity. In *Proceedings of the ACM conference on Recommender systems*, 257–260.

[19] B. Haeupler, V. S. Mirrokni, and M. Zadimoghaddam. 2011. Online stochastic weighted matching: improved approximation algorithms. In *Proceedings of the Internet and Network Economics*. Springer, 170–181.

[20] F. M. Harper and J. A. Konstan. 2015. The movielens datasets: history and context. *ACM Transactions on Interactive Intelligent Systems (TIIS)*, 5, 4, 1–19.

[21] C.-J. Ho and J. Vaughan. 2012. Online task assignment in crowdsourcing markets. In *Proceedings of the AAAI conference on artificial intelligence* number 1. Vol. 26, 45–51.

[22] P. Jaillet and X. Lu. 2013. Online stochastic matching: new algorithms with better bounds. *Mathematics of Operations Research*, 39, 3, 624–646.

[23] R. M. Karp, U. V. Vazirani, and V. V. Vazirani. 1990. An optimal algorithm for on-line bipartite matching. In *ACM symposium on Theory of computing*, 352–358.

[24] N. S. Lesmana, X. Zhang, and X. Bei. 2019. Balancing efficiency and fairness in on-demand ridesourcing. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 5309–5319.

[25] M. Lowalekar, P. Varakantham, and P. Jaillet. 2018. Online spatio-temporal matching in stochastic and dynamic domains. *Artificial Intelligence*, 261, 71–112.

[26] Y. Lu, S. Zhang, Y. Huang, L. Wang, X. Yu, Z. Zhao, and F. Wu. 2021. Future-aware diverse trends framework for recommendation. In *Proceedings of the Web Conference*, 2992–3001.

[27] W. Ma. 2014. Improvements and generalizations of stochastic knapsack and multi-armed bandit approximation algorithms. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 1154–1163.

[28] V. H. Manshadi, S. O. Gharan, and A. Saberi. 2012. Online stochastic matching: online actions based on offline statistics. *Mathematics of Operations Research*, 37, 4, 559–573.

[29] L. S. Marcolino, A. X. Jiang, and M. Tambe. 2013. Multi-agent team formation: diversity beats strength? In *Proceedings of the International Joint Conference on Artificial Intelligence*, 279–285.

[30] V. Nanda, P. Xu, K. A. Sankararaman, J. P. Dickerson, and A. Srinivasan. 2020. Balancing the tradeoff between profit and fairness in rideshare platforms during high-demand hours. In *Proceedings of the AAAI conference on artificial intelligence*, 2210–2217.

[31] B. Paudel and A. Bernstein. 2021. Random walks with erasure: diversifying personalized recommendations on social and information networks. In *Proceedings of the Web Conference 2021*, 2046–2057.

[32] J. Pilourdault, S. Amer-Yahia, S. B. Roy, and D. Lee. 2018. Task relevance and diversity as worker motivation in crowdsourcing. In *Proceedings of the International Conference on Data Engineering*, 365–376.

[33] S. A. Puthiya Parambath, N. Usunier, and Y. Grandvalet. 2016. A coverage-based approach to recommendation diversity on similarity graph. In *Proceedings of the ACM Conference on Recommender Systems*, 15–22.

[34] L. Qin and X. Zhu. 2013. Promoting diversity in recommendation by entropy regularizer. In *Proceedings of the International Joint Conference on Artificial Intelligence*. Vol. 13, 2698–2704.

[35] C. Sha, X. Wu, and J. Niu. 2016. A framework for recommending relevant and diverse items. In *Proceedings of the International Joint Conference on Artificial Intelligence*. Vol. 16, 3868–3874.

[36] H. Wang, E. Wang, Y. Yang, J. Wu, and F. Dressler. 2022. Privacy-preserving online task assignment in spatial crowdsourcing: a graph-based approach. In *Proceedings of the IEEE Conference on Computer Communications*. IEEE, 570–579.

[37] T. Wu, L. Chen, P. Hui, C. J. Zhang, and W. Li. 2015. Hear the whole story: towards the diversity of opinion in crowdsourcing markets. *Proceedings of the VLDB Endowment*, 8, 5, 485–496.

[38] G. Xiao, J. Zheng, and H. Dai. 2022. A unified model for bi-objective online stochastic bipartite matching with two-sided limited patience. In *Proceedings of the IEEE Conference on Computer Communications*. IEEE, 1079–1088.

[39] E. Y. Xu and P. Xu. 2024. Exploring the tradeoff between system profit and income equality among ride-hailing drivers. *Journal of Artificial Intelligence Research*, 79, 569–597.

[40] H. Yao, F. Wu, J. Ke, X. Tang, Y. Jia, S. Lu, P. Gong, J. Ye, and Z. Li. 2018. Deep multi-view spatial-temporal network for taxi demand prediction. In *Proceedings of the AAAI conference on artificial intelligence*, 2588–2595.

[41] J. Zhang, P. Lu, Z. Li, and J. Gan. 2018. Distributed trip selection game for public bike system with crowdsourcing. In *Proceedings of the IEEE Conference on Computer Communications*. IEEE, 2717–2725.

[42] B. Zhao, P. Xu, Y. Shi, Y. Tong, Z. Zhou, and Y. Zeng. 2019. Preference-aware task assignment in on-demand taxi dispatching: an online stable matching approach. In *Proceedings of the AAAI conference on artificial intelligence*, 2245–2252.

[43] C.-N. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen. 2005. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web*, 22–32.

Received 6 June 2024; accepted 14 August 2025