Learning Fair and Preferable Allocations through Neural Network

Extended Abstract

Ryota Maruo
Kyoto University
Kyoto, Japan
mryota@ml.ist.i.kyoto-u.ac.jp

Koh Takeuchi
Kyoto University
Kyoto, Japan
takeuchi@i.kyoto-u.ac.jp

Hisashi Kashima
Kyoto University
Kyoto, Japan
kashima@i.kyoto-u.ac.jp

ABSTRACT

The fair allocation of indivisible resources is a fundamental problem. Existing research has developed various allocation mechanisms or algorithms to satisfy different fairness notions. For example, round robin (RR) was proposed to meet the fairness criterion known as envy-freeness up to one good (EF1). Expert algorithms without mathematical formulations are used in real-world resource allocation problems to find preferable outcomes for users. Therefore, we aim to design mechanisms that strictly satisfy good properties with replicating expert knowledge. However, this problem is challenging because such heuristic rules are often difficult to formalize mathematically, complicating their integration into theoretical frameworks. Additionally, formal algorithms struggle to find preferable outcomes, and directly replicating these implicit rules can result in unfair allocations because human decision-making can introduce biases. In this paper, we aim to learn implicit allocation mechanisms from examples while strictly satisfying fairness constraints, specifically focusing on learning EF1 allocation mechanisms through supervised learning on examples of reported valuations and corresponding allocation outcomes produced by implicit rules. To address this, we developed a neural RR (NRR), a novel neural network that parameterizes RR. NRR is built from a differentiable relaxation of RR and can be trained to learn the agent ordering used for RR. We conducted experiments to learn EF1 allocation mechanisms from examples, demonstrating that our method outperforms baselines in terms of the proximity of predicted allocations and other metrics.

KEYWORDS

Fair Division; Deep Learning Architecture; Automated Mechanism Design

ACM Reference Format:
Ryota Maruo, Koh Takeuchi, and Hisashi Kashima. 2025. Learning Fair and Preferable Allocations through Neural Network: Extended Abstract. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

1 INTRODUCTION

The fair allocation of indivisible resources is a fundamental problem in both computer science and economics [1–3]. Research has predominantly focused on the fair division of *goods*, where n agents assign non-negative values on m indivisible items [6]. An example of the fair division of goods is course assignment [13].

Numerous allocation mechanisms or algorithms have been proposed for various fairness concepts. For example, the round robin (RR) mechanism [5] has been developed to find allocations that satisfy the fairness criterion called envy-free up to one good (EF1) [4, 10]. In this algorithm, an order of agents is defined, and each agent, in turn, selects their most preferred item from the remaining items.

Expert algorithms without mathematical formulations are used in real-world resource allocation problems to find preferable outcomes for users. Although their goodness is not formally proven, these algorithms can use implicit or empirical knowledge in various domains [8]. Therefore, we aim to design mechanisms that replicate expert knowledge, while strictly satisfying fairness to prevent biases in human judgments leading to unfair allocations [7, 9].

In this paper, we study learning implicit allocation rules from examples while strictly satisfying EF1 by extending the framework of Narasimhan et al. [12]. Given example pairs of reported valuations and allocations based on implicit rules, our goal is to train a parameterized EF1 mechanism by capturing the relationship between inputs (valuations) and outputs (allocations). To implement this approach, we introduced two novel techniques. First, we proposed a soft RR (SoftRR) algorithm that makes the discrete procedure of RR differentiable, enabling it to be used for back-propagation. Second, we constructed a novel neural network called a neural RR (NRR). We conducted experiments to learn EF1 allocations through examples, and confirmed that NRR outperforms the baselines in terms of the proximity of predicted allocations and other metrics.

2 PROBLEM SETTING

Fair Division. We consider the fair allocation of a set of indivisible goods $[m] = \{1, 2, ..., m\}$ to a set of agents $[n] = \{1, 2, \dots, n\}$. Each agent i has additive utilities, defined as $v_i(S) := \sum_{j \in S} v_{ij}$ for $S \subseteq [m]$. We denote a valuation profile by a matrix $\boldsymbol{V} := (v_{ij})_{i \in [n], j \in [m]} \in \mathbb{R}_{\geq 0}^{n \times m}$. An allocation $(A_i)_{i \in [n]}$ is denoted by a matrix $\boldsymbol{A} \in \{0, 1\}^{n \times m}$, where $A_{ij} = 1$ means i gets the j-th good. An allocation is EF1 if, for all $i, j \in [n]$, either $v_i(A_i) \geq v_i(A_j)$, or there exists a good $o \in A_j$ such that $v_i(A_i) \geq v_i(A_j \setminus \{o\})$. An allocation mechanism is denoted by $f : \boldsymbol{V} \mapsto \boldsymbol{A}$. f is said to be EF1 if f always outputs EF1 allocations for any profiles $\boldsymbol{V}$.

Learning Problem. Given an implicit allocation mechanism g, our goal is to find an allocation mechanism that approximates g, subject to EF1 constraint. We assume access to g through a dataset $S := \{(\boldsymbol{V}^1, \boldsymbol{A}^1), \dots, (\boldsymbol{V}^L, \boldsymbol{A}^L)\}$, where $\boldsymbol{V}^1, \dots, \boldsymbol{V}^L$ are examples of valuation profiles sampled from an unknown distribution over the set of all valuation profiles, and $\boldsymbol{A}^1 = g(\boldsymbol{V}^1), \dots, \boldsymbol{A}^L = g(\boldsymbol{V}^L)$ are

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

The full version is available at https://arxiv.org/abs/2410.17500

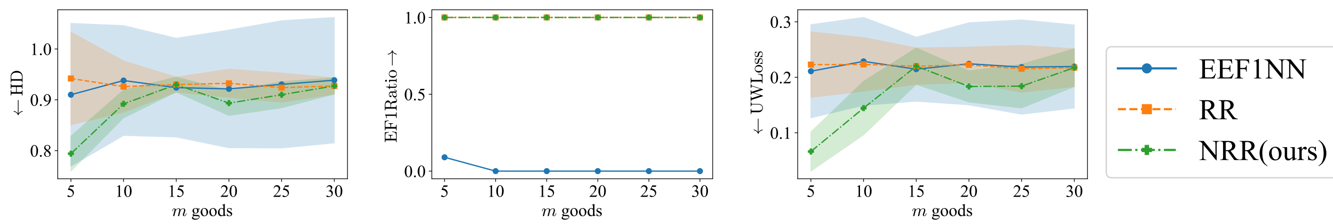

Figure 1: Evaluation metrics for varying numbers of goods. The horizontal axis represents the number of goods m. The vertical axes in each figure correspond to the following metrics: Hamming distance (leftmost), ratio of EF1 allocations (middle), and utilitarian welfare loss (rightmost). The symbols ↓ and ↑ indicate that the metric improves as the value decreases and increases, respectively.

Algorithm 1 Differentiable Relaxation of One Round

Input: A valuation profile $\boldsymbol{V} \in \mathbb{R}^{n\times m}_{\geq 0}$.
Output: A matrix $\boldsymbol{R} \in \mathbb{R}^{n\times m}$.

function $sr_\tau(\boldsymbol{V})$
  $\boldsymbol{R} \leftarrow \boldsymbol{O}_{n,m}$
  $\boldsymbol{c} \leftarrow \boldsymbol{1}_m$
  for $i = 1, \dots, n$ **do**
    $\boldsymbol{y} \leftarrow \mathrm{softmax}((\boldsymbol{V}[i] - \min(\boldsymbol{V}[i]) \cdot \boldsymbol{1} + 1) \odot \boldsymbol{c}/\tau)$
    $\boldsymbol{c} \leftarrow (\boldsymbol{1} - \boldsymbol{y}) \odot \boldsymbol{c}$
    $\boldsymbol{R}[i] \leftarrow \boldsymbol{y}$
  end for
  return $\boldsymbol{R}$
end function

Algorithm 2 SoftRR$_\tau$

Input: A valuation profile $\boldsymbol{V} \in \mathbb{R}^{n\times m}_{\geq 0}$
Output: A matrix $\boldsymbol{R} \in \mathbb{R}^{n\times m}$

$k \leftarrow \lceil m/n \rceil$
$\boldsymbol{V}_{\mathrm{rep}} \leftarrow \mathrm{repeat}(\boldsymbol{V}, k)$
$\boldsymbol{R}_{\mathrm{rep}} \leftarrow sr_\tau(\boldsymbol{V}_{\mathrm{rep}})$
Split $\boldsymbol{R}_{\mathrm{rep}}$ into k matrices: $\boldsymbol{R}_1 \leftarrow \boldsymbol{R}_{\mathrm{rep}}[1 : n, 1 : m], \boldsymbol{R}_2 \leftarrow \boldsymbol{R}_{\mathrm{rep}}[n+1 : 2n, 1 : m], \dots, \boldsymbol{R}_k \leftarrow \boldsymbol{R}_{\mathrm{rep}}[(k-1)n+1 : kn, 1 : m]$
$\boldsymbol{R} \leftarrow \sum_{r=1}^{k} \boldsymbol{R}_r$
return $\boldsymbol{R}$

the corresponding allocation outcomes determined by g. Given the dataset S, our goal is to find the EF1 allocation mechanism that best approximates g:

$$f_{\boldsymbol{\theta}^*} := \operatorname*{argmin}_{\boldsymbol{\theta}\in\Theta} \sum_{r=1}^{L} d(\boldsymbol{A}^r, f_{\boldsymbol{\theta}}(\boldsymbol{V}^r)), \tag{1}$$

where $d(\boldsymbol{A}, \boldsymbol{A}')$ is a function that calculates the discrepancy between two allocation outcomes $\boldsymbol{A}$ and $\boldsymbol{A}'$. We search for an EF1 mechanism over a parameterized subset of all EF1 mechanisms, denoted by $\mathcal{F} := \{f_{\boldsymbol{\theta}} \mid \boldsymbol{\theta} \in \Theta\} \subset \mathcal{F}_{\mathrm{EF1}}$ where $\boldsymbol{\theta}$ is a parameter from the parameter space Θ.

3 PROPOSED METHOD

To solve the above problem in Equation (1), we propose a parameterized family of mechanisms $\mathcal{F}$ based on RR. RR's output depends on the order of agents, and so we propose modeling $f_{\boldsymbol{\theta}}$ in Equation (1) through a neural network with a learnable parameter $\boldsymbol{\theta}$ that models the agent order. This network has a differentiable relaxation of RR and a sub-network to parametrize the agent order.

We first present the differentiable relaxation of RR in Algorithm 2. This algorithm uses the subroutine sr_τ from Algorithm 1, which simulates a single round by the softmax function with the temperature parameter τ. SoftRR$_\tau$ executes sr_τ on the repeated profile repeat$(\boldsymbol{V}, k)$ where $k = \lceil m/n \rceil$ to simulate multiple rounds.

We then construct the neural network that models $f_{\boldsymbol{\theta}}$. In this model, a sub-network first computes a permutation matrix $\hat{\boldsymbol{P}}$ that represents the agent order from the valuation profile $\boldsymbol{V}$. Then, we multiply $\hat{\boldsymbol{P}}$ by the input valuation to reorder the agents. Finally, we compute $\hat{\boldsymbol{P}}^\top \mathrm{SoftRR}(\hat{\boldsymbol{P}}\boldsymbol{V})$ and normalize it to output an allocation.

4 EXPERIMENTS

We conducted experiments to learn EF1 allocation mechanisms from examples. We synthesized datasets for good allocations by modeling i's valuation of good j as $v_{ij} = \mu_i + \varepsilon_{i,j}$, where $\mu_i \sim U[1, 2]$ and $\varepsilon_{i,j} \sim U[0, 0.01]$. We used the maximum utilitarian welfare (MUW) rule as an implicit allocation rule: $\mathrm{MUW}(\boldsymbol{V}) := \mathrm{argmax}_{\boldsymbol{A}} \sum_{i=1}^{n} v_i(A_i)$. We set $n = 15$ and tested on different values of m.

We compared our method by the original RR and EEF1NN, a neural network proposed by Mishra et al. [11]. The evaluation metrics included Hamming distance (HD) between predicted and correct allocations, the ratio of the number of EF1 allocations (EF1Ratio), and the utilitarian welfare loss (UWLoss) defined as $\mathrm{UWLoss}(\boldsymbol{V}, \hat{\boldsymbol{A}}) := 1 - (\sum_{i=1}^{n} v_i(\hat{A}_i))/\mathrm{MUW}(\boldsymbol{V})$.

We show experimental results in Figure 1. NRR yielded a lower HD as well as UWLoss compared to RR and EEF1NN. RR and NRR are EF1 allocation mechanisms by construction, and EF1Ratio remained 1.0 while EEF1NN failed to output EF1 allocations.

5 CONCLUSION

We studied learning EF1 allocation mechanisms through examples based on implicit rules. We first developed SoftRR for differentiable relaxation of RR, and proposed a neural network called NRR based on SoftRR. Experimental results show that our architecture can learn implicit rules by optimizing agent orders.

ACKNOWLEDGMENTS

KT is supported by JST FOREST Program, Grant Number JPMJFR232S.

REFERENCES

[1] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. 2023. Fair Division of Indivisible Goods: Recent Progress and Open Questions. *Artificial Intelligence* 322 (2023), 103965.

[2] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. 2019. Fair Allocation of Indivisible Goods and Chores. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, 53–59.

[3] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

[4] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.

[5] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. Association for Computing Machinery, New York, NY, USA, 305–322.

[6] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. 2017. Fair Public Decision Making. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. Association for Computing Machinery, New York, NY, USA, 629–646.

[7] Tom Gordon-Hecker, Shoham Choshen-Hillel, Shaul Shalvi, and Yoella Bereby-Meyer. 2017. *Resource Allocation Decisions: When Do We Sacrifice Efficiency in the Name of Equity?* Springer International Publishing, Chapter 6, 93–105.

[8] Ehud Gudes, Tsvi Kuflik, and Amnon Meisels. 1990. On Resource Allocation by an Expert System. *Engineering Applications of Artificial Intelligence* 3, 2 (1990), 101–109.

[9] Meng Li, Jeffrey Vietri, Alison P. Galvani, and Gretchen B. Chapman. 2010. How Do People Value Life? *Psychological Science* 21, 2 (2010), 163–167.

[10] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*. Association for Computing Machinery, New York, NY, USA, 125–131.

[11] Shaily Mishra, Manisha Padala, and Sujit Gujar. 2022. EEF1-NN: Efficient and EF1 Allocations Through Neural Networks. In *Pacific Rim International Conference on Artificial Intelligence 2022: Trends in Artificial Intelligence*. Springer Nature Switzerland, Cham, 388–401.

[12] Harikrishna Narasimhan, Shivani Agarwal, and David C. Parkes. 2016. Automated Mechanism Design without Money via Machine Learning. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. 433–439.

[13] Abraham Othman, Tuomas Sandholm, and Eric Budish. 2010. Finding Approximate Competitive Equilibria: Efficient and Fair Course Allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 873–880.