

# DISSECTING SUBMISSION LIMIT IN DESK-REJECTIONS: A MATHEMATICAL ANALYSIS OF FAIRNESS IN AI CONFERENCE POLICIES

Yuefan Cao\* Xiaoyu Li† Yingyu Liang‡ Zhizhou Sha§  
Zhenmei Shi¶ Zhao Song|| Jiahao Zhang\*\*

## ABSTRACT

As AI research surges in both impact and volume, conferences have imposed submission limits to maintain paper quality and alleviate organizational pressure. In this work, we examine the fairness of desk-rejection systems under submission limits and reveal that existing practices can result in substantial inequities. Specifically, we formally define the paper submission limit problem and identify a critical dilemma: when the number of authors exceeds three, it becomes impossible to reject papers solely based on excessive submissions without negatively impacting innocent authors. Thus, this issue may unfairly affect early-career researchers, as their submissions may be penalized due to co-authors with significantly higher submission counts, while senior researchers with numerous papers face minimal consequences. To address this, we propose an optimization-based fairness-aware desk-rejection mechanism and formally define two fairness metrics: individual fairness and group fairness. We prove that optimizing individual fairness is NP-hard, whereas group fairness can be efficiently optimized via linear programming. Through case studies, we demonstrate that our proposed system ensures greater equity than existing methods, including those used in CVPR 2025, offering a more socially just approach to managing excessive submissions in AI conferences.

## 1 INTRODUCTION

We are living in an era shaped by the unprecedented advancements of Artificial Intelligence (AI), where transformative breakthroughs have emerged across various domains in just a few years. A key driving force behind AI’s rapid progress is the prevalence of top conferences held frequently throughout the year, offering dynamic platforms to present many of the field’s most influential papers. For example, ResNet He et al. (2016), a foundational milestone in deep learning with over 250,000 citations, was first introduced at CVPR 2016. Similarly, the Transformer architecture Vaswani et al. (2017), the backbone of modern large language models, emerged at NeurIPS 2017. More recently, diffusion models Ho et al. (2020), which represent the state-of-the-art in image generation, were presented at NeurIPS 2020, while CLIP Radford et al. (2021), a leading model for image-text pretraining, was showcased at ICML 2021. These groundbreaking contributions from top-tier conferences have significantly accelerated the advancement of AI, enriching both theoretical insights and practical applications.

As AI continues to expand its applications in real-world domains such as dialogue systems Schulman et al. (2022); Achiam et al. (2023); Anthropic (2024), image generation Ho et al. (2020); Song et al. (2020); Wang et al. (2024b; 2023; 2024a), and video generation Ho et al. (2022); Blattmann et al.

\* ralph1997off@gmail.com. Zhejiang University.

† xiaoyu.li2@student.unsw.edu.au. University of New South Wales.

‡ yingyul@hku.hk. The University of Hong Kong. yliang@cs.wisc.edu. University of Wisconsin-Madison.

§ shazz20@mails.tsinghua.edu.cn. Tsinghua University.

¶ zhmeishi@cs.wisc.edu. University of Wisconsin-Madison.

|| magic.linuxkde@gmail.com. The Simons Institute for the Theory of Computing at UC Berkeley.

\*\* ml.jiahaozhang02@gmail.com. Independent Researcher.

(2023), its immense potential for commercialization has raised growing enthusiasm in AI research. This enthusiasm has led to a rapid, rocket-like increase in the number of AI-related papers in 2024, as witnessed by recent studies Stanford (2024). A direct consequence of this surge is the significant rise in submissions to AI conferences, which has placed a heavy burden on program committees tasked with selecting papers for acceptance. To address these challenges and maintain the quality of accepted papers, many leading conferences have introduced submission limits per author. In 2025, a wide range of leading AI conferences, including CVPR, ICCV, AAAI, WSDM, IJCAI, and KDD, have introduced submission limits per author in their guidelines, ranging from a maximum of 7 to 25. Table 1 provides an overview of these submission limits across major AI conferences.

Table 1: In this table, we summarize the submission limits of top conferences in recent years. For details of each conference website, we refer the readers to Section E in the Appendix. Some conferences (CVPR, ICCV, WSDM, KDD) employ a conventional desk-reject algorithm (Algorithm 2), where papers are desk-rejected once an author has registered more than  $x$  (submission limit) papers.

Conference Name	Year	Upper Bound
CVPR	2025	25
CVPR	2024	N/A
ICCV	2025	25
ICCV	2023	N/A
AAAI	2023-2025	10
AAAI	2022	N/A
WSDM	2021-2025	10
WSDM	2020	N/A
IJCAI	2021-2025	8
IJCAI	2020	6
IJCAI	2018-2019	10
IJCAI	2017	N/A
KDD	2024-2025	7
KDD	2023	N/A

However, such a desk-rejection mechanism may result in unintended negative societal impacts due to the Matthew effect in the research community Bol et al. (2018), as illustrated in Figure 1. Recent research has shown that the impact of a setback (e.g., a paper rejection) is often much greater for early-career researchers than for senior researchers Wang et al. (2019); Sun et al. (2023), which shows that the effect of a desk rejection can vary significantly depending on the author’s career stage. For instance, as illustrated in Figure 2, consider the case of a young student submitting their only draft to the conference, co-authored with a renowned researcher who submits numerous papers annually. If the paper is desk-rejected due to exceeding submission limits, the senior researcher might view this as a negligible inconvenience. In contrast, the rejection could have severe consequences for the student, as the paper might be crucial for applying to graduate programs, securing employment, or forming a chapter of their thesis. This disparity in the impact of desk rejections may worsen the Matthew effect in the AI community by disproportionately disadvantaging researchers with only one or two submitted papers, while having little effect on prolific senior researchers. Such outcomes raise important concerns about fairness and equity in current desk-rejection policies.

In response to the challenges posed by paper-limit-based desk-rejection systems, this work investigates an important and practical problem: ensuring fairness in desk-rejection systems for AI conferences under submission limits. As illustrated in Figure 3, our goal is to design a fair desk-rejection system that prioritizes rejecting submissions from authors with many papers while protecting those with fewer submissions, particularly early-career researchers. Our key contributions are as follows:

- We formally define the paper submission limit problem in desk-rejection systems and prove that an ideal system that rejects papers solely based on each author’s excessive submissions is mathematically impossible when there are more than three authors.
- We introduce two fairness metrics: individual fairness and group fairness. We formulate the fairness-aware paper submission limit problem as an integer programming problem. Specifically, we formally prove that optimizing individual fairness is NP-hard, while the

group fairness optimization problem can be solved efficiently using any off-the-shelf linear programming solver.

- Through case studies, we demonstrate that our proposed system achieves greater fairness compared to existing approaches used in top AI conferences such as CVPR 2025, promoting social justice and fostering a more inclusive ML research community.

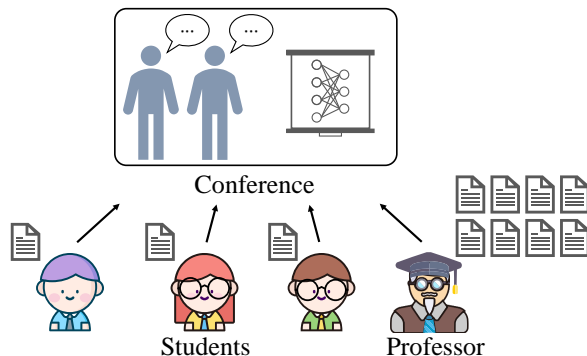


Figure 1: The Matthew Effect in the AI community. This figure illustrates the worsening Matthew Effect in the AI community, where senior researchers tend to have a significantly higher number of submissions, while junior researchers have relatively few.

**Roadmap.** Our paper is organized as follows: In Section 2, we review related literature. In Section 3, we present the key definition of the paper submission limit problem. In Section 4, we show that no algorithm can reach the ideal desk-rejection system without unfair collective punishments. In Section 5, we present our new fairness-aware desk-rejection system. In Section 6, we show by case studies that our system is better than existing systems. In Section 7, we present our conclusions.

## 2 RELATED WORKS

### 2.1 DESK REJECT MECHANISM

A wide range of desk-rejection mechanisms have been developed to reduce the human effort involved in the peer-review process Ansell & Samuels (2021). One of the most widely adopted desk-rejection rules is rejecting papers that violate anonymity requirements Jefferson et al. (2002); Tennant (2018). This rule is crucial for maintaining unbiased evaluations of researchers from diverse institutions and career levels while preventing conflicts of interest. Another common mechanism addresses duplicate and dual submissions Stone (2003); Leopold (2013), alleviating the duplication of reviewer efforts across multiple venues and upholding ethical publication standards. Additionally, plagiarism King & ChatGPT (2023); Elali & Rachid (2023) is a major concern in desk rejections at AI conferences, as it undermines the integrity of the academic community, violates intellectual property rights, and compromises the originality and credibility of research. In response to the growing number of submissions to AI conferences, new types of desk-rejection rules have recently emerged Leyton-Brown et al. (2024). For example, IJCAI 2020 and NeurIPS 2020 implemented a fast desk-rejection mechanism, allowing area chairs to reject papers based on a quick review of the abstract and main content to manage the review workload. However, this approach introduced noise and sometimes resulted in the rejection of generally good papers, leading to its reduced prevalence compared to more systematic mechanisms like enforcing submission limits, which is the main focus of this paper. To the best of our knowledge, limited literature has explored these emerging desk-rejection techniques, and our work is among the first to formally study the desk-rejection mechanism based on maximum submission limits.

### 2.2 THE COMPETITIVE RACE IN AI PUBLICATION

Due to the rapid increase in submissions to AI conferences in recent years Stanford (2024), concerns about the intense competition in these conferences are growing. As Bengio Yoshua noted Bengio

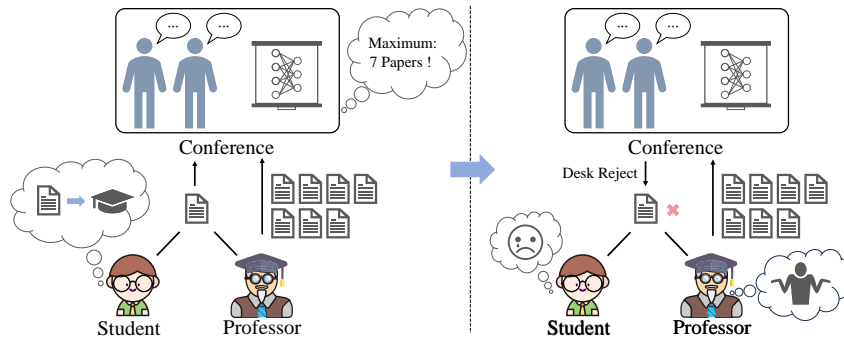


Figure 2: The unfairness of desk rejection based on submission limits. **Left: A careless mistake.** In this scenario, a young student submits the only paper, co-authored with a professor who submits numerous papers, and carelessly exceeds the submission limit. The paper, which may aim to apply to graduate programs, secure employment, or form a chapter of the thesis, is very important for the student but may not be for the professor. **Right: The desk rejection.** If the paper is desk-rejected due to submission limits, it poses a minor inconvenience to the professor, and the professor can shrug about it due to his remaining papers. However, it could have severe consequences for the students, as the paper is crucial for the student’s future plans.

(2020): “It is more competitive, everything is happening fast and putting a lot of pressure on everyone. The field has grown exponentially in size. Students are more protective of their ideas and in a hurry to put them out, by fear that someone else would be working on the same thing elsewhere, and in general, a PhD ends up with at least 50% more papers than what I gather it was 20 or 30 years ago.” Consequently, paper acceptance has become increasingly critical in AI job applications Ahmed (2022); Besiroglu et al. (2024), as having more papers is now the norm. Therefore, it is crucial to establish fair and practical guidelines for desk rejections Teixeira da Silva et al. (2018), ensuring that every group of authors is treated equitably in AI conferences.

### 2.3 FAIRNESS SYSTEM DESIGN

Fairness Francez (2012); Mehrabi et al. (2021) is a key principle of social justice, reflecting the absence of bias toward individuals or groups based on inherent characteristics. Due to its profound societal impact, fairness has become an essential consideration in the design of algorithms across various computer systems that interact with human factors. In recommender systems, fairness can manifest in various forms, such as item fairness Zhang et al. (2021); Ge et al. (2021), which ensures that items from different categories or with varying levels of prior exposure are recommended equitably, and user fairness Li et al. (2021a;b), which guarantees that all users, regardless of their backgrounds or preferences, have equal opportunities to access relevant content. These fairness measures help balance opportunities for both users and retailers, fostering equity in the recommendation process. In candidate selection systems Gilliland (1993); Wang et al. (2020), fairness ensures that all candidates are evaluated solely on merit, independent of factors such as race, gender, or socio-economic background, promoting equality and ensuring that the selection processes are inclusive. In information access systems Ekstrand et al. (2022), including job search Wu et al. (2022) and music discovery Melchiorre et al. (2021), fairness guarantees that all individuals can access the information they need without discrimination, ensuring equal opportunities for users to make informed decisions. Similarly, in dialog systems Guo et al. (2022); Gallegos et al. (2024), fairness ensures that language models avoid generating biased text or making inappropriate word-context associations related to social groups, supporting equitable and respectful interactions. Moreover, recent research has investigated group fairness in peer review processes for AI conferences, highlighting the importance of equitable evaluation for submissions Aziz et al. (2023). Despite the widespread focus on fairness in algorithmic design, the fairness of desk-rejection mechanisms remains an open question and serves as the primary motivation for this paper.

### 3 PRELIMINARY

In this section, we first introduce the notations in Section 3.1. Then, we present the general problem formulation in Section 3.2.

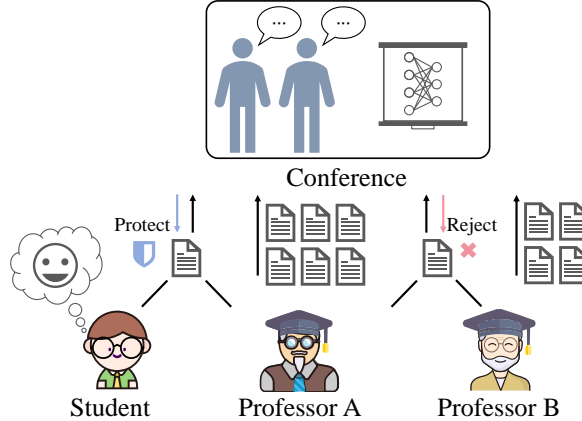


Figure 3: Our research objective. This figure presents the goal of our study: creating a more equitable desk-rejection system. Consider Professor A, who has carelessly submitted numerous papers exceeding the submission limit, collaborating with another senior researcher (Professor B) with many submissions, and a young student with only one paper. Our proposed system prioritizes desk-rejecting papers from authors with a large number of submissions first, thereby increasing the student’s chances of having their paper accepted. This approach aims to mitigate the disparity in the impact of desk rejections and promote fairness.

#### 3.1 NOTATIONS

For any positive integer  $n$ , we use  $[n]$  to denote the set  $\{1, 2, \dots, n\}$ . We use  $\mathbb{N}_+$  to represent the set of all positive integers. For two sets  $\mathcal{B}$  and  $\mathcal{C}$ , we denote the set difference as  $\mathcal{B} \setminus \mathcal{C} := \{x \in \mathcal{B} : x \notin \mathcal{C}\}$ . For a vector  $x \in \mathbb{R}^d$ ,  $\text{Diag}(d)$  denotes a diagonal matrix  $X \in \mathbb{R}^{d \times d}$ , where the diagonal entries satisfy  $X_{i,i} = x_i$  for all  $i \in [d]$ , and all off-diagonal entries are zero. We use  $\mathbf{1}_n$  to denote an  $n$ -dimensional column vector with all entries equal to one.

#### 3.2 PROBLEM FORMULATION

In this section, we further introduce the actual problem we will investigate in this paper, where we begin with introducing the definition for three kinds of authors that will appear later in our discussion.

**Definition 3.1** (Submission Limit Problem). *Let  $\mathcal{A} = \{a_1, a_2, \dots, a_n\}$  denote the set of  $n$  authors, and let  $\mathcal{P} = \{p_1, p_2, \dots, p_m\}$  denote the set of  $m$  papers. Each author  $a_i \in \mathcal{A}$  has a subset of papers  $P_i \subseteq \mathcal{P}$ , and each paper  $p_j \in \mathcal{P}$  is authored by a subset of authors  $A_j \subseteq \mathcal{A}$ . For each author,  $a_i \in \mathcal{A}$ , let  $C_i$  denote the set of all coauthors of  $a_i$  and let  $x \in \mathbb{N}_+$  denote the maximum number of papers each author can submit.*

The goal is to find a subset  $S \subseteq \mathcal{P}$  of papers (to keep) such that for every  $a_i \in \mathcal{A}$ ,

$$\underbrace{|\{p_j \in S : a_i \in A_j\}|}_{\text{\#remained papers of author } a_i} \leq x.$$

or equivalently find a subset  $\bar{S} \subseteq \mathcal{P}$  of papers (to reject) such that for every  $a_i \in \mathcal{A}$ ,

$$|P_i| - \underbrace{|\{j \in \bar{S} : i \in A_j\}|}_{\text{\#rejected papers of author } a_i} \leq x.$$

We now present several fundamental facts related to Definition 3.1, which can be easily verified through basic set theory.

**Fact 3.2.** For any author  $a_i \in \mathcal{A}$  and paper  $p_j \in \mathcal{P}$ ,  $a_i \in A_j$  if and only if  $p_j \in P_i$ .

**Fact 3.3.** For each author  $a_i \in \mathcal{A}$ , the number of papers submitted by the author can be formulated as:

$$|P_i| = |\{p_j \in \mathcal{P} : a_i \in A_j\}|.$$

**Fact 3.4.** For each paper  $j \in [m]$ , the number of authors of this paper can be formulated as:

$$|A_j| = |\{a_i \in \mathcal{A} : p_j \in P_i\}|.$$

**Fact 3.5.** For each author  $a_i \in \mathcal{A}$ , the set of coauthors for author  $a_i$  can be formulated as:

$$C_i = \left( \bigcup_{p_j \in P_i} A_j \right) \setminus \{a_i\}.$$

## 4 THE DESK REJECTION DILEMMA

In this section, we define the concept of an ideal desk-rejection system in Section 4.1 and formally demonstrate in Section 4.2 that no algorithm can achieve this ideal system.

### 4.1 IDEAL DESK-REJECTION

An ideal desk-rejection system should avoid unfairly rejecting papers from authors who either comply with the submission limit or exceed it by only one or two papers. Otherwise, authors may face consequences due to co-authors with an excessively high number of submissions.

To address this, we formally define the criteria for an ideal desk-rejection outcome for the problem in Definition 3.1, where rejections are based solely on an author’s excessive submissions, without unfairly penalizing others.

**Definition 4.1** (Ideal desk-rejection). *An ideal solution for the submission limit problem in Definition 3.1 is a paper subset  $S \subseteq \mathcal{P}$  such that every author has exactly  $\min\{x, |P_i|\}$  papers remaining after desk rejection.*

**Remark 4.2.** *The ideal desk-rejection in Definition 4.1 ensures that innocent authors with less than  $x$  submissions will retain all their papers, and a non-compliant author  $a_i$  with more than  $x$  submissions will be desk-rejected exact  $(|P_i| - x)$  papers.*

Thus, if there exists an algorithm that can reach the aforementioned ideal solution, we can ensure that no author is unfairly penalized due to their co-authors’ submission behavior, achieving both fairness and individual accountability.

### 4.2 HARDNESS OF IDEAL DESK-REJECTION

Unfortunately, we find that achieving an ideal desk-rejection system is fundamentally intractable. The main result regarding this hardness is presented in the following theorem:

**Theorem 4.3** (Hardness of Ideal Desk-Rejection). *Let  $n = |\mathcal{A}|$  denote the number of authors in Definition 3.1. We can show that*

- **Part 1:** For  $n \leq 2$ , there always exists an algorithm that can achieve the ideal desk-rejection in Definition 4.1.
- **Part 2:** For  $n \geq 3$ , there exists at least one problem instance where no algorithm can guarantee achieving the ideal desk-rejection in Definition 4.1.

*Proof.* For **Part 1**, the result follows directly from Lemma B.3 and Lemma B.5. For **Part 2**, the result is established using Lemma B.6 and Lemma B.7. Detailed technical proofs for these lemmas are provided in Appendix B.  $\square$

Therefore, since an ideal desk-rejection system is not achievable, it is inevitable that some authors may face excessive desk-rejections due to collective punishments. This challenge is particularly concerning for early-career researchers with only one or two submissions, motivating the need to seek an approximate solution that optimizes fairness in desk-rejection systems.

## 5 FAIRNESS-AWARE DESK-REJECTION

In this section, we first introduce two fairness metrics in Section 5.1, and then present the hardness result on minimizing one of them in Section 4.2. In Section 5.3, we show our optimization-based fairness-aware desk-rejection framework.

### 5.1 FAIRNESS METRICS

As discussed earlier, achieving an ideal desk-rejection system is practically infeasible, as unintended rejections due to collective punishments are unavoidable. To address this, we relax the ideal system into an approximate form, where some unfair desk-rejections are permitted, while these rejections should be proportional to each author’s total number of submissions.

Specifically, we introduce a cost function for each author, which estimates the impact of desk-rejection on each author:

**Definition 5.1** (Cost Function). *Considering the submission limit problem in Definition 3.1, we define the cost function  $c : [n] \times 2^{[m]} \rightarrow [0, 1]$  for a specific author  $a_i$  and a set of remaining paper  $S$  as*

$$c(a_i, S) := \frac{|P_i| - |\{p_j \in S : a_i \in A_j\}|}{|P_i|}.$$

**Remark 5.2.** *The cost function  $c(a_i, S)$  measures the proportion of papers authored by  $a_i$  that are rejected, prioritizing fairness for early-career authors with fewer submissions and aiming to reduce setbacks for them.*

To further demonstrate how this author-wise cost function could benefit fairness, we present the following example:

**Example 5.3.** *Consider a submission limit problem with  $x = 10$  and  $n = 2$ . Suppose author  $a_1$  submits papers  $p_1, p_2, \dots, p_{11}$ , and author  $a_2$  submits only paper  $p_{11}$ . Rejecting paper  $p_{11}$  (i.e.,  $S = \mathcal{P} \setminus \{p_{11}\}$ ) results in a cost of  $c(a_1, S) = 1/11$  for  $a_1$  but a cost of  $c(a_2, S) = 1$  for  $a_2$ , which is unfair to  $a_2$ . On the other hand, if we reject paper  $p_1$  (i.e.,  $S' = \mathcal{P} \setminus \{p_1\}$ ), the cost for  $a_1$  remains  $c(a_1, S') = 1/11$ , while the cost for  $a_2$  becomes  $c(a_2, S') = 0$ . This minimizes both the highest cost and the total cost. This example demonstrates that our cost function encourages rejecting papers from authors with many submissions while protecting authors with few submissions.*

To ensure fair treatment for all authors and avoid imposing excessive setbacks on early-career researchers, we introduce two fairness metrics based on our cost function. These metrics are inspired by the principles of utilitarian social welfare and egalitarian social welfare Aziz et al. (2024). We begin by defining individual fairness, which is a strict worst-case fairness metric that aligns with the egalitarian social welfare framework by estimating the individual cost among all authors.

**Definition 5.4** (Individual Fairness). *Let  $c : [n] \times 2^{[m]} \rightarrow [0, 1]$  be the cost function defined in Definition 5.1. We define function  $\zeta_{\text{ind}} : 2^{[m]} \rightarrow [0, 1]$  to measure the individual fairness:*

$$\zeta_{\text{ind}}(S) := \max_{i \in [n]} c(a_i, S).$$

Next, we present the concept of group fairness, which aligns with utilitarian social welfare and measures the total cost across all authors.

**Definition 5.5** (Group Fairness). *Let  $c : [n] \times 2^{[m]} \rightarrow [0, 1]$  be the cost function defined in Definition 5.1. We define function  $\zeta_{\text{group}} : 2^{[m]} \rightarrow [0, 1]$  to measure the group fairness:*

$$\zeta_{\text{group}}(S) := \frac{1}{n} \sum_{i \in [n]} c(a_i, S).$$

To show the relationship between these two fairness metrics, we have the following proposition:

**Proposition 5.6** (Relationship of Fairness Metrics, informal version of Proposition C.1 in Appendix C). *For any solution  $S \subseteq \mathcal{P}$  to the submission limit problem in Definition 3.1, we have*

$$\zeta_{\text{ind}}(S) \leq \zeta_{\text{group}}(S).$$

## 5.2 HARDNESS OF INDIVIDUAL FAIRNESS-AWARE SUBMISSION LIMIT PROBLEM

After presenting fairness metrics for the desk-rejection system, we introduce an optimization-based framework to address these metrics. We first study the individual fairness-aware submission limit problem to minimize the individual fairness measure  $\zeta_{\text{ind}}$  in Definition 5.4.

**Definition 5.7** (Individual Fairness-Aware Submission Limit Problem). *We consider the following optimization problem:*

$$\begin{aligned} & \min_{S \subseteq \mathcal{P}} \zeta_{\text{ind}}(S) \\ & \text{s.t. } |\{p_j \in S : a_i \in A_j\}| \leq x, \quad \forall a_i \in \mathcal{A}. \end{aligned}$$

To represent the fairness metric minimization problem in matrix form, we introduce Definition 5.8.

**Definition 5.8** (Author-Paper Matrix). *Let  $W \in \{0, 1\}^{n \times m}$  denote the author-paper matrix for the author set  $\mathcal{A}$  and paper set  $\mathcal{P}$ . Then, we define  $W_{i,j} = 1$  if author  $a_i$  is a coauthor of paper  $p_j$ , and  $W_{i,j} = 0$  otherwise.*

Therefore, we present a more tractable integer programming form of the original problem and prove its equivalence to the original formulation:

**Definition 5.9** (Individual Fairness-Aware Submission Limit Problem, Matrix Form). *We consider the following integer optimization problem:*

$$\begin{aligned} & \min_{r \in \{0,1\}^m} \|\mathbf{1}_n - D^{-1}Wr\|_{\infty} \\ & \text{s.t. } (Wr)/x \leq \mathbf{1}_n \end{aligned}$$

where  $D = \text{Diag}(|P_1|, \dots, |P_n|)$ , and the rejection vector  $r \in \{0, 1\}^m$  is a 0-1 vector, with  $r_j = 1$  indicating that paper  $p_j$  is remained, and  $r_j = 0$  indicating that it is desk-rejected.

**Proposition 5.10** (Matrix Form Equivalence for  $\zeta_{\text{ind}}$ , informal version of Proposition C.3 in Appendix C). *The individual fairness-aware submission limit problem in Definition 5.7 and the matrix form integer programming problem in Definition 5.9 are equivalent.*

Unfortunately, solving this integer programming problem is highly non-trivial, which means it may not yield a feasible solution within a reasonable time for large-scale conference submission systems. We establish the computational hardness of this problem in the following theorem:

**Theorem 5.11** (Hardness, informal version of Theorem C.7 in Appendix C.2). *The Individual Fairness-Aware Submission Limit Problem defined in Definition 5.7 is NP-hard.*

Since minimizing individual fairness is computationally intractable, our fairness-aware desk-rejection system instead focuses on minimizing group fairness.

## 5.3 GROUP FAIRNESS OPTIMIZATION

Given the inherent hardness of individual fairness optimization, we address the fairness problem using an alternative yet equally important metric: group fairness, as defined in Definition 5.5. This metric is not only a crucial fairness measure in its own right but also serves as a lower bound for individual fairness as stated in Proposition 5.6, potentially improving individual fairness implicitly.

Following a similar approach in Section 5.2, we first formulate the submission limit problem with respect to group fairness and derive a tractable integer programming formulation in matrix form:

**Definition 5.12** (Group Fairness-Aware Submission Limit Problem). *We consider the following optimization problem:*

$$\begin{aligned} & \min_{S \subseteq \mathcal{P}} \zeta_{\text{group}}(S) \\ & \text{s.t. } |\{p_j \in S : a_i \in A_j\}| \leq x, \quad \forall a_i \in \mathcal{A}. \end{aligned}$$

**Definition 5.13** (Group Fairness-Aware Submission Limit Problem, Matrix Form). *We consider the following integer programming problem:*

$$\max_{r \in \{0,1\}^m} \mathbf{1}_n^{\top} D^{-1}Wr$$



**Algorithm 1** Fairness-Aware Desk-Reject Algorithm

---

```

1: /*  $\mathcal{A}$  denotes the set of  $n$  authors. */
2: /*  $\mathcal{P}$  denote the set of  $m$  papers. */
3: /* Author  $a_i \in \mathcal{N}$  has a subset of papers  $P_i \subset \mathcal{P}$ . */
4: /* Paper  $p_j \in \mathcal{P}$  is coauthored by a subset of authors  $A_j \subseteq \mathcal{A}$ .*/
5: /*  $x$  represents the submission limit for each author.*/
6: procedure FAIRDESKREJECT( $\mathcal{A}, \mathcal{P}, x$ )
7:   /* Initialize the constants of the problem. */
8:   for  $i \in [n], j \in [m]$  do
9:     if  $p_j \in \mathcal{A}_i$  then
10:       $W_{i,j} \leftarrow 1$ 
11:     else
12:       $W_{i,j} \leftarrow 0$ 
13:     end if
14:   end for
15:    $D \leftarrow \text{Diag}(|P_1|, \dots, |P_n|)$ 
16:   /* Solve the linear programming problem in Definition 5.15. */
17:    $r^* \leftarrow \text{LPSolver}(W, D, x, r^0)$ 
18:   /* Transform the solution. */
19:    $S \leftarrow \emptyset$ 
20:   for  $j \in [m]$  do
21:     if  $r_j = 1$  then
22:        $S \leftarrow S \cup \{p_j\}$ 
23:     end if
24:   end for
25:   return  $S$ 
26: end procedure

```

---

$$\text{s.t. } (Wr)/x \leq \mathbf{1}_n,$$

where  $D = \text{Diag}(|P_1|, \dots, |P_n|)$ , and the rejection vector  $r \in \{0, 1\}^m$  is a 0-1 vector, with  $r_j = 1$  indicating that paper  $p_j$  is remained, and  $r_j = 0$  indicating that it is desk-rejected.

**Proposition 5.14** (Matrix Form Equivalence for  $\zeta_{\text{group}}$ , informal version of Proposition C.8 in Appendix C). *The fairness-aware submission limit problem in Definition 5.12 and the matrix form integer programming problem in Definition 5.13 are equivalent.*

However, solving integer programming problems is practically challenging. To this end, we first relax the feasible region of  $r$  to  $[0, 1]^m$ , and then analyze the resulting relaxed problem.

**Definition 5.15** (Group Fairness-Aware Submission Limit Problem, Relaxation). *We consider the optimization problem*

$$\begin{aligned} & \max_{r \in [0, 1]^m} \mathbf{1}_n^\top D^{-1} W r \\ & \text{s.t. } (Wr)/x \leq \mathbf{1}_n, \end{aligned}$$

where  $D = \text{Diag}(|P_1|, \dots, |P_n|)$ , and the rejection vector  $r \in \{0, 1\}^m$  is a 0-1 vector, with  $r_j = 1$  indicating that paper  $p_j$  is remained, and  $r_j = 0$  indicating that it is desk-rejected.

Fortunately, the relaxed problem is a linear program, which can be efficiently solved using standard linear programming solvers. Moreover, its optimal solution is equivalent to that of the original integer programming problem, and this result is formalized in the following theorem:

**Theorem 5.16** (Optimal Solution Equivalence of the Relaxed Problem, informal version of Theorem C.9). *The optimal solution of the relaxed linear programming problem in Definition 5.15 is equivalent to the optimal solution of the original integer programming problem in Definition 5.13.*

This theoretical result is significant as we formally establish that the group fairness-aware submission problem in Definition 5.12 reduces to a linear programming (LP) problem with guaranteed optimality, solvable using off-the-shelf LP solvers. We formalize this procedure in Algorithm 1, where

LPSolver denotes any standard LP solver, including but not limited to the simplex method Bar-tels & Golub (1969), interior-point path-finding methods Lee & Sidford (2014), and state-of-the-art stochastic central path methods Cohen et al. (2019); Jiang et al. (2021).

**Remark 5.17.** *The time complexity of our fairness-aware desk-rejection algorithm in Algorithm 1 aligns with modern linear programming solvers. For instance, using the stochastic central path method Cohen et al. (2021); Jiang et al. (2021); van den Brand et al. (2020); Song & Yu (2021), it achieves a time complexity of  $O^*(m^{2.37} \log(m/\delta))$ , where  $\delta$  represents the relative accuracy corresponding to a  $(1 + \delta)$ -approximation guarantee.*

**Remark 5.18.** *In practice, major AI conferences routinely process submissions at the scale of  $m \sim 10^4$  Stanford (2024). Given this regime, our algorithm guarantees efficient computation, enabling fairness-aware desk rejection within tractable timeframes, even for large-scale conferences.*

## 6 CASE STUDY

Since desk-rejection data from top AI conferences is not publicly available, and fully open-review conferences like ICLR do not impose submission limits, evaluating real-world conference submissions is impractical. Therefore, we present a case study to demonstrate how our proposed desk-rejection algorithm more effectively addresses fairness issues. Additional case studies are provided in Appendix D. Let the paper subscript  $j$  in  $p_j \in \mathcal{P}$  denote the submission order. We analyze the widely used desk-rejection system (e.g., CVPR 2025) in Algorithm 2, which rejects all papers submitted after an author’s  $x$ -th submission. To highlight its limitations, we present a minimal example:

**Example 6.1.** *Consider a submission limit problem as defined in Definition 3.1 with  $n = 2$ ,  $x = 25$ , and  $m = 26$ . Author  $a_1$  submits all papers  $p_1, \dots, p_{26}$ , while author  $a_2$  submits only  $p_{26}$ .*

Given the ideal desk-rejection criteria in Definition 4.1, it is evident that we can reject any papers in  $\{p_1, p_2, \dots, p_{25}\}$  following the techniques in Lemma B.4. After rejection, since  $a_1$  retains 25 papers and  $a_2$  retains 1 paper, the fairness metrics are  $\zeta_{\text{ind}}(S) = \max\{1/26, 0\} = 1/26$  and  $\zeta_{\text{group}}(S) = \frac{1}{2}(1/26 + 0) = 1/52$ . On the other hand, the CVPR 2025 algorithm, as described in Algorithm 2, rejects  $p_{26}$ , retaining  $S = \{p_1, \dots, p_{25}\}$ . This unfairly penalizes  $a_2$ , resulting in  $\zeta_{\text{ind}}(S) = \max\{1/26, 1\} = 1$  and  $\zeta_{\text{group}}(S) = \frac{1}{2}(1/26 + 1) = 27/52$ , which is much worse compared with the ideal results. In contrast, our method in Algorithm 1 solves the linear program and recovers the ideal solution, achieving the same fairness metrics as the optimal case. A simple workaround to mitigate unfairness in conventional desk-rejection systems is the roulette algorithm, which randomly rejects papers from non-compliant authors like  $a_1$  until the submission limit  $x$  is reached. However, this heuristic cannot fully prevent the rejection of the undesirable paper  $p_{26}$  and results in suboptimal fairness outcomes compared to our fairness-aware rejection, since the expected fairness metrics under the roulette algorithm satisfy  $\mathbb{E}[\zeta_{\text{ind}}] = (25/26) \cdot (1/26) + (1/26) \cdot 1 \leq 1/26$  and  $\mathbb{E}[\zeta_{\text{group}}] = (25/26) \cdot (1/52) + (1/26) \cdot (27/52) \leq 27/52$ . Thus, this example illustrates that conventional desk-rejection systems in top conferences such as CVPR can suffer from severe fairness issues, whereas our proposed method effectively mitigates these problems. Additionally, this example also highlights another noteworthy consequence of the conventional desk-rejection system. Specifically, authors collaborating with senior researchers who have numerous submissions may have to compete for earlier submission slots to avoid desk rejection. However, the submission order should not influence whether a paper is accepted, which reveals the unintended implications of the order-based desk-rejection system.

## 7 CONCLUSION

In this work, we identify the fairness issue in the desk-rejection mechanisms of AI conferences under submission limits. Our theoretical analysis proves that an ideal system that rejects papers solely based on authors’ non-compliance, without unfairly penalizing others due to collective punishment, is impossible. We further consider an optimization-based fairness-aware desk-rejection system to alleviate the unfairness problem. In this system, we considered two fairness metrics: individual fairness and group fairness. We formally established that optimizing individual fairness in desk-rejection is NP-hard, while optimizing group fairness can be reduced to a linear programming problem that can be solved highly efficiently. Through case studies, we showed that the proposed method outperforms the existing desk-rejection system with submission limits in top conferences.

## REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Nur Ahmed. Scientific labor market and firm-level appropriation strategy in artificial intelligence research. Technical report, MIT Sloan Working Paper, 2022.
- Ben W Ansell and David J Samuels. Desk rejecting: A better use of your time. *PS: Political Science & Politics*, 54(4):686–689, 2021.
- Anthropic. The claude 3 model family: Opus, sonnet, haiku, 2024. URL <https://www-cdn.anthropic.com>.
- Haris Aziz, Evi Micha, and Nisarg Shah. Group fairness in peer review. *Advances in Neural Information Processing Systems*, 36, 2023.
- Haris Aziz, Evi Micha, and Nisarg Shah. Group fairness in peer review. *Advances in Neural Information Processing Systems*, 36, 2024.
- Richard H Bartels and Gene H Golub. The simplex method of linear programming using lu decomposition. *Communications of the ACM*, 12(5):266–268, 1969.
- Yoshua Bengio. Time to rethink the publication process in machine learning. <https://yoshuabengio.org/2020/02/26/time-to-rethink-the-publication-process-in-machine-learning/>, 2020.
- Tamay Besiroglu, Sage Andrus Bergerson, Amelia Michael, Lennart Heim, Xueyun Luo, and Neil Thompson. The compute divide in machine learning: A threat to academic contribution and scrutiny? *arXiv preprint arXiv:2401.02452*, 2024.
- Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.
- Thijs Bol, Mathijs De Vaan, and Arnout van de Rijt. The matthew effect in science funding. *Proceedings of the National Academy of Sciences*, 115(19):4887–4890, 2018.
- Michael B Cohen, Yin Tat Lee, and Zhao Song. Solving linear programs in the current matrix multiplication time. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 938–942, 2019.
- Michael B Cohen, Yin Tat Lee, and Zhao Song. Solving linear programs in the current matrix multiplication time. *Journal of the ACM (JACM)*, 68(1):1–39, 2021.
- Michael D Ekstrand, Anubrata Das, Robin Burke, Fernando Diaz, et al. Fairness in information access systems. *Foundations and Trends® in Information Retrieval*, 16(1-2):1–177, 2022.
- Faisal R Elali and Leena N Rachid. Ai-generated research paper fabrication and plagiarism in the scientific community. *Patterns*, 4(3), 2023.
- Nissim Francez. *Fairness*. Springer Science & Business Media, 2012.
- Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K Ahmed. Bias and fairness in large language models: A survey. *Computational Linguistics*, pp. 1–79, 2024.
- Michael R Garey and David S Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*, volume 174. freeman San Francisco, 1979.
- Yingqiang Ge, Shuchang Liu, Ruoyuan Gao, Yikun Xian, Yunqi Li, Xiangyu Zhao, Changhua Pei, Fei Sun, Junfeng Ge, Wenwu Ou, et al. Towards long-term fairness in recommendation. In *Proceedings of the 14th ACM international conference on web search and data mining*, pp. 445–453, 2021.

- Stephen W Gilliland. The perceived fairness of selection systems: An organizational justice perspective. *Academy of management review*, 18(4):694–734, 1993.
- Yue Guo, Yi Yang, and Ahmed Abbasi. Auto-debias: Debiasing masked language models with automated biased prompts. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1012–1023, 2022.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022.
- Tom Jefferson, Philip Alderson, Elizabeth Wager, and Frank Davidoff. Effects of editorial peer review: a systematic review. *Jama*, 287(21):2784–2786, 2002.
- Shunhua Jiang, Zhao Song, Omri Weinstein, and Hengjie Zhang. A faster algorithm for solving general lps. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 823–832, 2021.
- Richard M Karp. Reducibility among combinatorial problems. *Complexity of Computer Computations*, pp. 85–103, 1972.
- Michael R King and ChatGPT. A conversation on artificial intelligence, chatbots, and plagiarism in higher education. *Cellular and molecular bioengineering*, 16(1):1–2, 2023.
- Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in  $o(\text{vrnk})$  iterations and faster algorithms for maximum flow. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pp. 424–433. IEEE, 2014.
- Seth S Leopold. Duplicate submission and dual publication: what is so wrong with them?, 2013.
- Kevin Leyton-Brown, Yatin Nandwani, Hedayat Zarkoob, Chris Cameron, Neil Newman, Dinesh Raghu, et al. Matching papers and reviewers at large conferences. *Artificial Intelligence*, 331: 104119, 2024.
- Yunqi Li, Hanxiong Chen, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. User-oriented fairness in recommendation. In *Proceedings of the web conference 2021*, pp. 624–632, 2021a.
- Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, and Yongfeng Zhang. Towards personalized fairness based on causal notion. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1054–1063, 2021b.
- David G Luenberger, Yinyu Ye, et al. *Linear and nonlinear programming*, volume 2. Springer, 1984.
- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6):1–35, 2021.
- Alessandro B Melchiorre, Navid Rekabsaz, Emilia Parada-Cabaleiro, Stefan Brandl, Oleg Lesota, and Markus Schedl. Investigating gender fairness of recommendation algorithms in the music domain. *Information Processing & Management*, 58(5):102666, 2021.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.

- John Schulman, Barret Zoph, Christina Kim, Jacob Hilton, Jacob Menick, Jiayi Weng, Juan Felipe Ceron Uribe, Liam Fedus, Luke Metz, Michael Pokorny, et al. Chatgpt: Optimizing language models for dialogue. *OpenAI blog*, 2(4), 2022.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- Zhao Song and Zheng Yu. Oblivious sketching-based central path method for linear programming. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9835–9847, 2021.
- Stanford. Research and development - ai index, 2024. URL [https://aiindex.stanford.edu/wp-content/uploads/2024/04/HAI\\_AI-Index-Report-2024\\_Chapter1.pdf](https://aiindex.stanford.edu/wp-content/uploads/2024/04/HAI_AI-Index-Report-2024_Chapter1.pdf).
- WR Stone. Plagiarism, duplicate publication, and duplicate submission: They are all wrong! *IEEE Antennas and Propagation Magazine*, 45(4):47–49, 2003.
- Ye Sun, Fabio Caccioli, and Giacomo Livan. Ranking mobility and impact inequality in early academic careers. *Proceedings of the National Academy of Sciences*, 120(34):e2305196120, 2023.
- Jaime A Teixeira da Silva, Aceil Al-Khatib, Vedran Katavić, and Helmar Bornemann-Cimenti. Establishing sensible and practical guidelines for desk rejections. *Science and Engineering Ethics*, 24:1347–1365, 2018.
- Jonathan P Tennant. The state of the art in peer review. *FEMS Microbiology letters*, 365(19):fny204, 2018.
- Jan van den Brand, Yin Tat Lee, Aaron Sidford, and Zhao Song. Solving tall dense linear programs in nearly linear time. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 775–788, 2020.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Ruotong Wang, F Maxwell Harper, and Haiyi Zhu. Factors influencing perceived fairness in algorithmic decision-making: Algorithm outcomes, development procedures, and individual differences. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, pp. 1–14, 2020.
- Yang Wang, Benjamin F Jones, and Dashun Wang. Early-career setback and future career impact. *Nature communications*, 10(1):4331, 2019.
- Yilin Wang, Zeyuan Chen, Liangjun Zhong, Zheng Ding, Zhizhou Sha, and Zhuowen Tu. Dolphin: Diffusion layout transformers without autoencoder. *arXiv preprint arXiv:2310.16305*, 2023.
- Yilin Wang, Haiyang Xu, Xiang Zhang, Zeyuan Chen, Zhizhou Sha, Zirui Wang, and Zhuowen Tu. Omnicontrolnet: Dual-stage integration for conditional image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7436–7448, 2024a.
- Zirui Wang, Zhizhou Sha, Zheng Ding, Yilin Wang, and Zhuowen Tu. Tokencompose: Text-to-image diffusion with token-level supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8553–8564, 2024b.
- Haolun Wu, Bhaskar Mitra, Chen Ma, Fernando Diaz, and Xue Liu. Joint multisided exposure fairness for recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on research and development in information retrieval*, pp. 703–714, 2022.
- Yang Zhang, Fuli Feng, Xiangnan He, Tianxin Wei, Chonggang Song, Guohui Ling, and Yongdong Zhang. Causal intervention for leveraging popularity bias in recommendation. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pp. 11–20, 2021.

## APPENDIX

**Roadmap.** In Section A, we give the conventional desk-reject algorithm. In Section B, we supplement the missing proofs in Section 4. In Section B, we supplement the missing proofs in Section 4. In Section C, we present the missing proofs in Section 5. In Section D, we show additional case studies. In Section E, we provide the details related to conference submission limits.

## A CONVENTIONAL DESK-REJECT ALGORITHM

**Algorithm 2** Conventional Desk-Reject Algorithm

---

```

1: procedure DESKREJECT( $\mathcal{A}, \mathcal{P}, x$ )
2:   /* Initialize registered paper set for each author. */
3:   for  $i = 1 \rightarrow n$  do
4:      $R_i \leftarrow \emptyset$ 
5:   end for
6:   /* Initialize the subset of remaining papers. */
7:    $S \leftarrow \mathcal{P}$ 
8:   /* Process each paper in submission order.*/
9:   for  $j = 1 \rightarrow m$  do
10:    for  $i \in A_j$  do
11:      /* If author  $a_i$  has reached the submission limit, the paper will be rejected.*/
12:      if  $|R_i| \geq x$  then
13:         $S \leftarrow S \setminus \{p_j\}$ 
14:        break
15:      end if
16:    end for
17:    /* If paper  $p_j$  is not rejected, we add it to each co-author's registered paper set.*/
18:    if  $p_j \in S$  then
19:      for  $i \in A_j$  do
20:         $R_i \leftarrow R_i \cup \{j\}$ 
21:      end for
22:    end if
23:  end for
24:  return  $S$ 
25: end procedure

```

---

## B MISSING PROOFS IN SECTION 4

In this section, we provide the complete technical proofs for Theorem 4.3 in Section 4. In Section B.1, we first introduce key definitions that will be useful. To structure our analysis, we in the subsequent proofs. We then establish positive results for the cases where  $n \leq 2$  in Section B.2, followed by negative results for  $n \geq 3$  in Section B.3.

## B.1 BASIC DEFINITIONS

To systematically analyze the desk-rejection problem, we begin by classifying authors based on their submission behavior and their relationship to co-authors. This classification will help us organize and present the proofs in a more structured and readable manner.

**Definition B.1** (Author Categories). *For any author  $a_i \in \mathcal{A}$ , we define the following categories:*

- **Non-compliant:** An author  $a_i$  is non-compliant if they have submitted more than  $x$  papers, i.e.,  $|P_i| > x$ . Such authors exceed the submission limit and are subject to desk-rejection under the policy.
- **Vulnerable:** An author  $a_i$  is vulnerable if they have submitted no more than  $x$  papers ( $|P_i| \leq x$ ) but have at least one non-compliant co-author, i.e.,  $\exists k \in C_i$  such that  $|P_k| > x$ .

Although these authors comply with the submission limit, they are at risk of being unfairly penalized due to their co-authors' non-compliance.

- **Safe:** An author  $a_i$  is safe if they have submitted no more than  $x$  papers ( $|P_i| \leq x$ ) and all their co-authors are also compliant, i.e.,  $\forall k \in C_i, |P_k| \leq x$ . These authors are guaranteed to retain all their submissions, as neither they nor their co-authors violate the submission limit.

Next, we formalize the notion of achievability for the ideal desk-rejection system.

**Definition B.2** (Achievability). *Given a submission limit problem instance as defined in Definition 3.1:*

- **Positive result:** A problem instance is a positive result if there exists an algorithm that can achieve the ideal desk-rejection as defined in Definition 4.1.
- **Negative result:** A problem instance is a negative result if, under proper conditions, no algorithm can achieve the ideal desk-rejection as defined in Definition 4.1.

In the following sections, we will use these definitions to systematically prove the positive results for small numbers of authors ( $n \leq 2$ ) and the negative results for larger numbers of authors ( $n \geq 3$ ), which covers two cases in Theorem 4.3.

## B.2 POSITIVE RESULTS

In this subsection, we present two positive results that support the  $n \leq 2$  case in Theorem 4.3. We begin with the positive result for  $n = 1$  and any  $x \in \mathbb{N}_+$ .

**Lemma B.3** (Positive result for  $n = 1$  and any  $x \in \mathbb{N}_+$ , general case). *If the following conditions hold:*

- Let  $n = 1$  denote the number of authors as defined in Definition 3.1.
- Let  $x \in \mathbb{N}_+$  denote the maximum number of submissions allowed for each author in the conference.

Then, there exists an algorithm that achieves the ideal desk-rejection as defined in Definition 4.1.

*Proof.* We consider the three cases for the only author  $a_1$ : non-compliant, vulnerable, and safe, as defined in Definition B.1.

**Case 1: Non-compliant author.** If author  $a_1$  is non-compliant, we desk-reject ( $|P_1| - x$ ) papers. This ensures that exactly  $x$  papers remain, satisfying the ideal desk-rejection condition.

**Case 2: Vulnerable author.** Since  $n = 1$  and there is only one author, author  $a_1$  has no co-authors to make itself vulnerable. Therefore, this case cannot happen.

**Case 3: Safe author.** If author  $a_1$  is safe, no papers need to be rejected. The ideal desk-rejection condition is trivially satisfied.

In all possible cases, we can achieve the ideal desk-rejection. Thus, the proof is finished.  $\square$

To present the positive result for  $n = 2$  and any  $x \in \mathbb{N}_+$ , we first discuss a specific case where all authors are non-compliant.

**Lemma B.4** (Positive result for  $n = 2$  and any  $x \in \mathbb{N}_+$ , non-compliant author only case). *If the following conditions hold:*

- Let  $n = 1$  denote the number of authors as defined in Definition 3.1.
- All the authors are non-compliant authors as defined in Definition B.1.

- Let  $x \in \mathbb{N}_+$  denote the maximum number of submissions allowed for each author in the conference.

Then, there exists an algorithm that achieves the ideal desk-rejection as defined in Definition 4.1.

*Proof.* Let  $c \in \mathbb{N}$  denote the number of papers co-authored by both author  $a_1$  and author  $a_2$ . For  $i \in \{1, 2\}$ , let  $b_i \in \mathbb{N}$  denote the number of single-authored papers by author  $a_i$ .

We then have:

$$b_1 + c = |P_1|$$

and

$$b_2 + c = |P_2|.$$

**Case 1:**  $c \leq x$ . In this case, we have  $b_1 \geq |P_1| - x$  and  $b_2 \geq |P_2| - x$ . Since  $b_i$  represents the number of single-authored papers by author  $a_i$ , we can desk reject exactly  $(|P_i| - x)$  papers from author  $a_i$ .

**Case 2:**  $c > x$ . Here, we have  $b_1 < |P_1| - x$  and  $b_2 < |P_2| - x$ . We first desk reject all  $b_1$  single-authored papers from author  $a_1$  and all  $b_2$  single-authored papers from author  $a_2$ . Next, we desk reject  $(c - x)$  co-authored papers from both authors. This ensures that the remaining  $x$  papers are co-authored by both  $a_1$  and  $a_2$ . Thus, we have successfully rejected exactly  $(|P_i| - x)$  papers from each author  $a_i$ .

By combining the two cases above, the proof is complete.  $\square$

With the help of Lemma B.4, we now establish the positive result for  $n = 2$  and any  $x \in \mathbb{N}_+$ .

**Lemma B.5** (Positive result for  $n = 2$  and any  $x \in \mathbb{N}_+$ , general case). *If the following conditions hold:*

- Let  $n = 1$  denote the number of authors as defined in Definition 3.1.
- Let  $x \in \mathbb{N}_+$  denote the maximum number of submissions allowed for each author in the conference.

Then, there exists an algorithm that achieves the ideal desk-rejection as defined in Definition 4.1.

*Proof.* We consider two authors,  $a_1$  and  $a_2$ . Without loss of generality, we assume that  $a_1$  has at least as many papers as  $a_2$ , i.e.,  $|P_1| \geq |P_2|$ . By exhaustively enumerating all possible compositions of author types (i.e., non-compliant, vulnerable, or safe) for  $a_1$  and  $a_2$ , we observe that the vulnerable-safe composition is impossible. This is because a vulnerable author must co-author at least one paper with a non-compliant author. After excluding this case, we analyze the remaining possible scenarios as follows:

**Case 1: Both  $a_1$  and  $a_2$  are safe authors.** In this case, no papers need to be rejected, and the ideal desk-rejection trivially holds.

**Case 2:  $a_1$  is a non-compliant author and  $a_2$  is a safe author.** Since rejecting papers from  $a_1$  does not affect  $a_2$ 's submissions, we can simply reject  $(|P_1| - x)$  papers from  $a_1$  to achieve the ideal desk-rejection.

**Case 3:  $a_1$  is a non-compliant author and  $a_2$  is a vulnerable author.** By Definition B.1, we have  $|P_1| > x$  and  $|P_2| \leq x$ . Let  $c := |\{p_j \in S : p_j \in P_1, p_j \in P_2\}|$  denote the number of co-authored papers by  $a_1$  and  $a_2$ . From basic set theory, we know that  $c \leq |P_2|$ . Since  $|P_2| \leq x$ , it follows that  $c \leq x$ . Therefore, we have:

$$\underbrace{|P_1| - c}_{\text{Individual papers of } a_1} \geq \underbrace{|P_1| - x}_{\text{Excess papers of } a_1},$$

which implies that the number of individual papers authored solely by  $a_1$  exceeds the number of over-limit papers for  $a_1$ . Thus, we can first reject  $a_1$ 's individual papers without affecting  $a_2$ 's submissions, thereby achieving the desired ideal desk-rejection.



**Case 4: Both  $a_1$  and  $a_2$  are non-compliant authors.** This case directly follows from Lemma B.4. Combining all the cases above, we conclude that the ideal desk-rejection can always be achieved, which finishes the proof.  $\square$

### B.3 NEGATIVE RESULTS

In this subsection, we present two positive results that support the  $n \geq 3$  case in Theorem 4.3. We commence by showing the negative result for  $n = 3$  and  $x = 1$ .

**Lemma B.6** (Negative result for  $n = 3$  and  $x = 1$ ). *If the following conditions hold:*

- Let  $n = 3$  denote the number of authors as defined in Definition 3.1.
- Let  $x = 1$  denote the maximum number of submissions allowed for each author in the conference.

*Then, under proper conditions, no algorithm can achieve the ideal desk-rejection as defined in Definition 4.1.*

*Proof.* Let all the authors be non-compliant authors as defined in Definition B.1, and let the number of papers be  $m = 3$ . We suppose the three papers  $p_1, p_2$ , and  $p_3$  have the following authorship:

- Paper  $p_1$  is co-authored by  $a_1$  and  $a_2$ .
- Paper  $p_2$  is co-authored by  $a_1$  and  $a_3$ .
- Paper  $p_3$  is co-authored by  $a_2$  and  $a_3$ .

From the authors’ perspective, the relationships are as follows:

- Author  $a_1$  has papers  $p_1$  and  $p_2$ .
- Author  $a_2$  has papers  $p_1$  and  $p_3$ .
- Author  $a_3$  has papers  $p_2$  and  $p_3$ .

We enumerate all possible rejection plans and their outcomes in Table 2.

Table 2: Remaining number of papers for each author after desk rejection.

Rejected Papers	Author $a_1$	Author $a_2$	Author $a_3$
N/A	2	2	2
$p_1$	1	1	2
$p_2$	1	2	1
$p_3$	2	1	1
$p_1, p_2$	0	1	1
$p_1, p_3$	1	0	1
$p_2, p_3$	1	1	0
$p_1, p_2, p_3$	0	0	0

First, suppose we desk reject paper  $p_3$ . Then, authors  $a_2$  and  $a_3$  each have one paper remaining, but author  $a_1$  still has two papers. To satisfy the constraint  $x = 1$ , we must reject one of  $p_1$  or  $p_2$ .

If we reject  $p_1$ , author  $a_2$  is left with no papers, which is unfair. If we reject  $p_2$ , author  $a_3$  is left with no papers, which is also unfair.

Thus, no rejection plan satisfies the ideal desk rejection condition for all authors. This completes the proof.  $\square$

Next, we present the negative result for any  $n \geq 3$  and  $x = n - 2$ .

**Lemma B.7** (Negative result for any  $n \geq 3$  and  $x = n - 2$ ). *If the following conditions hold:*

- *Let  $n \geq 3$  denote the number of authors as defined in Definition 3.1.*
- *Let  $x = n - 2$  denote the maximum number of submissions allowed for each author in the conference.*

*Then, under proper conditions, no algorithm can achieve the ideal desk-rejection as defined in Definition 4.1.*

*Proof.* In this negative problem instance, we choose the number of papers to be the same as the number of authors, i.e.,  $m = n$ , and we assume all the  $n$  authors are non-compliant authors as defined in Definition B.1.

For each of the  $n$  papers  $p_i \in \mathcal{P}$ , we let  $i$ -th paper  $p_i$  contain  $n - 1$  authors, excluding only the  $i$ -th author  $a_i$ . Specifically, we have:

- The first paper  $p_1$  has authors  $a_2, a_3, \dots, a_n$ .
- The second paper  $p_2$  has authors  $a_1, a_3, a_4, \dots, a_n$ .
- $\dots$
- The  $(n - 1)$ -th paper has authors  $a_1, a_2, \dots, a_{n-2}, a_n$ .
- The  $n$ -th paper has authors  $a_1, a_2, \dots, a_{n-2}, a_{n-1}$ .

Since each author is allowed to submit at most  $x = n - 2$  papers, we must desk-reject at least two papers. We analyze the process of desk-rejecting these two papers step by step.

**Step 1: Desk-reject the first paper.**

Without loss of generality, we consider rejecting paper  $p_1$  first. After this operation, authors  $a_2, a_3, \dots, a_n$ , will have  $n - 2$  submitted papers, while author  $a_1$  will have  $n - 1$  submitted papers.

**Step 2: Desk-reject the second paper.**

Without loss of generality, we consider rejecting paper  $p_2$  next. After this operation, authors  $a_3, a_4, \dots, a_n$ , will have  $n - 3$  submitted papers, while author  $a_1$  and  $a_2$  will have  $n - 2$  submitted papers.

At this point, it is impossible for authors  $a_3, a_4, a_5, \dots, a_n$  to have exactly  $(n - 2)$  submitted papers. Therefore, no algorithm can achieve the ideal desk-rejection under the given conditions. This completes the proof.  $\square$

## C MISSING PROOFS IN SECTION 5

In this section, we first present the missing proofs for fairness metrics in Section C.1, and then present the supplementary proofs for the hardness of individual fairness optimization in Section C.2. Finally, we show the additional proofs for the group fairness optimization problem in Section C.3.

### C.1 FAIRNESS METRICS

We present the relationship between the fairness metrics.

**Proposition C.1** (Relationship of Fairness Metrics, formal version of Proposition 5.6 in Section 5.1). *For any solution  $S \subseteq \mathcal{P}$  for the submission limit problem in Definition 3.1, we have*

$$\zeta_{\text{group}}(S) \leq \zeta_{\text{ind}}(S).$$

*Proof.* By Definition 5.5, we have:

$$\begin{aligned}
\zeta_{\text{group}}(S) &= \frac{1}{n} \sum_{i \in [n]} c(a_i, S) \\
&\leq \frac{1}{n} \sum_{i \in [n]} \max_{i \in [n]} c(a_i, S) \\
&= \frac{1}{n} \cdot n \cdot \max_{i \in [n]} c(a_i, S) \\
&= \zeta_{\text{ind}}(S),
\end{aligned}$$

where the first equality directly follows from Definition 5.5, the second and the third inequality follow from basic algebra, and the last equality follows from Definition 5.4. Thus, we complete the proof.  $\square$

## C.2 HARDNESS OF INDIVIDUAL FAIRNESS-AWARE SUBMISSION LIMIT PROBLEM

Before proving the theoretical results in Section 5.2, we first introduce a useful fact that serves as a foundation for the subsequent proofs.

**Fact C.2.** *For each author  $a_i \in \mathcal{A}$ , the number of papers after desk-rejection (i.e.,  $|\{p_j \in S : a_i \in A_j\}|$ ) can be written as  $W_i^\top r$ .*

*Proof.* This simply follows from:

$$\begin{aligned}
W_i^\top r &= \sum_{j \in [m]} W_{i,j} \cdot r_j \\
&= |\{j \in [m] : W_{i,j} = 1, r_j = 1\}| \\
&= |\{p_j \in \mathcal{P} : a_i \in A_j, p_j \in S\}| \\
&= |\{p_j \in S : a_i \in A_j\}|,
\end{aligned}$$

where the first and the second equality follow from basic algebra and set theory, and the third and the fourth equality follow from Definition 3.1.  $\square$

With the help of the aforementioned fact, we now prove the equivalence of the matrix form for the individual fairness problem.

**Proposition C.3** (Matrix Form Equivalence for  $\zeta_{\text{ind}}$ , formal version of Proposition 5.10 in Section 5.2). *The individual fairness-aware submission limit problem in Definition 5.7 and the matrix form integer programming problem in Definition 5.9 are equivalent.*

*Proof.* In Definition 5.7, the paper set  $\mathcal{P}$  consists of  $m$  papers, each of which can either be maintained or desk-rejected. Thus, the subset of maintained papers,  $S$ , can be represented by a 0-1 vector  $r \in \{0, 1\}^m$ , where  $r_j = 1$  indicates that paper  $p_j$  is maintained, and  $r_j = 0$  indicates that it is desk-rejected. We now establish the equivalence of both the objective function and the constraints in these two formulations.

**Part 1: Optimization Objective.** We first consider the objective function  $\mathbf{1}_n^\top D^{-1} W r$  in Definition 5.9:

$$\begin{aligned}
\min_{r \in \{0,1\}^m} \|\mathbf{1}_n - D^{-1} W r\|_\infty &= \min_{r \in \{0,1\}^m} \max_{i \in [n]} (1 - (D^{-1} W r)_i) \\
&= \min_{r \in \{0,1\}^m} \max_{i \in [n]} (1 - (W_i^\top r)_i / D_{i,i}) \\
&= \min_{r \in \{0,1\}^m} \max_{i \in [n]} (1 - (W_i^\top r)_i / |P_i|) \\
&= \min_{r \in \{0,1\}^m} \max_{i \in [n]} (1 - |\{p_j \in S : a_i \in A_j\}| / |P_i|) \\
&= \min_{r \in \{0,1\}^m} \max_{i \in [n]} c(a_i, S)
\end{aligned}$$

$$= \min_{r \in \{0,1\}^m} \zeta_{\text{ind}}(S),$$

where the first equality follows from the definition of infinity norm, the second equality follows from basic algebra, the third equality follows from Definition 5.9, the fourth equality follows from Fact C.2, the fifth equality follows from Definition 5.1, and the last equality follows from Definition 5.4. By decoding  $r$  back into the paper subset  $S$ , we recover the original optimization objective in Definition 5.7.

**Part 2: Constraints.** The constraint in Definition 5.9 can be rewritten using basic algebra as:

$$W_i \cdot r \leq x, \quad \forall i \in [n].$$

By applying Fact C.2, we see that this constraint is equivalent to its counterpart in Definition 5.7.

Since both the objective function and constraints in Definition 5.7 and Definition 5.9 are equivalent, the proof is complete.  $\square$

To show the hardness of the individual fairness problem, we first present a classical set cover problem with well-established hardness.

**Definition C.4** (Set Cover Problem Karp (1972); Garey & Johnson (1979)). *The Set Cover problem is the following:*

- **Input:** A universe  $U = \{1, \dots, n\}$ , a family of sets  $\{S_1, \dots, S_m\} \subseteq 2^U$ , and a integer  $K > 0$ .
- **Question:** Is there a subfamily  $\{S_j : j \in J\}$  for some  $J \subseteq \{1, \dots, m\}$  and  $|J| \leq K$  that covers  $U$ , i.e.,  $\bigcup_{j \in J} S_j = U$ ?

**Lemma C.5** (Folklore Karp (1972); Garey & Johnson (1979)). *The Set Cover problem defined in Definition C.4 is NP-hard.*

Additionally, we also present a technical lemma which is useful for showing the hardness of the individual fairness problem.

**Lemma C.6.** *For any  $r \in \{0, 1\}^m$ , the following two statements are equivalent:*

- **Part 1.**  $\|\mathbf{1}_n - D^{-1}Wr\|_\infty \leq 1 - \frac{1}{\min_{i \in [n]} |P_i|}$ .
- **Part 2.**  $\min_{i \in [n]} (Wr)_i \geq 1$ .

*Proof.* We first show that Part 1 implies Part 2. Suppose that

$$\|\mathbf{1}_n - D^{-1}Wr\|_\infty \leq 1 - \frac{1}{\min_{i \in [n]} |P_i|}.$$

By the definition of the infinity norm, we have

$$1 - \frac{(Wr)_{i'}}{|P_{i'}|} \leq 1 - \frac{1}{\min_{i \in [n]} |P_i|}, \quad \forall i' \in [n].$$

Rearranging gives

$$(Wr)_{i'} \geq \frac{|P_{i'}|}{\min_{i \in [n]} |P_i|} \geq 1, \quad \forall i' \in [n].$$

Since for all  $i' \in [n]$ , we have  $(Wr)_{i'} \geq 1$ , we can conclude that  $\min_{i \in [n]} (Wr)_i \geq 1$ .

Now we show that that Part 2 implies Part 1. Suppose that  $\min_{i \in [n]} (Wr)_i \geq 1$ , then we have  $(Wr)_i \geq 1$  for all  $i \in [n]$ , which implies that for all  $i \in [n]$ ,

$$1 - \frac{(Wr)_i}{|P_i|} \leq 1 - \frac{1}{|P_i|} \leq 1 - \frac{1}{\max_{i' \in [n]} |P_{i'}|}.$$

Hence

$$\|\mathbf{1}_n - D^{-1}Wr\|_\infty \leq 1 - \frac{1}{\min_{i \in [n]} |P_i|}.$$

Thus the proof is complete.  $\square$

**Theorem C.7** (Hardness, formal version of Theorem 5.11 in Section 5.2). *The Individual Fairness-Aware Submission Limit Problem defined in Definition 5.7 is NP-hard.*

*Proof.* By Proposition 5.10, it suffices to reduce Set Cover problem to the integer optimization problem of the matrix form in Definition 5.9.

Given an instance of Set Cover, we build the matrix  $W \in \{0, 1\}^{n \times m}$  by defining  $W_{i,j} = 1$  if element  $i \in S_j$ , and 0 otherwise. Now set  $|P_i| = \sum_{j \in [m]} W_{i,j}$  for every row  $i \in [n]$ . Finally, we choose  $x = m$ . We reduce the Set Cover problem to the following optimization problem:

$$\begin{aligned} & \min_{r \in \{0,1\}^m} \|\mathbf{1}_n - D^{-1}Wr\|_\infty \\ & \text{s.t. } Wr \leq m\mathbf{1}_n, \\ & \|r\|_1 \leq K. \end{aligned}$$

Note that this problem is easier than the optimization problem defined in Definition 5.7. The constraint  $Wr \leq m\mathbf{1}_n$  is always satisfied, so we can drop it out. Now, it suffices to consider the decision problem:

$$\begin{aligned} & \text{Find } r \in \{0, 1\}^m \\ & \text{s.t. } \|\mathbf{1}_n - D^{-1}Wr\|_\infty \leq 1 - \frac{1}{\min_{i \in [n]} |P_i|}, \\ & \|r\|_1 \leq K. \end{aligned}$$

Note that  $\|\mathbf{1}_n - D^{-1}Wr\|_\infty \leq 1 - \frac{1}{\min_{i \in [n]} |P_i|}$  is equivalent to  $\min_{i \in [n]} (Wr)_i \geq 1$  by Lemma C.6.

Hence the problem is equivalent to

$$\text{Find } r \in \{0, 1\}^m \text{ s.t. } \min_{i \in [n]} (Wr)_i \geq 1 \text{ and } \|r\|_1 \leq K.$$

It is not hard to see that the Set Cover problem has a solution if and only if the above problem has a solution. Requiring  $\min_{i \in [n]} (Wr)_i > 1$  exactly means that each element  $i$  in the universe is covered by at least set  $S_j$ . The constraint  $\|r\|_1 \leq K$  means that the size of cover is at most  $K$ . In other words, there exists a subfamily of size at most  $K$  covering all elements if and only if there is an  $r \in \{0, 1\}^m$  with  $\min_{i \in [n]} (Wr)_i > 1$  and  $\|r\|_1 \leq K$ .

Therefore, by Lemma C.5, the individual fairness-aware submission limit problem is NP-hard.  $\square$

### C.3 GROUP FAIRNESS OPTIMIZATION

Now, we present the missing proofs on both matrix form equivalence and linear programming optimal solution equivalence for the group fairness optimization problem.

**Proposition C.8** (Matrix Form Equivalence for  $\zeta_{\text{group}}$ , formal version of Proposition 5.14 in Section 5.3). *The problem in Definition 5.12 and the problem in Definition 5.13 are equivalent.*

*Proof.* In Definition 5.12, there are  $m$  papers in  $\mathcal{P}$ , where each paper can either be maintained or rejected. Thus, we can encode the paper subset  $S$  using a binary vector  $r \in \{0, 1\}^m$ , where  $r_j = 1$  indicates that paper  $p_j$  is maintained, and  $r_j = 0$  indicates that it is desk-rejected. We now demonstrate that both the objective function and the constraints are equivalent.

**Part 1: Optimization Objective.** We first examine the objective function  $\mathbf{1}_n^\top D^{-1}Wr$  in Definition 5.13:

$$\begin{aligned} \mathbf{1}_n^\top D^{-1}Wr &= \sum_{i \in [n]} (D^{-1}Wr)_i \\ &= \sum_{i \in [n]} (W \cdot r)_i / |P_i| \\ &= \sum_{i \in [n]} (W_i^\top \cdot r) / |P_i| \end{aligned}$$

$$\begin{aligned}
&= \sum_{i \in [n]} \frac{|\{p_j \in S : a_i \in A_j\}|}{|P_i|} \\
&= \sum_{i \in [n]} (1 - c(a_i, S)),
\end{aligned}$$

where the first equality follows from basic algebra, the second follows from Definition 5.13, the third follows from matrix-vector multiplication, the fourth follows from Fact C.2, and the final equality follows from Definition 5.1. Consequently, the maximization problem in Definition 5.13 can be rewritten as:

$$\max_{r \in \{0,1\}^m} \sum_{i \in [n]} (1 - c(a_i, S)).$$

Since maximizing this objective is equivalent to minimizing  $\sum_{i \in [n]} c(a_i, S)$ , we can reformulate it as:

$$\min_{r \in \{0,1\}^m} \sum_{i \in [n]} c(a_i, S).$$

By decoding  $r$  back into the paper subset  $S$ , we recover the original optimization objective in Definition 5.12.

**Part 2: Constraints.** Since the constraint is identical to that in the individual fairness minimization problem in Definition 5.7, this result follows directly from Part 2 in the proof of Proposition C.3.

Since both the objective function and constraints in Definition 5.12 and Definition 5.13 are equivalent, the proof is complete.  $\square$

**Theorem C.9** (Optimal Solution Equivalence of the Relaxed Problem, formal version of Theorem 5.16 in Section 5.3). *The optimal solution of the relaxed problem in Definition 5.15 is equivalent to the optimal solution of the original problem in Definition 5.13.*

*Proof.* The problem in Definition 5.15 is a linear program since it has a linear objective function  $\mathbf{1}_n^\top D^{-1} W r$  and linear constraints: the box constraint  $r \in [0, 1]^m$  and a linear inequality constraint  $(W r)/x \leq \mathbf{1}_n$ .

Furthermore, the problem is convex because the objective function is linear, the constraint  $(W r)/x \leq \mathbf{1}_n$  is affine, and the feasible region defined by  $r \in [0, 1]^m$  is a convex set.

By the fundamental theorem of linear programming (see Page 23 of Luenberger et al. (1984)), the optimal solution must occur at an extreme point of the convex polytope defined by the constraints. This implies that for all  $i \in [m]$ , we must have either  $r_i = 0$  or  $r_i = 1$ . Therefore, the optimal solution of the relaxed linear program coincides with that of the original integer program, which finishes the proof.  $\square$

## D ADDITIONAL CASE STUDIES

As discussed in Section 4.2, optimizing the individual fairness metric is computationally challenging. Therefore, we minimize the group fairness metric, which serves as a lower bound for individual fairness, as a practical alternative. In this subsection, we present case studies demonstrating the relationship between both types of fairness metrics.

**Example D.1.** *Consider a submission limit problem as defined in Definition 3.1 with  $x = 2$ ,  $n = 3$  authors, and  $m = 6$  papers. Let author  $a_1$  submit four papers  $p_1, p_2, p_3, p_4$ , author  $a_2$  submit two papers  $p_3, p_5$ , and author  $a_3$  submit two papers  $p_4, p_6$ .*

In this case, the ideal desk-rejection criteria in Definition 4.1 reject  $p_1$  and  $p_2$  (i.e.,  $S = \{p_3, p_4, p_5, p_6\}$ ), yielding fairness metrics  $\zeta_{\text{ind}}(S) = \max\{1/2, 0, 0\} = 1/2$  and  $\zeta_{\text{group}}(S) = \frac{1}{3}(1/2 + 0 + 0) = 1/6$ . By applying an LP solver to minimize group fairness using Algorithm 1 and enumerating all rejection strategies to verify individual fairness minimization, we observe that

minimizing group fairness in this case aligns with minimizing individual fairness as defined in Definition 5.9. This case illustrates that minimizing group fairness can sometimes benefit individual fairness.

However, group fairness and individual fairness are not always consistent. In some cases, prioritizing group fairness may disproportionately burden certain individuals. To illustrate this, we consider the following example.

**Example D.2.** Consider a submission limit problem as defined in Definition 3.1 with  $x = 2$ ,  $n = 5$  authors, and  $m = 4$  papers. Let author  $a_1$  submit four papers  $p_1, p_2, p_3, p_4$ , author  $a_2$  submit two papers  $p_1, p_2$ , and authors  $a_3, a_4, a_5$  be coauthors of papers  $p_3, p_4$ .

In this scenario, an ideal desk-rejection is impossible because  $a_1$  must have two papers rejected, but rejecting any papers would cause at least one of the authors in  $a_2, \dots, a_5$  to fall below the submission limit of  $x = 2$ . Here, group fairness and individual fairness diverge: Algorithm 1 minimizes group fairness by rejecting  $p_1$  and  $p_2$  (i.e.,  $S = \{p_3, p_4\}$ ), which unfairly excludes all of  $a_2$ 's submissions. This results in fairness metrics  $\zeta_{\text{group}}(S) = \frac{1}{4}(1/2 + 1 + 0 + 0) = 3/8$  and  $\zeta_{\text{ind}}(S) = \max\{1/2, 1, 0, 0\} = 1$ .

Conversely, the individual fairness minimization problem in Definition 5.9 rejects one paper from  $a_1, a_2$  and another from  $a_3, a_4$ , leading to  $\zeta_{\text{group}}(S) = \frac{1}{4}(1/2 + 1/2 + 1/2 + 1/2) = 1/2$  and  $\zeta_{\text{ind}}(S) = \max\{1/2, 1/2, 1/2, 1/2\} = 1/2$ .

This example highlights an unintended consequence of minimizing group fairness: it may unfairly penalize authors with fewer coauthors, as rejecting their papers incurs a smaller total cost. On the other hand, optimizing individual fairness inevitably spreads rejections across a broader set of authors, potentially leading to a higher overall fairness cost. Balancing individual and group fairness remains an open challenge, which we leave for future work.

## E SUMMARY OF CONFERENCE LINKS

In the introduction, Table 1 only gives a brief summary of the conference year and its limitation of per-author submission. Thus, we provide a detailed list of conferences in each year in this section, and then summarize the submission limits in Table. 3.

- CVPR
  - 2025, <https://cvpr.thecvf.com/Conferences/2025/CVPRChanges>
  - 2024, <https://cvpr.thecvf.com/Conferences/2024/AuthorGuidelines>
- ICCV
  - 2025, <https://iccv.thecvf.com/Conferences/2025/AuthorGuidelines>
  - 2023, <https://iccv2023.thecvf.com/policies-361500-2-20-15.php>
- AAAI
  - 2025, <https://aaai.org/conference/aaai/aaai-25/submission-instructions/>
  - 2024, <https://aaai.org/aaai-24-conference/submission-instructions/>
  - 2023, <https://aaai-23.aaai.org/submission-guidelines/>
  - 2022, <https://aaai.org/conference/aaai/aaai-22/>
- WSDM
  - 2025, <https://www.wsdm-conference.org/2025/call-for-papers/>
  - 2024, <https://www.wsdm-conference.org/2024/call-for-papers/>

- 2023, <https://www.wsdm-conference.org/2023/calls/call-papers/>
- 2022, <https://www.wsdm-conference.org/2022/calls/>
- 2021, <https://www.wsdm-conference.org/2021/call-for-papers.php>
- 2020, <https://www.wsdm-conference.org/2020/call-for-papers.php>
- **IJCAI**
  - 2025, <https://2025.ijcai.org/call-for-papers-main-track/>
  - 2024, <https://ijcai24.org/call-for-papers/>
  - 2023, <https://ijcai-23.org/call-for-papers/>
  - 2022, <https://ijcai-22.org/calls-papers>
  - 2021, <https://ijcai-21.org/cfp/index.html>
  - 2020, <https://ijcai20.org/call-for-papers/index.html>
  - 2019, <https://www.ijcai19.org/call-for-papers.html>
  - 2018, <https://www.ijcai-18.org/cfp/index.html>
  - 2017, <https://ijcai-17.org/MainTrackCFP.html>
- **KDD**
  - 2025, <https://kdd2025.kdd.org/research-track-call-for-papers/>
  - 2024, <https://kdd2024.kdd.org/research-track-call-for-papers/>
  - 2023, <https://kdd.org/kdd2023/call-for-research-track-papers/index.html>



Table 3: In this table, we summarize the submission limits of top conferences in recent years. For details of each conference website, we refer the readers to Section E in Appendix.

<b>Conference Name</b>	<b>Year</b>	<b>Upper Bound</b>
CVPR	2025	25
CVPR	2024	N/A
ICCV	2025	25
ICCV	2023	N/A
AAAI	2025	10
AAAI	2024	10
AAAI	2023	10
AAAI	2022	N/A
WSDM	2025	10
WSDM	2024	10
WSDM	2023	10
WSDM	2022	10
WSDM	2021	10
WSDM	2020	N/A
IJCAI	2025	8
IJCAI	2024	8
IJCAI	2023	8
IJCAI	2022	8
IJCAI	2021	8
IJCAI	2020	6
IJCAI	2019	10
IJCAI	2018	10
IJCAI	2017	N/A
KDD	2025	7
KDD	2024	7
KDD	2023	N/A