

Fourier Analysis-based Iterative Combinatorial Auctions*

Jakob Weissteiner^{1,4†}, Chris Wendler^{2†}, Sven Seuken^{1,4}, Ben Lubin³ and Markus Püschel²

¹University of Zurich

²ETH Zurich

³Boston University

⁴ETH AI Center

weissteiner@ifi.uzh.ch, chris.wendler@inf.ethz.ch, seuken@ifi.uzh.ch, blubin@bu.edu, pueschel@inf.ethz.ch

Abstract

Recent advances in Fourier analysis have brought new tools to efficiently represent and learn set functions. In this paper, we bring the power of Fourier analysis to the design of combinatorial auctions (CAs). The key idea is to approximate bidders' value functions using Fourier-sparse set functions, which can be computed using a relatively small number of queries. Since this number is still too large for practical CAs, we propose a new hybrid design: we first use neural networks (NNs) to learn bidders' values and then apply Fourier analysis to the learned representations. On a technical level, we formulate a Fourier transform-based winner determination problem and derive its mixed integer program formulation. Based on this, we devise an iterative CA that asks Fourier-based queries. We experimentally show that our hybrid ICA achieves higher efficiency than prior auction designs, leads to a fairer distribution of social welfare, and significantly reduces runtime. With this paper, we are the first to leverage Fourier analysis in CA design and lay the foundation for future work in this area. Our code is available on GitHub: <https://github.com/marketdesignresearch/FA-based-ICAs>.

1 Introduction

Combinatorial auctions (CAs) are used to allocate multiple heterogeneous items to bidders. CAs are particularly useful in domains where bidders' preferences exhibit *complementarities* and *substitutabilities* as they allow bidders to submit bids on *bundles* of items rather than on individual items.

Since the bundle space grows exponentially in the number of items, it is impossible for bidders to report values for all bundles in settings with more than a modest number of items. Thus, parsimonious preference elicitation is key for the practical design of CAs. For general value functions, Nisan and Segal [2006] have shown that to guarantee full efficiency, exponential communication in the number of items is needed.

Thus, practical CAs cannot provide efficiency guarantees in large domains. Instead, recent proposals have focused on *iterative combinatorial auctions (ICAs)*, where the auctioneer interacts with bidders over rounds, eliciting a *limited* amount of information, aiming to find a highly efficient allocation.

ICAs have found widespread application; most recently, for the sale of licenses to build offshore wind farms [Ausubel and Cramton, 2011]. For the sale of spectrum licenses, the combinatorial clock auction (CCA) [Ausubel *et al.*, 2006] has generated more than \$20 billion in total revenue [Ausubel and Baranov, 2017]. Thus, increasing the efficiency by only 1–2% points translates into monetary gains of millions of dollars.

1.1 Machine Learning-based Auction Design

Recently, researchers have used machine learning (ML) to improve the performance of CAs. Early work by Blum *et al.* [2004] and Lahaie and Parkes [2004] first studied the relationship between learning theory and preference elicitation in CAs. Dütting *et al.* [2019], Shen *et al.* [2019] and Rahme *et al.* [2021] used neural networks (NNs) to learn whole auction mechanisms from data. Brero *et al.* [2019] introduced a Bayesian ICA using probabilistic price updates to achieve faster convergence. Shen *et al.* [2020] use reinforcement learning for dynamic pricing in sponsored search auctions. Most related to the present paper is the work by Brero *et al.* [2018; 2021], who developed a value-query-based ML-powered ICA using support vector regressions (SVRs) that achieves even higher efficiency than the CCA. In follow-up work, Weissteiner and Seuken [2020] extended their ICA to NNs, further increasing the efficiency. In work subsequent to the first version of this paper, Weissteiner *et al.* [2022] proposed Monotone-Value Neural Networks (MVNNs), which are particularly well suited to learning value functions in combinatorial assignment domains. However, especially in large domains, it remains a challenge to find the efficient allocation while keeping the elicitation cost low. Thus, even state-of-the-art approaches suffer from significant efficiency losses and often result in unfair allocations, highlighting the need for better preference elicitation algorithms.

1.2 Combining Fourier Analysis and CAs

The goal of preference elicitation in CAs is to learn bidders' value functions using a small number of queries. Mathematically, value functions are *set functions*, which are in general

*Full paper including appendix available at arXiv: <https://arxiv.org/abs/2009.10749>

†These authors contributed equally to this paper.

exponentially large and notoriously hard to represent or learn. To address this complexity, we leverage Fourier analysis for set functions [Bernasconi *et al.*, 1996; O’Donnell, 2014; Püschel and Wendler, 2020]. In particular, we consider *Fourier-sparse approximations*, which are represented by few parameters. These parameters are the non-zero Fourier coefficients (FCs) obtained by a base change with the Fourier transform (FT). We use the framework by Püschel and Wendler [2020], which contains new FTs beyond the classical Walsh-Hadamard transform (WHT) [Bernasconi *et al.*, 1996], providing more flexibility. Until recently, methods for learning Fourier-sparse set functions focused on the WHT, and they placed assumptions on bidders’ value functions that are too restrictive for CAs [Stobbe and Krause, 2012]. However, recently, Amrollahi *et al.* [2019] proposed a new algorithm that can approximate general set functions by WHT-sparse ones, which is suitable for large CAs and we use it in this work.

1.3 Our Contribution

Our main contribution in this paper is to bring the power of Fourier analysis to CA design (Section 3). In particular, we formulate *FT-based winner determination problems (WDPs)* and derive corresponding mixed integer programs (MIPs) for several FTs (Section 4). Our MIPs allow for the efficient solution of the FT-based WDP and provide the foundation for using Fourier-sparse approximations in auction design.

We first experimentally show that the WHT performs best among the FTs in terms of induced level of sparsity (Section 5.1) and reconstruction error (Section 5.2). As an initial approach, we develop a WHT-based allocation rule (Section 5.3). However, this requires too many queries for direct use in CAs. To overcome this, we propose a practical hybrid ICA based on NNs and Fourier analysis (Section 6.1). The key idea is to compute Fourier-sparse approximations of NN-based bidder representations, enabling us to keep the number of queries small. The advantage of the NN-based representations is that they capture key aspects of the bidders’ value functions and can be queried arbitrarily often (Section 6.2).

Our efficiency experiments show that our hybrid ICA achieves higher efficiency than state-of-the-art mechanisms, leads to a significant computational speedup, and yields fairer allocations (Section 6.3). This shows that leveraging Fourier analysis in CA design is a promising new research direction.

2 Preliminaries

In this section, we present our formal model and review the MLCA mechanism, which our hybrid ICA builds upon.

2.1 Formal Model for ICAs

We consider a CA with n bidders and m indivisible items. Let $N = \{1, \dots, n\}$ and $M = \{1, \dots, m\}$ denote the set of bidders and items, respectively. We denote with $x \in \mathcal{X} = \{0, 1\}^m$ a bundle of items represented as an indicator vector, where $x_j = 1$ iff item $j \in M$ is contained in x . Bidders’ true preferences over bundles are represented by their (private) value functions $v_i : \mathcal{X} \rightarrow \mathbb{R}_+$, $i \in N$, i.e., $v_i(x)$ represents bidder i ’s true value for bundle x . By $a = (a_1, \dots, a_n) \in \mathcal{X}^n$ we denote an allocation of bundles to bidders, where a_i is

the bundle bidder i obtains. We denote the set of *feasible* allocations by $\mathcal{F} = \{a \in \mathcal{X}^n : \sum_{i \in N} a_{ij} \leq 1, \forall j \in M\}$. The (true) *social welfare* of an allocation a is defined as $V(a) = \sum_{i \in N} v_i(a_i)$. We let $a^* \in \arg\max_{a \in \mathcal{F}} V(a)$ be a social-welfare maximizing, i.e., *efficient*, allocation. The efficiency of any $a \in \mathcal{F}$ is measured by $V(a)/V(a^*)$. We assume that bidders’ have quasilinear utilities u_i , i.e., for a payments $p \in \mathbb{R}_+^n$ it holds that $u_i(a_i, p) = v_i(a_i) - p_i$.

An ICA *mechanism* defines how the bidders interact with the auctioneer and how the final allocation and payments are determined. We denote a bidder’s (possibly untruthful) reported value function by $\hat{v}_i : \mathcal{X} \rightarrow \mathbb{R}_+$. In this paper, we consider ICAs that ask bidders iteratively to report their value $\hat{v}_i(x)$ for particular bundles x selected by the mechanism (for early work on value queries see [Hudson and Sandholm, 2003]). A finite set of such reported bundle-value pairs of bidder i is denoted as $R_i = \{(x^{(l)}, \hat{v}_i(x^{(l)}))\}$, $x^{(l)} \in \mathcal{X}$. Let $R = (R_1, \dots, R_n)$ denote the tuple of reported bundle-value pairs obtained from all bidders. We define the *reported social welfare* of an allocation a given R as $\hat{V}(a|R) = \sum_{i \in N: (a_i, \hat{v}_i(a_i)) \in R_i} \hat{v}_i(a_i)$, where $(a_i, \hat{v}_i(a_i)) \in R_i$ ensures that only values for reported bundles contribute. Finally, the optimal allocation $a_R^* \in \mathcal{F}$ given the reports R is defined as

$$a_R^* \in \arg\max_{a \in \mathcal{F}} \hat{V}(a|R). \quad (1)$$

The final allocation $a_R^* \in \mathcal{F}$ and payments $p(R) \in \mathbb{R}_+^n$ are computed based on the elicited reports R only.

As the auctioneer can only ask each bidder i a limited number of queries $|R_i| \leq Q^{\max}$, the ICA needs a smart preference elicitation algorithm, with the goal of finding a highly efficient a_R^* with a limited number of value queries.

2.2 A Machine Learning-powered ICA

We now review the *machine learning-powered combinatorial auction (MLCA)* by Brero *et al.* [2021]. Interested readers are referred to Appendix A.1, where we present MLCA in detail.

MLCA starts by asking each bidder value queries for Q^{init} randomly sampled initial bundles. Next, MLCA proceeds in rounds until a maximum number of value queries per bidder $Q^{\max}$ is reached. In each round, for each bidder $i \in N$, it trains an ML algorithm $\mathcal{A}_i$ on the bidder’s reports R_i . Next, MLCA generates new value queries $q^{\text{new}} = (q_i^{\text{new}})_{i=1}^n$ with $q_i^{\text{new}} \in \mathcal{X} \setminus R_i$ by solving a ML-based WDP $q^{\text{new}} \in \arg\max_{a \in \mathcal{F}} \sum_{i \in N} \mathcal{A}_i(a_i)$. The idea is the following: if $\mathcal{A}_i$ are good surrogate models of the bidders’ true value functions then q^{new} should be a good proxy of the efficient allocation a^* and thus provide valuable information.

At the end of each round, MLCA receives reports R^{new} from all bidders for the newly generated q^{new} and updates R . When $Q^{\max}$ is reached, MLCA computes an allocation a_R^* maximizing the *reported* social welfare (eq. (1)) and determines VCG payments $p(R)$ (see Appendix A.2).

2.3 Incentives of MLCA and Hybrid ICA

A key concern in the design of ICAs are bidders’ incentives. However, the seminal result by Nisan and Segal [2006] discussed above implies that practical ICAs cannot simply use

VCG to achieve strategyproofness. And in fact, no ICA deployed in practice is *strategyproof* – including the famous SMRA and CCA auctions used to conduct spectrum auctions. Instead, auction designers have designed mechanisms that, while being manipulable, have “good incentives in practice” (see [Cramton, 2013; Milgrom, 2007]).

Naturally, the MLCA mechanism is also not strategyproof, and Brero *et al.* [2021] provide a simple example of a possible manipulation. The idea behind the example is straightforward: if the ML algorithm does not learn a bidder’s preferences perfectly, a sub-optimal allocation may result. Thus, a bidder may (in theory) benefit from misreporting their preferences with the goal of “correcting” the ML algorithm, so that, with the misreported preferences, the mechanism actually finds a preferable allocation.

However, MLCA has two features that mitigate manipulations. First, MLCA explicitly queries each bidder’s marginal economy, which implies that the marginal economy term of the final VCG payment is practically independent of bidder i ’s bid (for experimental support see [Brero *et al.*, 2021]). Second, MLCA enables bidders to “push” information to the auction which they deem useful. This mitigates certain manipulations of the main economy term in the VCG payment rule, as it allows bidders to increase the social welfare directly by pushing (useful) truthful information, rather than attempting to manipulate the ML algorithm. Brero *et al.* [2021] argued that with these two design features, MLCA exhibits very good incentives in practice. They performed a computational experiment, testing whether an individual bidder (equipped with more information than he would have in a real auction) can benefit from deviating from truthful bidding, while all other bidders are truthful. In their experiments, they could not identify a beneficial manipulation strategy. While this does not rule out that some (potentially more sophisticated) beneficial manipulations do exist, it provides evidence to support the claim that MLCA has good incentives in practice.

With two additional assumptions, one also obtains a theoretical incentive guarantee for MLCA. Assumption 1 requires that, if all bidders bid truthfully, then MLCA finds an efficient allocation (we show in Appendix D.3 that in two of our domains, we indeed find the efficient allocation in the majority of cases). Assumption 2 requires that, for all bidders i , if all other bidders report truthfully, then the social welfare of bidder i ’s marginal economy is independent of his value reports. If both assumptions hold, then bidding truthfully is an ex-post Nash equilibrium in MLCA.

Our hybrid ICA (Algorithm 1 in Section 6.1) is built upon MLCA, leaving the general framework in place, and only changing the algorithm that generates new queries each round. Given this design, the incentive properties of MLCA extend to the hybrid ICA. Specifically, our hybrid ICA is also not strategyproof, but it also has the same design features (including push-bids) to mitigate manipulations.

In future work, it would be interesting to evaluate experimentally whether the improved performance of the hybrid ICA translates into better manipulation mitigation compared to MLCA. However, such an analysis is beyond the scope of the present paper, which focuses on the ML algorithm that is integrated into the auction mechanism.

3 Fourier Analysis of Value Functions

We now show how to apply Fourier analysis to value functions providing the theoretical foundation of FT-based WDPs.

Classic Fourier analysis decomposes an audio signal or image into an orthogonal set of sinusoids of different frequencies. Similarly, the classical Fourier analysis for *set functions* (i.e., functions mapping subsets of a discrete set to a scalar) decomposes a set function into an orthogonal set of Walsh functions [Bernasconi *et al.*, 1996], which are piecewise constant with values 1 and -1 only. Recent work by Püschel and Wendler [2020] extends the Fourier analysis for set functions with several novel forms of set Fourier transforms (FTs). Importantly, because bidders’ value functions are set functions, they are amenable to this type of Fourier analysis, and it is this connection that we will leverage in our auction design.

Sparsity. The motivation behind our approach is that we expect bidders’ value functions to be *sparse*, i.e., they can be described with much less data than is contained in the exponentially-sized full value function. While this sparsity may be difficult to uncover when looking at bidders’ value reports directly, it may reveal itself in the Fourier domain (where then most FCs are zero). As all FTs are changes of basis, each FT provides us with a new *lens on the bidder’s value function*, revealing structure and thus potentially reducing dimensionality.

Set function Fourier transform. We now provide a formal description of FTs for reported value functions $\hat{v}_i$. To do so, we represent $\hat{v}_i$ as a vector $(\hat{v}_i(x))_{x \in \mathcal{X}}$. Each known FT is a change of basis and thus can be represented by a certain matrix $F \in \{-1, 0, 1\}^{2^m \times 2^m}$ with the form:

$$\phi_{\hat{v}_i}(y) = (F\hat{v}_i)(y) = \sum_{x \in \mathcal{X}} F_{y,x} \hat{v}_i(x). \quad (2)$$

There is exactly one Fourier coefficient per bundle, this follows from the theory presented by Püschel and Wendler [2020]. The corresponding inverse transform F^{-1} is thus:

$$\hat{v}_i(x) = (F^{-1}\phi_{\hat{v}_i})(x) = \sum_{y \in \mathcal{X}} F_{x,y}^{-1} \phi_{\hat{v}_i}(y). \quad (3)$$

$\phi_{\hat{v}_i}$ is again a set function and we call $\phi_{\hat{v}_i}(y)$ the *Fourier coefficient* at frequency y . A value function is *Fourier-sparse* if $|\text{supp}(\phi_{\hat{v}_i})| = |\{y : \phi_{\hat{v}_i}(y) \neq 0\}| \ll 2^m$. We call $\text{supp}(\phi_{\hat{v}_i})$ the *Fourier support* of $\hat{v}_i$.

Classically, the WHT is used as F [Bernasconi *et al.*, 1996; O’Donnell, 2014], but we also consider two recently introduced FTs (FT3, FT4) due to their information-theoretic interpretation given in [Püschel and Wendler, 2020]:

$$\mathbf{FT3}: F_{y,x} = (-1)^{|y|-|x|} \mathbb{I}_{\min(x,y)=x}, \quad (4)$$

$$\mathbf{FT4}: F_{y,x} = (-1)^{|\min(x,y)|} \mathbb{I}_{\max(x,y)=1_m}, \quad (5)$$

$$\mathbf{WHT}: F_{y,x} = \frac{1}{2^m} (-1)^{|\min(x,y)|}. \quad (6)$$

Here, $\min$ is the elementwise minimum (intersection of sets), $\max$ analogously, $|\cdot|$ is the set size, 1_m denotes the m -dimensional vector of 1s, and the indicator function $\mathbb{I}_P$ is equal to 1 if the predicate P is true and 0 otherwise.

Notions of Fourier-sparsity. In recent years, the notion of Fourier-sparsity has gained considerable attention, leading to highly efficient algorithms to compute FTs [Stobbe and Krause, 2012; Amrollahi *et al.*, 2019; Wendler *et al.*, 2021]. Many classes of set functions are Fourier-sparse (e.g., graph cuts, hypergraph valuations and decision trees [Abraham *et al.*, 2012]) and can thus be learned efficiently. The benefit of considering multiple FTs is that they offer different, non-equivalent notions of sparsity as illustrated by the following example.

Example 1. Consider the set of items $M = \{1, 2, 3\}$ and the associated reported value function $\hat{v}_i$ shown in Table 1 (where we use 001 as a shorthand notation for $(0, 0, 1)$), together with the corresponding FCs $\phi_{\hat{v}_i}$. This bidder ex-

	000	100	010	001	110	101	011	111
$\hat{v}_i$	0	1	1	1	3	3	3	5
FT3	0	1	1	1	1	1	1	-1
FT4	5	-2	-2	-2	0	0	0	1
WHT	17/8	-7/8	-7/8	-7/8	1/8	1/8	1/8	1/8

Table 1: Example with different induced notions of sparsity of all considered FTs.

hibits complementary effects for each bundle containing more than one item, as can be seen, e.g., from $3 = \hat{v}_i(110) > \hat{v}_i(100) + \hat{v}_i(010) = 2$ and $5 = \hat{v}_i(111) > \hat{v}_i(100) + \hat{v}_i(010) + \hat{v}_i(001) = 3$. Observe that while this value function is sparse in FT4, i.e., $\phi_{\hat{v}_i}(110) = \phi_{\hat{v}_i}(101) = \phi_{\hat{v}_i}(011) = 0$, it is neither sparse in FT3 nor WHT. Note that the coefficients $\phi_{\hat{v}_i}(100)$, $\phi_{\hat{v}_i}(010)$, and $\phi_{\hat{v}_i}(001)$ capture the value of single items and thus cannot be zero.

The induced spectral energy distributions for each FT, i.e., for each cardinality (i.e., number of items) d from 0 to $m = 3$, we compute $\sum_{y \in \mathcal{X}: |y|=d} \phi_{\hat{v}_i}(y)^2 / \sum_{y \in \mathcal{X}} \phi_{\hat{v}_i}(y)^2$, are shown in Table 2.

	$d = 0$	$d = 1$	$d = 2$	$d = 3$
FT3	0.00	42.86	42.86	14.28
FT4	65.79	31.58	0.00	2.63
WHT	65.69	33.41	0.68	0.22

Table 2: Spectral energy in % for each cardinality (i.e., number of items) d from 0 to $m = 3$ of all considered FTs.

Fourier-sparse approximations. In practice, $\hat{v}_i$ may only be approximately sparse. Meaning that while not being sparse, it can be approximated well by a Fourier-sparse function $\tilde{v}_i$. Formally, let $\mathcal{S}_i = \text{supp}(\phi_{\hat{v}_i})$ with $|\mathcal{S}_i| = k$, we call

$$\tilde{v}_i(x) = \sum_{y \in \mathcal{S}_i} F_{x,y}^{-1} \phi_{\hat{v}_i}(y) \text{ for all } x \in \mathcal{X} \quad (7)$$

such that $\|\tilde{v}_i - \hat{v}_i\|_2$ is small a k -Fourier-sparse approximation of $\hat{v}_i$. We denote the vector of FCs by $\phi_{\tilde{v}_i|\mathcal{S}_i} = (\phi_{\tilde{v}_i}(y))_{y \in \mathcal{S}_i}$.

4 Fourier Transform-based WDPs

To leverage Fourier analysis for CA design, we represent bidders' value functions using Fourier-sparse approximations. A

key step in most auction designs is to find the social welfare-maximizing allocation given bidder's reports, which is known as the *Winner Determination Problem* (WDP). To apply FTs, we need to be able to solve the WDP efficiently. Accordingly, we next derive MIPs for each of the FTs.

For each bidder $i \in N$, let $\tilde{v}_i : \mathcal{X} \rightarrow \mathbb{R}_+$ be a Fourier-sparse approximation of the bidders' reported value function $\hat{v}_i$. Next, we define the *Fourier transform-based WDP*.

Definition 1. (FOURIER TRANSFORM-BASED WDP)

$$\arg\max_{a \in \mathcal{F}} \sum_{i \in N} \tilde{v}_i(a_i). \quad (\text{FT-WDP})$$

For $x, y \in \mathbb{R}^d$, let $x \leq y$, $\max(x, y)$ and $(-1)^x$ be defined component-wise, and let $\langle \cdot, \cdot \rangle$ denote the Euclidean scalar product. First, we formulate succinct representations of $\tilde{v}_i$.

Lemma 1. For $i \in N$ let $\mathcal{S}_i = \{y^{(1)}, \dots, y^{(k)}\}$ be the support of a k -Fourier-sparse approximation $\tilde{v}_i$ and $W_i \in \{0, 1\}^{k \times m}$ be defined as $(W_i)_{l,j} = \mathbb{I}_{y_j^{(l)}=1}$. Then it holds that

$$\text{FT3: } \tilde{v}_i(x) = \langle \phi_{\tilde{v}_i|\mathcal{S}_i}, \max(0_k, 1_k - W_i(1_m - x)) \rangle \quad (8)$$

$$\text{FT4: } \tilde{v}_i(x) = \langle \phi_{\tilde{v}_i|\mathcal{S}_i}, \max(0_k, 1_k - W_i x) \rangle \quad (9)$$

$$\text{WHT: } \tilde{v}_i(x) = \langle \phi_{\tilde{v}_i|\mathcal{S}_i}, (-1)^{W_i x} \rangle. \quad (10)$$

See Appendix B.1 for the proof. With Lemma 1 and rewriting $\max(\cdot, \cdot)$ and $(-1)^{\cdot}$ as linear constraints, we next encode (FT-WDP) as a MIP (see Appendix B.2 for the proof).

Theorem 1. (FT-BASED MIPs) Let $\tilde{v}_i : \mathcal{X} \rightarrow \mathbb{R}$ be a k -Fourier-sparse approximation from (8), (9), or (10). Then there exists a $C > 0$ s.t. the MIP defined by the objective

$$\arg\max_{a \in \mathcal{F}, \beta_i \in \{0,1\}^k} \sum_{i \in N} \langle \phi_{\tilde{v}_i|\mathcal{S}_i}, \alpha_i \rangle, \quad (11)$$

and for $i \in N$ one set of transform specific constraints (12)–(14), or (15)–(17), or (18)–(20), is equivalent to (FT-WDP).

$$\text{FT3: } \text{s.t. } \alpha_i \geq 1_k - W_i(1_m - a_i) \quad (12)$$

$$\alpha_i \leq 1_k - W_i(1_m - a_i) + C\beta_i \quad (13)$$

$$0_k \leq \alpha_i \leq C(1_k - \beta_i) \quad (14)$$

$$\text{FT4: } \text{s.t. } \alpha_i \geq 1_k - W_i a_i \quad (15)$$

$$\alpha_i \leq 1_k - W_i a_i + C\beta_i \quad (16)$$

$$0_k \leq \alpha_i \leq C(1_k - \beta_i) \quad (17)$$

$$\text{WHT: } \text{s.t. } \alpha_i = -2\beta_i + 1_k \quad (18)$$

$$\beta_i = W_i a_i - 2\gamma_i \quad (19)$$

$$\gamma_i \in \mathbb{Z}^k \quad (20)$$

5 Analyzing the Potential of a FT-based CA

In this section, we experimentally evaluate the FTs and propose an FT-based allocation rule that motivates our practical hybrid ICA mechanism presented later in Section 6.

For our experiments, we use the spectrum auction test suite (SATS) [Weiss *et al.*, 2017].¹ SATS enables us to generate

¹We used SATS version 0.6.4 for our experiments. The implementations of GSVM and LSVM have changed slightly in newer SATS versions. This must be considered when comparing the performance of different mechanisms in those domains.




Figure 1: Spectral energy distribution in LSVM for all FTs. For each cardinality (x-axis), we plot the spectral energy of all frequencies of that cardinality normalized by the total spectral energy (y-axis).

synthetic CA instances in different domains. We have access to each bidder’s *true* full value function v_i and the efficient allocation a^* . When simulating bidders, we assume truthful bidding (i.e., $\hat{v}_i = v_i$). We consider three domains:

The Global Synergy Value Model (GSVM) [Goeree and Holt, 2010] has 18 items, 6 *regional* and 1 *national bidder*.

The Local Synergy Value Model (LSVM) [Scheffel *et al.*, 2012] consists of 18 items, 5 *regional* and 1 *national bidder*. Complementarities arise from spatial proximity of items.

The Multi-Region Value Model (MRVM) [Weiss *et al.*, 2017] has 98 items and 10 bidders (categorized as *local*, *regional*, or *national*) and models large US spectrum auctions.

5.1 Notions of Fourier Sparsity

We first experimentally show that different notions of FT lead to different types of sparsity in LSVM (for other domains see Appendix C.1). For this we first compute the FTs of all bidders and then calculate their *spectral energy* distribution. That is, for each cardinality d (#items) from 0 to m , we compute $\sum_{y \in \mathcal{X}: |y|=d} \phi_{\hat{v}_i}(y)^2 / \sum_{y \in \mathcal{X}} \phi_{\hat{v}_i}(y)^2$. In Figure 1 we present the mean over 30 LSVM instances and bidder types.

Figure 1, shows that while the energy is spread among FCs of various degrees in FT3 and FT4, in WHT the low degree (≤ 2) FCs contain most of the energy, i.e., the WHT has much fewer dominant FCs that accurately describe each value function. As the WHT is orthogonal, learning low degree WHT-sparse approximations leads to low reconstruction error. Low degree WHT-sparse approximations can be learnt efficiently and accurately from a small number of queries using compressive sensing [Stobbe and Krause, 2012].

Note that the FT3 is identical to the classical polynomial value function representation [Lahaie, 2010] defined as

$$\hat{v}_i^{\text{poly}}(x) = \sum_{l=1}^m \sum_{j=\{j_1, \dots, j_l\} \subseteq M} x_{j_1} \cdot \dots \cdot x_{j_l} \cdot c_j^{(i)}. \quad (21)$$

where the coefficient $c_j^{(i)}$ is equal to the FT3 FC at frequency y with $y_i = 1$ for $i \in \{j_1, \dots, j_l\}$ and $y_i = 0$ else.² E.g. for $M = \{1, 2\}$, $\hat{v}_i^{\text{poly}}(x) = x_1 c_{\{1\}}^{(i)} + x_2 c_{\{2\}}^{(i)} + x_1 x_2 c_{\{1,2\}}^{(i)}$. Thus, converting $\hat{v}_i^{\text{poly}}$ into another FT basis (here WHT) can indeed be very helpful for the design of ML-based CAs.

²This can be seen by calculating the inverse in (4), i.e., $F_{y,x}^{-1} = \mathbb{I}_{\min(x,y)=x}$, and plug in $F_{y,x}^{-1}$ into (3).

DOMAIN	K	BIDDER	FT3	FT4	WHT	NN
GSVM	100	NAT.	11.3 ± 0.7	14.2 ± 0.8	1.8 ± 0.1	9.0 ± 1.8
		REG.	0.0	1.4 ± 0.2	0.4 ± 0.1	7.2 ± 0.9
	200	NAT.	0.0	0.0	0.0	5.7 ± 0.4
		REG.	0.0	0.0	0.0	5.2 ± 0.8
LSVM	100	NAT.	78.4 ± 1.0	580.2 ± 7.9	31.2 ± 0.4	48.7 ± 1.2
		REG.	28.2 ± 2.3	48.5 ± 2.7	6.8 ± 0.3	17.8 ± 0.7
	200	NAT.	95.8 ± 1.2	639.0 ± 10.0	26.2 ± 0.3	40.6 ± 0.7
		REG.	25.8 ± 2.0	43.1 ± 2.4	5.3 ± 0.3	15.3 ± 0.9

Table 3: Reconstruction error with a 95%-CI of k -Fourier-sparse approximations $\tilde{v}_i$ and NNs trained on k randomly selected bundles. Winners are marked in grey.

5.2 Reconstruction Error of Fourier Transforms

Next we validate the FT approach by comparing the reconstruction error of the FTs in the medium-sized GSVM and LSVM, where we can still compute the full FT (in contrast to MRVM). For now, we assume that we have access to bidders’ full $\hat{v}_i$. In Procedure 1, we determine the best k -Fourier-sparse approximation $\tilde{v}_i$ (see Appendix C.2 for details).

Procedure 1. (BEST FCS GIVEN FULL ACCESS TO $\hat{v}_i$)

Compute $\tilde{v}_i$ using the k absolutely largest FCs $\phi_{\hat{v}_i|S_i}$ from the full FT for each bidders’ reported value function $\phi_{\hat{v}_i} = F\hat{v}_i$.

Remark 1. Since the WHT is orthogonal and the simulated auction data is noise-free, its approximation error is exactly equal to the residual of the FCs. Thus, Procedure 1 is optimal for the WHT. This is not the case for FT3 and FT4 because they are not orthogonal.

We then calculate the RMSE ($\frac{1}{2m} \sum_{x \in \mathcal{X}} (\hat{v}_i(x) - \tilde{v}_i(x))^2$)^{1/2} averaged over 100 instances and bidder types. In Table 3, we present the RMSEs for the three FTs and for NNs, where we used the architectures from Weissteiner and Seuken [2020].

For GSVM, we observe that we can perfectly reconstruct $\hat{v}_i$ with the 200 best FCs, which shows that GSVM is 200-sparse. In contrast, LSVM is non-sparse, and we do not achieve perfect reconstruction with 200 FCs. Overall, we observe that the WHT outperforms FT3 and FT4. Moreover, we see that, if we could compute the k best FCs of the WHT from k training points, the WHT would outperform the NNs.

However, in practice, we do not have access to full value functions. Instead, we must use an algorithm that computes the best FCs using a reasonable number of value queries.

Remark 2. Thanks to its orthogonality the WHT has strong theoretical guarantees for sparse recovery from few samples using compressive sensing (see [Stobbe and Krause, 2012]). Thus, we focus on the WHT in the remainder of this paper.

5.3 A Fourier Transform-based Allocation Rule

We now present an FT-based allocation rule using the *robust sparse WHT algorithm (RWHT)* by Amrollahi *et al.* [2019]. RWHT learns a Fourier-sparse approximation $\tilde{v}_i$ of $\hat{v}_i$ from *value queries*. Procedure 2 finds the allocation $\tilde{a}$.

Procedure 2. (WHT-BASED ALLOCATION RULE)

- i. Use RWHT to compute k -sparse approximations $\tilde{v}_i, i \in N$.
- ii. Solve $\tilde{a} \in \arg\max_{a \in \mathcal{F}} \sum_{i \in N} \tilde{v}_i(a_i)$ using Theorem 1.




Figure 2: Efficiency of Procedure 2 in GSVM.

In Figure 2, we present the efficiency of $\tilde{a}$ on 50 GSVM instances for various values of k . We see that RWHT achieves a median efficiency of 100% for 90 or more FCs. Nevertheless, the main practical issue with this approach is the number of value queries required. As we can see, RWHT needs 102,000 value queries (39% of all bundles) to find the 90 best FCs. For a practical ICA mechanism this is far too many.

6 A Practical Hybrid ICA Mechanism

In this section, we introduce and experimentally evaluate a practical hybrid ICA mechanism, based on FTs and NNs.

6.1 The Hybrid ICA Mechanism

The main issue of the FT-based allocation rule in Section 5.3 is the large number of queries, which we now address. The idea is the following: instead of directly applying a sparse FT algorithm (like RWHT) to bidders, we apply it to a NN-based representation. In this way, we query NNs instead of bidders. Based on the FCs of the NNs, we determine a Fourier-sparse approximation $\tilde{v}_i$ with only *few value queries*, where the idea is that the FCs of each NN concentrate on the most dominant FCs of its respective value function. Indeed, recent evidence suggests that a NN trained by SGD can learn the Fourier-support [Rahaman *et al.*, 2019]. We analyze our NN support discovery rule in Section 6.2. We now present HYBRID ICA, leaving details of the sub-procedures to Appendix D.1.

HYBRID ICA (Algorithm 1) consists of 3 phases: the *MLCA*, the *Fourier reconstruction*, and the *Fourier allocation phase*. It is parameterized by an FT F and the numbers $\ell_1, \ell_2, \ell_3, \ell_4$ of different query types. In total, it asks each bidder $\sum_{i=1}^4 \ell_i$ queries: ℓ_1 random initial, ℓ_2 MLCA, ℓ_3 Fourier reconstruction, and ℓ_4 Fourier allocation queries.

1. MLCA Phase. We first run MLCA such that the NNs can then be trained on “meaningfully” elicited reports. In MLCA, we request reports for ℓ_1 random initial bundles and for ℓ_2 MLCA queries (Lines 1-2).

2. Fourier Reconstruction Phase. Next, we compute a Fourier-sparse approximation $\tilde{v}_i$. For this, we first fit a NN $\mathcal{N}_i$ on the reports R_i (Line 4). Then we compute the best FCs of the fitted NNs (Line 5, **Procedure 3**) in order to discover which FCs are important to represent the bidders. Based on these FCs, we determine ℓ_3 Fourier reconstruction queries $\tilde{S}_i$ (Line 6, **Procedure 4**), send them to the bidders and fit $\tilde{v}_i$ to the reports R_i received so far (Line 7, **Procedure 5**).

Algorithm 1: HYBRID ICA

Params: F Fourier transform; $\ell_1, \ell_2, \ell_3, \ell_4$ query split

- 1 Set $Q^{\text{init}} = \ell_1$ and $Q^{\text{max}} = \ell_1 + \ell_2$ ▷ MLCA phase
- 2 Run MLCA($Q^{\text{init}}, Q^{\text{max}}$); get $\ell_1 + \ell_2$ reports R
- 3 **foreach** bidder $i \in N$ **do** ▷ Fourier reconstr. phase
- 4 Fit NN $\mathcal{N}_i$ to R_i
- 5 Determine the best FCs of $\mathcal{N}_i$ ▷ Proc. 3
- 6 Compute ℓ_3 reconstruction queries $\tilde{S}_i \subseteq \mathcal{X}$ ▷ Proc. 4
- 7 Ask $\tilde{S}_i$, add reports to R_i , and fit $\tilde{v}_i$ to R_i ▷ Proc. 5
- 8 **for** $l = 1, \dots, \ell_4$ **do** ▷ Fourier alloc. phase
- 9 Solve $q \in \arg\max_{a \in \mathcal{F}} \sum_{i \in N} \tilde{v}_i(a_i)$ (FT-WDP)
- 10 **foreach** bidder $i \in N$ **do**
- 11 **if** $q_i \in R_i$ **then** ▷ Bundle already queried
- 12 Define $\mathcal{F}' = \{a \in \mathcal{F} : a_i \neq x, \forall x \in R_i\}$
- 13 Resolve $q' \in \arg\max_{a \in \mathcal{F}'} \sum_{i \in N} \tilde{v}_i(a_i)$
- 14 Update $q_i = q'_i$
- 15 Query bidder i 's value for q_i and add report to R_i
- 16 Fit $\tilde{v}_i$ to R_i ▷ Proc. 5
- 17 From R compute: a_R^* as in eq. (1), VCG payments $p(R)$
- 18 **return** Final allocation a_R^* and VCG payments $p(R)$

3. Fourier Allocation Phase. We use the fitted $\tilde{v}_i$ to generate ℓ_4 Fourier allocation queries. Here, we solve the FT-based WDP (Line 9) to get candidate queries q , ensure that all queries are new (Lines 11–14), add the received reports to R_i (Line 15) and refit $\tilde{v}_i$ (Line 16). Finally in Line 17, HYBRID ICA computes based on all reports R a welfare-maximizing allocation a_R^* and VCG payments $p(R)$ (see Appendix A.2).

Experiment Setup. For HYBRID ICA and MLCA, we use the NN architectures from Weissteiner and Seuken [2020] and set a total query budget of 100 (GSVM, LSVM) and 500 (MRVM). For HYBRID ICA, we optimized the FTs and query parameters ℓ_i on a training set of CA instances. Table 4 shows the best configurations.

	NN ARCHITECTURES	FT	ℓ_1	ℓ_2	ℓ_3	ℓ_4
GSVM	R:[32, 32] N:[10, 10]	WHT	30	21	20	29
LSVM	R:[32, 32] N:[10, 10, 10]	WHT	30	30	10	30
MRVM	L,R,N:[16, 16]	WHT	30	220	0	250

 Table 4: Best configuration of HYBRID ICA. R:[d_1, d_2] denotes a 3-hidden-layer NN for the regional bidder with d_1 , and d_2 nodes.

6.2 NNs Support Discovery Experiments

In HYBRID ICA we use the NNs for support discovery where it is key that the FCs of these NNs concentrate on the dominant FCs of its value function, i.e. $\text{supp}(\phi_{\mathcal{N}_i}) \approx \text{supp}(\phi_{\tilde{v}_i})$.

To evaluate the NN-based support discovery (Line 5), we consider the spectral *energy ratio* obtained by dividing the spectral energies of the k frequencies selected from the NN and the k best frequencies (for the WHT the best FCs are the ones with the largest absolute value). Formally, for each bidder i , let the k frequencies selected from the NN be $\tilde{S}_i = \{\tilde{y}^{(1)}, \dots, \tilde{y}^{(k)}\}$ and the best ones be $S_i^* = \{y^{*(1)}, \dots, y^{*(k)}\}$. Then, bidder i 's energy ratio is given by $\sum_{\tilde{y} \in \tilde{S}_i} \phi_{\tilde{v}_i}(\tilde{y})^2 / \sum_{y \in S_i^*} \phi_{\tilde{v}_i}(y)^2 \in [0, 1]$ (see Appendix D.2 for details). This ratio is equal to one if $\tilde{S}_i = S_i^*$. Figure 3

MECHANISM	GSVM					LSVM					MRVM					
	EFFICIENCY IN %	REGIONAL IN %	NATIONAL IN %	REV IN %	HRS/ INST.	EFFICIENCY IN %	REGIONAL IN %	NATIONAL IN %	REV IN %	HRS/ INST.	EFFICIENCY IN %	LOCAL IN %	REGIONAL IN %	NATIONAL IN %	REV IN %	HRS/ INST.
HYBRID ICA	99.97 ± 0.03	94.72	5.25	81	0.81	98.74 ± 0.43	89.09	9.65	78	1.95	96.63 ± 0.31	0.00	1.19	95.44	36	23.88
MLCA	99.17 ± 0.37	98.11	1.06	79	4.65	99.14 ± 0.42	93.40	5.75	77	6.09	95.32 ± 0.32	0.00	0.53	94.79	41	43.26
HYBRID ICA (NO FR)	98.30 ± 0.49	97.94	0.36	75	0.93	98.16 ± 0.60	93.83	4.33	72	2.03	96.63 ± 0.31	0.00	1.19	95.44	36	23.88
HYBRID ICA (NO FR/FA)	98.16 ± 0.50	97.47	0.69	75	0.71	97.75 ± 0.63	92.78	5.27	72	1.86	93.91 ± 0.36	0.01	0.42	93.48	42	14.68
Efficient Allocation		94.75	5.25				84.03	13.7				0.00	2.11	97.89		

Table 5: HYBRID ICA vs. MLCA, HYBRID ICA (NO FR), and HYBRID ICA (NO FR/FA). All results are averages over a test set of 100 (GSVM and LSVM) and 30 (MRVM) CA instances. For efficiency we give a 95% confidence interval and mark the best mechanisms in grey.




Figure 3: Average energy ratio (y-axis) with 97.5% and 2.5% empirical quantiles for a number of selected frequencies k (x-axis) over 30 instances in GSVM and LSVM and over 5 instances in MRVM.

shows that the NN-based supports are almost on par with the best supports given a fixed budget of k frequencies.

6.3 Efficiency Experiments

We now evaluate the efficiency of HYBRID ICA vs MLCA.

Results. Table 5 contains our main results in all domains.³ We show efficiency, distribution of efficiency to bidder types, revenue ($\sum_{i \in N} p(R)_i / V(a^*)$), and runtime. First, we see that HYBRID ICA *statistically significantly outperforms* MLCA *w.r.t. efficiency* in GSVM and MRVM and performs on par in LSVM. Second, it also leads to a computational speedup ($\times 6$ GSVM, $\times 3$ LSVM, $\times 2$ MRVM). The reason for this computational speedup is that the generation of the $\ell_3 + \ell_4$ Fourier queries (estimating the superset of the support using RWHT on the NNs, fitting the FT models using compressive sensing and solving our new FT-based MIPs) is faster than the generation of the NN-based MLCA allocation queries (training NNs and solving the NN-based MIP). Third, it distributes the welfare more evenly (= fairer) to bidder types.⁴ This also leads to a distribution that more closely resembles that of the efficient allocation (see *Efficient Allocation*). We present full efficiency path plots for the different phases of HYBRID ICA in Appendix D.3.

Fourier queries. To verify the importance of the ℓ_3 Fourier reconstruction and ℓ_4 Fourier allocation queries, we also present HYBRID ICA (NO FR) and HYBRID ICA (NO FR/FA), which use random queries in place of the ℓ_3 Fourier reconstruction and the $\ell_3 + \ell_4$ Fourier-based queries. As we see in Table 5, using the Fourier queries leads to significantly better efficiency and HYBRID ICA (NO FR/FA) does not achieve a fairer efficiency distribution. A comparison of HYBRID ICA to HYBRID ICA (NO FR) reveals that, in GSVM and LSVM, the Fourier reconstruction queries cause the fairer distribution. We empirically verified that these queries are

³All experiments were conducted on machines with Intel Xeon E5 v4 2.20GHz processors with 24 cores and 128GB RAM or with Intel E5 v2 2.80GHz processors with 20 cores and 128GB RAM.

⁴We consider an allocation to be “fairer” if its social welfare is more evenly distributed among bidder types. This is similar (but not identical) to the standard notion of egalitarian social welfare.

composed of larger bundles (i.e., 17 items c.p. to 4 in MLCA queries) and thus allocate large bundles to bidders that would have been overlooked. In MRVM, the optimal query split for HYBRID ICA uses $\ell_3 = 0$ Fourier reconstruction queries such that HYBRID ICA is equal to HYBRID ICA (NO FR). Thus, in MRVM, HYBRID ICA’s increased efficiency and fairer distribution results from the Fourier allocation queries.

Overall, we see that our Fourier-based auction is especially powerful in sparse domains. In practice, bidders are often limited by their cognitive abilities [Scheffel *et al.*, 2012] or use a low-dimensional computational model to represent their value function. Thus, their reported preferences typically exhibit only a limited degree of substitutability and complementarity, which is captured well by Fourier-sparsity.

7 Conclusion

We have introduced Fourier analysis for the design of CAs. The main idea was to represent value functions using Fourier-sparse approximations, providing us with a new lens on bidder’s values in the Fourier domain.

On a technical level, we have derived succinct MIPs for the Fourier transform-based WDPs, which makes computing Fourier-based allocation queries practically feasible. We have leveraged this to design a new hybrid ICA that uses NN and Fourier-queries. Our experiments have shown that our approach leads to higher efficiency, a computational speedup and a fairer distribution of social welfare than state-of-the-art.

With this paper, we have laid the foundations for future work leveraging Fourier analysis for representing and eliciting preferences in combinatorial settings.

Acknowledgments

We thank the anonymous reviewers for helpful comments. This paper is part of a project that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 805542). Furthermore, this material is based upon work supported by the National Science Foundation under grant no. CMMI-1761163.

References

- [Abraham *et al.*, 2012] I. Abraham, M. Babaioff, S. Dughmi, and T. Roughgarden. Combinatorial auctions with restricted complements. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 2012.
- [Amrollahi *et al.*, 2019] A. Amrollahi, A. Zandieh, M. Kapralov, and A. Krause. Efficiently learning

- Fourier sparse set functions. In *Advances in Neural Information Processing Systems*, 2019.
- [Ausubel and Baranov, 2017] L. M. Ausubel and O. Baranov. A practical guide to the combinatorial clock auction. *Economic Journal*, 127(605):F334–F350, 2017.
- [Ausubel and Cramton, 2011] L. Ausubel and P. Cramton. Auction design for wind rights. *Report to Bureau of Ocean Energy Management, Regulation and Enforcement*, 2011.
- [Ausubel et al., 2006] L. Ausubel, P. Cramton, and P. Milgrom. The clock-proxy auction: A practical combinatorial auction design. In *Combinatorial Auctions*, pages 115–138. MIT Press, 2006.
- [Bernasconi et al., 1996] A. Bernasconi, B. Codenotti, and J. Simon. On the Fourier analysis of Boolean functions. *preprint*, pages 1–24, 1996.
- [Blum et al., 2004] A. Blum, J. Jackson, T. Sandholm, and M. Zinkevich. Preference elicitation and query learning. *Journal of Machine Learning Research*, 5:649–667, 2004.
- [Brero et al., 2018] G. Brero, B. Lubin, and S. Seuken. Combinatorial auctions via machine learning-based preference elicitation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, 2018.
- [Brero et al., 2019] G. Brero, S. Lahaie, and S. Seuken. Fast iterative combinatorial auctions via bayesian learning. In *Proceedings of the 33rd AAAI Conference of Artificial Intelligence*, 2019.
- [Brero et al., 2021] G. Brero, B. Lubin, and S. Seuken. Machine learning-powered iterative combinatorial auctions. *arXiv:1911.08042*, Jan 2021.
- [Cramton, 2013] P. Cramton. Spectrum auction design. *Review of Industrial Organization*, 42(2):161–190, 2013.
- [Dütting et al., 2019] P. Dütting, Z. Feng, H. Narasimhan, D. Parkes, and S. Ravindranath. Optimal auctions through deep learning. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- [Goeree and Holt, 2010] J. K. Goeree and C. A. Holt. Hierarchical package bidding: A paper & pencil combinatorial auction. *Games and Economic Behavior*, 2010.
- [Hudson and Sandholm, 2003] B. Hudson and T. Sandholm. Using value queries in combinatorial auctions. In *Proceedings of the 4th ACM Conference on Electronic commerce*, pages 226–227, 2003.
- [Lahaie and Parkes, 2004] S. Lahaie and C. Parkes. Applying learning algorithms to preference elicitation. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, 2004.
- [Lahaie, 2010] S. Lahaie. Kernel methods for revealed preference analysis. In *ECAI*, pages 439–444, 2010.
- [Milgrom, 2007] P. Milgrom. Package auctions and exchanges. *Econometrica*, 75(4):935–965, 2007.
- [Nisan and Segal, 2006] N. Nisan and I. Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 2006.
- [O’Donnell, 2014] R. O’Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.
- [Püschel and Wendler, 2020] M. Püschel and C. Wendler. Discrete signal processing with set functions. *IEEE Transactions on Signal Processing*, 69:1039–1053, 2020.
- [Rahaman et al., 2019] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, and A. Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, 2019.
- [Rahme et al., 2021] J. Rahme, S. Jelassi, J. Bruna, and M. Weinberg. A permutation-equivariant neural network architecture for auction design. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, 2021.
- [Scheffel et al., 2012] T. Scheffel, G. Ziegler, and M. Bichler. On the impact of package selection in combinatorial auctions: an experimental study in the context of spectrum auction design. *Experimental Economics*, 2012.
- [Shen et al., 2019] W. Shen, P. Tang, and S. Zuo. Automated mechanism design via neural networks. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 2019.
- [Shen et al., 2020] W. Shen, B. Peng, H. Liu, M. Zhang, R. Qian, Y. Hong, Z. Guo, Z. Ding, P. Lu, and P. Tang. Reinforcement mechanism design: With applications to dynamic pricing in sponsored search auctions. In *Proceedings of the 34th AAAI Conference*, 2020.
- [Stobbe and Krause, 2012] P. Stobbe and A. Krause. Learning Fourier sparse set functions. In *Artificial Intelligence and Statistics*, 2012.
- [Weiss et al., 2017] M. Weiss, B. Lubin, and S. Seuken. Sats: A universal spectrum auction test suite. In *Proceedings of the 16th Conference on Autonomous Agents and Multi-Agent Systems*, 2017.
- [Weissteiner and Seuken, 2020] Jakob Weissteiner and Sven Seuken. Deep learning-powered iterative combinatorial auctions. In *Proceedings of the 34th AAAI Conference of Artificial Intelligence*, 2020.
- [Weissteiner et al., 2022] Jakob Weissteiner, Jakob Heiss, Julien Siems, and Sven Seuken. Monotone-value neural networks: Exploiting preference monotonicity in combinatorial assignment. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence*, 2022.
- [Wendler et al., 2021] C. Wendler, A. Amrollahi, B. Seifert, A. Krause, and M. Püschel. Learning set functions that are sparse in non-orthogonal Fourier bases. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, 2021.