



Title	Deep False-Name-Proof Auction Mechanisms
Author(s)	Sakurai, Yuko; Oyama, Satoshi; Guo, Mingyu et al.
Citation	PRIMA 2019: Principles and Practice of Multi-Agent Systems, 11873, 594-601 https://doi.org/10.1007/978-3-030-33792-6_45
Issue Date	2019-10-21
Doc URL	https://hdl.handle.net/2115/88721
Rights	This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature' s AM terms of use, but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: http://dx.doi.org/10.1007/978-3-030-33792-6_45
Type	journal article
File Information	sakurai-prima2019.pdf



Deep False-name-proof Auction Mechanisms

Yuko Sakurai¹, Satoshi Oyama², Mingyu Guo³ and Makoto Yokoo⁴

¹ National Institute of Advanced Industrial Science and Technology

yuko.sakurai@aist.go.jp

² Hokkaido University/RIKEN

oyama@ist.hokudai.ac.jp

³ University of Adelaide

mingyu.guo@adelaide.edu.au

⁴ Kyushu University/RIKEN

yokoo@inf.kyuhu-u.ac.jp

Abstract. We explore an approach to designing false-name-proof auction mechanisms using deep learning. While multi-agent systems researchers have recently proposed data-driven approaches to automatically designing auction mechanisms through deep learning, false-name-proofness, which generalizes strategy-proofness by assuming that a bidder can submit multiple bids under fictitious identifiers, has not been taken into account as a property that a mechanism has to satisfy. We extend the RegretNet neural network architecture to incorporate false-name-proof constraints and then conduct experiments demonstrating that the generated mechanisms satisfy false-name-proofness.

Keywords: Mechanism design · Deep learning · False-name-proofness

1 Introduction

Mechanism design, a subfield of microeconomic theory and game theory, focuses on designing mechanisms that result in desirable outcomes even if the agents act strategically. One desirable property that mechanisms have to satisfy is *strategy-proofness*: for a bidder, declaring her true valuation is a dominant strategy, i.e., an optimal strategy regardless of the other bidders' actions. The Vickrey-Clarke-Groves (VCG) mechanism is well-known to be a strategy-proof mechanism that can be applied to combinatorial auctions, in which multiple items are simultaneously offered, and a bidder can bid on any bundle of items. In the VCG mechanism, an allocation is determined that maximizes the social surplus, i.e., the sum of all participants' utilities including that of the auctioneer. A winner pays the smallest amount she would have had to bid to win her bundle of items.

The problem is, in anonymous settings such as the Internet, a bidder can pretend to be multiple bidders. We refer to such a manipulation as *false-name bidding*. False-name bids are bids submitted under fictitious identifiers, e.g., multiple e-mail addresses. It is difficult to detect false-name bids since identifying each participant on the Internet is virtually impossible. We say a mechanism is *false-name-proof* if, for each bidder, declaring her true valuations by using a

single identifier is a dominant strategy. Unfortunately, Yokoo *et al.* [9] showed that the VCG mechanism is not false-name-proof and that no false-name-proof mechanism satisfies Pareto efficiency. Thus, several false-name-proof mechanisms have been proposed [6, 8].

We consider the design of *false-name-proof* auctions through deep learning. Several multi-agent systems researchers recently used deep learning in the automated design of optimal auction mechanisms [3, 4, 7]. Conitzer and Sandholm [1, 2] introduced the automated mechanism design (AMD) approach in which the problem of finding a mechanism to satisfy desirable properties is formulated as a linear program. However, Guo and Conitzer [5] showed that the AMD approach does not have sufficient scalability in terms of memory requirement and computational time. Thus, methods based on the AMD approach apply limited and specialized problem settings with a small number of agents and items. To overcome the scalability problem, Dütting *et al.* [3] recently proposed a data-driven approach to using deep neural networks called the *RegretNet framework* for the AMD problem of optimal auctions to maximize the expected revenue.

We have extended the RegretNet framework to incorporate false-name-proofness into designing combinatorial auctions that maximize the expected revenue. As far as the authors know, this is the first attempt to use machine learning for the design of false-name-proof auction mechanisms. Many of the existing manually designed mechanisms have been criticized for their relatively low revenue. It is thus important to examine how much revenue the machine-learning generated mechanisms can attain. In our experiments, we generated mechanisms for two problem settings. We found that when bidders' valuations are limited, the generated mechanism is closely similar to the Adaptive Reserve Price mechanism [6].

2 Preliminaries

2.1 Model

Let $N = \{1, 2, \dots, n\}$ be the set of bidders and let $M = \{1, 2, \dots, m\}$ be the set of items. A bidder $i \in N$ has a valuation function $v_i : 2^M \rightarrow \mathbb{R}_+$; i.e., $v_i(B)$ denotes bidder i 's valuation for a bundle of items $B \subseteq M$. V_i denotes the space of a possible valuation function for bidder i . $v = (v_1, \dots, v_n)$ denotes a profile of valuations, and $v_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$ denotes the profile of valuations except for bidder i . We assume that a valuation function v_i normalized by $v_i(\emptyset) = 0$ satisfies *free disposal*, i.e., $v_i(B') \geq v_i(B)$ for all $B' \supseteq B$. We also assume that each bidder is single minded; i.e., she has at most one minimal bundle with a positive value. Here, minimal bundle B for bidder i with $v_i(B)$ satisfies $v_i(B') < v_i(B)$ for $\forall B' \subset B$. Bidder i 's valuation function v_i is drawn independently from distribution F_i . We assume that an auctioneer knows the distributions $F = (F_1, \dots, F_n)$.

Each bidder reports her bid $b_i(B)$ for any bundle of items $B \subseteq M$. $v_i(B) = b_i(B)$ is not guaranteed since a bidder might report her bid b_i untruthfully. Let $b = (b_1, \dots, b_n)$ be the profile of bids and $b_{-i} = (b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$ be the

profile of bids except for bidder i . We consider a randomized mechanism for a combinatorial auction. A combinatorial auction mechanism $\mathcal{M}(a, p)$ consists of a randomized allocation rule a and a payment rule p . When a set of n bidders participates, the randomized allocation rule is defined as $a : \mathbb{R}^{nm} \rightarrow [0, 1]^{nm}$, and the payment rule is defined as $p : \mathbb{R}^{nm} \rightarrow \mathbb{R}_+^n$. $a_i(B) \in [0, 1]$ denotes the probability that bidder i obtains bundle B and $p_i(B)$ is bidder i 's payment for bundle B .

To satisfy the allocation feasibility requirement, the following conditions must be satisfied: (1) the probability that item $j \in M$ is allocated to a set of bidders N is at most 1 and (2) the total allocation to agent $i \in N$ is at most 1.

$$\sum_{i \in N} \sum_{B \subseteq M: j \in B} a_i(B) \leq 1, \forall j \in M \quad (1)$$

$$\sum_{B \subseteq M} a_i(B) \leq 1, \forall i \in N \quad (2)$$

The expected utility of bidder i with valuation function v_i is given by

$$u_i(v_i, b) = \sum_{B \subseteq M} v_i(B) \cdot a_i(B) - p_i(b). \quad (3)$$

Next, let us introduce three properties of mechanisms.

Strategy-Proofness (SP): A mechanism $\mathcal{M}(a, p)$ is strategy-proof if it maximizes a bidder's utility regardless of the other bidders' reports; i.e., $\forall i \in N, \forall b_i, \forall v_i, u_i(v_i, (v_i, b_{-i})) \geq u_i(v_i, (b_i, b_{-i}))$.

Individual Rationality (IR): A mechanism $\mathcal{M}(a, p)$ is *individually rational* if no bidder suffers any loss; i.e., $\forall N, \forall i \in N, \forall v_i, \forall b_{-i}, u_i(v_i, (v_i, b_{-i})) \geq 0$ holds.

False-Name-Proofness (FNP): A mechanism $\mathcal{M}(a, p)$ is false-name-proof if it maximizes a bidder's utility by reporting a true valuation function using a single identifier; i.e., if for all $k+1$ valuation functions of $v_i, b_{id_1}, \dots, b_{id_k}$ where b_{id_j} is a false-name bid and $k \leq n$, $u_i(v_i, (v_i, b_{-i})) \geq u_i(v_i, (b_{id_1}, \dots, b_{id_k}, b_{-i}))$. We assume that the number of false-name bids k is at most the number of items n . This is a reasonable assumption because false-name bids are made for obtaining items.

Example 1. Consider a combinatorial auction with two items. We denote $b_i = ((b_i(\{1\}), b_i(\{2\}), b_i(\{1, 2\})))$ and $a_i = ((a_i(\{1\}), a_i(\{2\}), a_i(\{1, 2\})))$.

Case 1: Bidders 1 and 2 submit bids $b_1 = (0, 0, 10)$ and $b_2 = (0, 0, 8.4)$, respectively. A mechanism $\mathcal{M}(a, p)$ outputs $a_1 = (0, 0, 1)$ and $a_2 = (0, 0, 0)$ as an allocation rule and $p_1 = 8.4$ and $p_2 = 0$ as a payment rule. If we assume that the bidders reported their valuations truthfully, the expected utility of bidder 1 is $10 \times 1 - 8.4 = 1.6$ and the expected utility of bidder 2 is 0.

Case 2: Bidders 1 and 2 submit bids $b_1 = (0, 0, 9)$ and $b_2 = (0, 0, 8.4)$, respectively. A mechanism $\mathcal{M}(a, p)$ outputs $a_1 = (0, 0, 0.9)$, $a_2 = (0, 0, 0)$, $p_1 = 8.4$, and $p_2 = 0$. If we assume that bidder 1 misreported her valuation and that her true valuation is her bid in Case 1, her expected utility is $10 \times 0.9 - 8.4 = 0.6 < 1.6$.

Case 3: Bidders 1, 2, and 3 submit $(5, 0, 5)$, $(0, 0, 8.4)$, and $(0, 5, 5)$, respectively.

A mechanism $\mathcal{M}(a, p)$ outputs $a_1 = (0.9, 0, 0)$, $a_2 = (0, 0, 0)$, $a_3 = (0, 0.9, 0)$, $p_1 = 4$, $p_2 = 0$, and $p_3 = 4$. If we assume that bidders 1 and 3 are false-name bids from bidder 1 in Case 1, the probability that bundle $\{1, 2\}$ is allocated to her is $a_1(\{1\}) \cdot a_2(\{2\}) + a_1(\{2\}) \cdot a_2(\{1\}) + a_1(\{1, 2\}) + a_2(\{1, 2\})$. Thus, her expected utility is $v_1(\{1, 2\}) \times 0.81 - (p_1 + p_2) = 10 \times 0.81 - 8 = 0.01 < 1.6$.

Although we show only three cases, we can say that mechanism $\mathcal{M}(a, p)$ satisfies SP and FNP if it is robust against all possible misreports and false-name manipulations.

2.2 Existing false-name-proof mechanisms

The existing false-name-proof combinatorial auction mechanisms were manually developed [6, 8, 9]. We introduce two representative mechanisms.

Minimal Bundle (MB) [8]: First, $B \subseteq M$ is allocated to bidder i , where B is a minimal bundle of i . Then, $B^* \subseteq M \setminus B$ is allocated to another bidder i' who has the highest remaining valuation, where B^* is a minimal bundle of i' , and so on. The payment for an allocated bundle B is equal to the highest valuation of another bidder for a bundle that is minimal and conflicting with B .

Adaptive Reserve Price (ARP) [6]: The basic idea of ARP is to base the reserve prices on the other bidders' bids. The reserve price on the set of all items is determined by doubling the second highest bid among ones for each single item. If a bidder makes the highest bid for the set of all items that exceeds this reserve price, she wins. Otherwise, the reserve price for single items with the highest and the second highest bids is set as half of the highest bid for the set of all items. If the highest or/and second highest bids for any single item exceed the reserve price, she/they win. No other items are allocated.

3 RegretNet Framework

The *RegretNet* framework proposed by Dütting *et al.* [3] comprises two separate networks for the allocation and payment rules. Both networks are simultaneously trained using samples from the value distribution by maximizing expected revenue subject to SP.

Let $(a^w, p^w) \in \mathcal{M}$ be an auction with parameters $w \in \mathbb{R}^d$ and some $d \in \mathbb{N}$. The *loss function* is defined as the negated expected revenue $\mathcal{L}(a, p) = -\mathbf{E}_{v \sim F}[\sum_{i \in N} p_i^w(v)]$. With the other bids fixed, the *expected ex post regret* of SP rgt_sp_i for bidder i is defined as the maximum excess in her utility, considering all possible misreports of her valuation functions: $\text{rgt_sp}_i(a^w, p^w) = \mathbf{E}[\max_{b_i \in V_i} u_i^w(v_i, (b_i, v_{-i})) - u_i^w(v_i, (v_i, v_{-i}))]$. An auction satisfies SP if and only if $\text{rgt_sp}_i(a^w, p^w) \leq 0$ for any $i \in N$. For IR, Dütting *et al.* incorporated the IR constraint in the networks.

In practice, $\mathcal{L}(a, p)$ and $\text{rgt_sp}_i(a^w, p^w)$ can be estimated from a sample of valuation profiles $S = \{v^{(1)}, \dots, v^{(L)}\}$ drawn independently from F . Thus, the learning problem is defined as

$$\min_{w \in \mathbb{R}^d} \widehat{\mathcal{L}}(a^w, p^w) \quad \text{s.t.} \quad \widehat{\text{rgt_sp}_i}(a^w, p^w) = 0, \quad \forall i \in N, \quad (4)$$

where

$$\widehat{\mathcal{L}}(a^w, p^w) = -\frac{1}{L} \sum_{l=1}^L \sum_{i=1}^n p_i^w(v^{(l)}), \quad (5)$$

$$\widehat{\text{rgt-sp}}_i(a^w, p^w) = \frac{1}{L} \sum_{l=1}^L \max_{b_i \in V_i} u_i^w(v_i^{(l)}, (b_i, v_{-i}^{(l)})) - u_i^w(v_i^{(l)}, v^{(l)}). \quad (6)$$

Dütting *et al.* used the *augmented Lagrangian method* to solve this learning problem. The Lagrangian function for the optimization problem with a strategy-proof constraint is defined as

$$\widehat{\mathcal{L}}(a^w, p^w) + \sum_{i \in N} \lambda_i \widehat{\text{rgt-sp}}_i(a^w, p^w) + \frac{\rho}{2} \left(\sum_{i \in N} \widehat{\text{rgt-sp}}_i(a^w, p^w) \right)^2, \quad (7)$$

where $\lambda \in \mathbb{R}^n$ is a vector of Lagrangian multipliers and $\rho > 0$ is a fixed parameter used to control the weight of the quadratic penalty.

4 Introducing false-name-proof constraints


In neural network training, the size of the input must be constant. This creates difficulties in introducing false-name-proof constraints since false-name bids change the number of bids an auction receives. If the maximum number of bids has already been received, a false-name constraint cannot be generated because more bids cannot be accepted.

To overcome this problem, we use subsets of the actual bids as *virtual bids* and then generate false-name bids for them. For example, if the number of actual bidders n is 3, we use subsets of size 2 and generate false-name bids for each of the two bidders. We randomly generate a certain number of false-name bids for each virtual bid and introduce the false-name-proof constraint for the virtual bids. Let us assume that bidder i submits k false-name bids by using $id_1, \dots, id_k$. We restate $v = (v_1, \dots, v_i, v_0, \dots, v_0, \dots, v_n)$ and $b = (b_1, \dots, b_{id_1}, \dots, b_{id_k}, \dots, b_n)$, where v_0 is a *null bidder* whose valuation for any bundle is zero; i.e., $v_0(B) = 0$, for any $B \subseteq M$. We define the expected regret for FNP as

$$\text{rgt_fnp}_i(a^w, p^w) = \mathbf{E} \left[\max_{b_{id_i} \in V_i} u_i(v_i, (b_{id_1}, \dots, b_{id_k}, b_{-i})) - u_i(v_i, (v_i, b_{-i})) \right]. \quad (8)$$

An auction satisfies FNP in expectation if and only if $\text{rgt_fnp}_i(a^w, p^w) \leq 0$ for any $i \in N$.

Figure 1 illustrates how false-name constraints are generated in an auction with two items and three bidders. We repeat the same process for six cases because there are three possible choices of two bidders and two possible choices of a bidder who makes false-name bids. The false-name bids (b_{id_1}, b_{id_2}) are randomly sampled from V_i . We sample a certain number of false-name bids for each case and assume that the expected regret for FNP is not positive; i.e., $\text{rgt_fnp}_i(a^w, p^w) \leq 0$.

**Fig. 1.** Generating false-name-proof constraints**Table 1.** Results with existing mechanisms, where A denotes results of allocation for bidders 1, 2, and 3 and P denotes results of payment for bidders 1, 2, and 3

Bids			ARP		MB		VCG	
Bidder 1	Bidder 2	Bidder 3	A	P	A	P	A	P
(0, 0, 10)	(0, 0, 0)	(0, 0, 8.4)	({1, 2}, ∅, ∅)	(8.4, 0, 0)	({1, 2}, ∅, ∅)	(8.4, 0, 0)	({1, 2}, ∅, ∅)	(8.4, 0, 0)
(5, 0, 5)	(0, 5, 5)	(0, 0, 8.4)	({1}, {2}, ∅)	(4.2, 4.2, 0)	(∅, ∅, {1, 2})	(0, 0, 5)	({1}, {2}, ∅)	(3.4, 3.4, 0)
(4, 0, 4)	(0, 5, 5)	(0, 0, 8.4)	(∅, ∅, {1, 2})	(0, 0, 5)	(∅, ∅, {1, 2})	(0, 0, 5)	({1}, {2}, ∅)	(3.4, 4.4, 0)

We define the Lagrangian function for the optimization problem with strategy-proof and false-name-proof constraints as

$$\begin{aligned}
& \widehat{\mathcal{L}}(a^w, p^w) + \sum_{i \in N} \lambda_i \widehat{\text{rgt_sp}}_i(a^w, p^w) + \frac{\rho}{2} \left(\sum_{i \in N} \widehat{\text{rgt_sp}}_i(a^w, p^w) \right)^2 \\
& + \sum_{i \in N} \mu_i \widehat{\text{rgt_fnp}}_i(a^w, p^w) + \frac{\sigma}{2} \left(\sum_{i \in N} \widehat{\text{rgt_fnp}}_i(a^w, p^w) \right)^2, \quad (9)
\end{aligned}$$

where $\widehat{\text{rgt_fnp}}_i(a^w, p^w) = \frac{1}{L} \sum_{l=1}^L \max_{b_{id_i} \in V_i} u_i^w(v_i^{(l)}, (b_{id_1}, \dots, b_{id_k}, v_{-i}^{(l)}) - u_i^w(v_i^{(l)}, v^{(l)})$, $\lambda, \mu \in \mathbb{R}^n$ is a vector of Lagrangian multipliers, and $\rho, \sigma > 0$ is a fixed parameter to control the weight of the quadratic penalty.

5 Experiments

We implemented a learning algorithm for our false-name-proof mechanisms that maximize the expected revenue in the RegretNet framework [3]⁵. Specifically, we extend the training algorithm by introducing false-name-proof constraints into the objective function defined in Section 4. In our experiments, we focused on combinatorial auctions with two items and three bidders for simplicity. We considered two different valuation settings: a discretized valuation setting and a uniform distribution setting. We used sample-based optimization for both misreports and false-name bids and generated 100 random misreports and 100 false-name bids for each valuation profile. The batch size and number of batches were set to 128 and 5000, respectively. The number of training iterations was 400,000.

⁵ <https://github.com/saisrivatsan/deep-opt-auctions>

Table 2. Results for discretized setting

Bids			Bidder 1		Bidder 2		Bidder 3	
Bidder 1	Bidder 2	Bidder 3	a_1	p_1	a_2	p_2	a_3	p_3
(0, 0, 10)	(0, 0, 0)	(0, 0, 8.4)	(0.000, 0.000, 1.000)	8.630	(0.000, 0.000, 0.000)	0.000	(0.000, 0.000, 0.000)	0
(5, 0, 5)	(0, 5, 5)	(0, 0, 8.4)	(0.858, 0.000, 0.000)	3.919	(0.000, 0.796, 0.000)	3.559	(0.000, 0.000, 0.142)	1.189
(4, 0, 4)	(0, 5, 5)	(0, 0, 8.4)	(0.320, 0.000, 0.000)	1.281	(0.000, 0.361, 0.000)	1.603	(0.000, 0.000, 0.639)	5.370

Table 3. Results for uniform distribution setting

Bids			Bidder 1		Bidder 2		Bidder 3	
Bidder 1	Bidder 2	Bidder 3	a_1	p_1	a_2	p_2	a_3	p_3
(0, 0, 1)	(0, 0, 0)	(0, 0, 0.84)	(0.035, 0.000, 0.770)	0.580	(0.000, 0.000, 0.000)	0.000	(0.000, 0.000, 0.000)	0.000
(0.5, 0, 0.5)	(0, 0.5, 0.5)	(0, 0, 0.84)	(0.253, 0.000, 0.030)	0.135	(0.000, 0.454, 0.000)	0.220	(0.000, 0.000, 0.000)	0.001
(0.4, 0, 0.4)	(0, 0.5, 0.5)	(0, 0, 0.84)	(0.000, 0.000, 0.000)	0.001	(0.000, 0.472, 0.000)	0.230	(0.000, 0.000, 0.000)	0.003

5.1 Discretized setting

The valuations of bidders were uniformly sampled from a finite valuation set:

$$V = \{(0, 0, 0), (4, 0, 4), (0, 4, 4), (5, 0, 5), (0, 5, 5), (0, 0, 8.4), (0, 0, 10),$$

where each tuple contains $(v_i(\{1\}), v_i(\{2\}), v_i(\{1, 2\}))$.

The average expected social surplus was 7.425, and the average expected revenue was 7.030 for independently generated test data. To clarify the property of the generated mechanism, we present the results for three existing mechanisms (VCG, MB, and ARP). While the VCG mechanism is vulnerable to false-name manipulation, MB and ARP satisfy FNP, as shown in Table 1. We chose three bid cases, as shown in Tables 1 and 2. The second and third cases can be considered as the situation in which bidder 1 in the first case (0, 0, 10) submitted false-name bids by using bidders 1 and 2. In Table 2, we show the results with three-decimal accuracy. We checked all cases of possible false-name manipulation and found that the generated mechanism satisfied FNP. For example, the utility of bidder 1 in the first case (0, 0, 10) was 1.37, but her utility when she submitted (5, 0, 5) and (0, 5, 5) became negative ($10 \times 0.858 \times 0.796 - 3.919 - 3.669 = -0.648$).

While we cannot exactly compare the generated mechanism with the existing mechanisms since the latter are deterministic, we can see that the results of the former are closely similar to those of the ARP mechanism.

5.2 Uniform distribution setting

The bidder valuations were real numbers sampled from finite intervals. We first uniformly sampled a class of agents from three bidder classes: single-minded bidder for item 1, item 2, and bundle $\{1, 2\}$, respectively. For the first class of agents, valuation was in the form $(v_i(\{1\}), 0, v_i(\{1\}))$, where $v_i(\{1\}) \sim U[0, 1]$. The second class of agents had valuations $(0, v_i(\{2\}), v_i(\{2\}))$, where $v_i(\{2\}) \sim U[0, 1]$. The third class of agents had valuations $(0, 0, v_i(\{1, 2\}))$, where $v_i(\{1, 2\}) \sim U[0, 2]$. The average expected social surplus was 1.021, and the average expected revenue was 0.755 for the test set. Table 3 shows the results for three cases of bids

by the generated mechanism. It satisfied FNP. The allocation and the price for each agent were lower than the results in the discretized setting. This is because the space of possible valuations was wider than that in the discretized setting.

6 Conclusion

We explored an approach to designing false-name-proof combinatorial auction mechanisms using deep learning techniques. We extended the existing RegretNet framework to handle false-name-proof constraints and then evaluated the generated mechanisms, demonstrating that they satisfy false-name-proofness. In future work we will extend our approach to more complicated settings.

Acknowledgments

This work was partially supported by JSPS KAKENHI Grant Numbers JP17H00761, JP17KK0008 and JP18H03337, by the Kayamori Foundation of Informational Science Advancement, and by the Telecommunications Advancement Foundation. We thank Paul Dütting and his coauthors for sharing the source code for the RegretNet framework.

References

1. Conitzer, V., Sandholm, T.: Complexity of mechanism design. In: UAI. pp. 103–110 (2002)
2. Conitzer, V., Sandholm, T.: Self-interested automated mechanism design and implications for optimal combinatorial auctions. In: ACM EC. pp. 132–141 (2004)
3. Dütting, P., Feng, Z., Narasimhan, H., Parkes, D., Ravindranath, S.S.: Optimal auctions through deep learning. In: ICML. pp. 1706–1715 (2019)
4. Feng, Z., Narasimhan, H., Parkes, D.C.: Deep learning for revenue-optimal auctions with budgets. In: AAMAS. pp. 354–362 (2018)
5. Guo, M., Conitzer, V.: Computationally feasible automated mechanism design: General approach and case studies. In: AAAI. pp. 1676–1679 (2010)
6. Iwasaki, A., Conitzer, V., Omori, Y., Sakurai, Y., Todo, T., Guo, M., Yokoo, M.: Worst-case efficiency ratio in false-name-proof combinatorial auction mechanisms. In: AAMAS. pp. 633–640 (2010)
7. Manisha, P., Jawahar, C.V., Gujar, S.: Learning optimal redistribution mechanisms through neural networks. In: AAMAS. pp. 345–353 (2018)
8. Yokoo, M.: Characterization of strategy/false-name proof combinatorial auction protocols: Price-oriented, rationing-free protocol. In: IJCAI. pp. 733–739 (2003)
9. Yokoo, M., Sakurai, Y., Matsubara, S.: Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence Journal* **130**(2), 167–181 (2001)