# Deep Learning-Powered Iterative Combinatorial Auctions 

Jakob Weissteiner<br>University of Zurich<br>weissteiner@ifi.uzh.ch

Sven Seuken<br>University of Zurich<br>seuken@ifi.uzh.ch


#### Abstract

In this paper, we study the design of deep learning-powered iterative combinatorial auctions (ICAs). We build on prior work where preference elicitation was done via kernelized support vector regressions (SVRs). However, the SVR-based approach has limitations because it requires solving a machine learning (ML)-based winner determination problem (WDP). With expressive kernels (like gaussians), the MLbased WDP cannot be solved for large domains. While linear or quadratic kernels have better computational scalability, these kernels have limited expressiveness. In this work, we address these shortcomings by using deep neural networks (DNNs) instead of SVRs. We first show how the DNNbased WDP can be reformulated into a mixed integer program (MIP). Second, we experimentally compare the prediction performance of DNNs against SVRs. Third, we present experimental evaluations in two medium-sized domains which show that even ICAs based on relatively small-sized DNNs lead to higher economic efficiency than ICAs based on kernelized SVRs. Finally, we show that our DNN-powered ICA also scales well to very large CA domains.


## 1 Introduction

Combinatorial auctions (CAs) are used to allocate multiple heterogeneous items to bidders in domains where these items may be substitutes or complements. Specifically, in a CA, bidders are allowed to submit bids on bundles of items rather than on individual items. CAs are widely used in practice, including for the sale of airport landing and take-off slots (Rassenti, Smith, and Bulfin 1982), in industrial procurement (Bichler et al. 2006), and for the sale of spectrum licenses (Cramton 2013).

One of the main challenges in large CAs is that the bundle space grows exponentially in the number of items. This typically makes it impossible for the bidders to report their full value function, even for medium-sized domains. Thus, careful preference elicitation is needed in CAs.

Nisan and Segal (2006) have shown that to achieve full efficiency and support general value functions, exponential communication in the number of items is needed in the worst case. Thus, practical auction designs cannot provide

[^0]efficiency guarantees in large CA domains. Instead, many recent proposals for CAs have focused on iterative combinatorial auctions (ICAs) where the auctioneer interacts with bidders over multiple rounds, eliciting a limited amount of information, aiming to find a highly efficient allocation.

ICAs have found widespread application in practice. For example, just between 2008 and 2014, the combinatorial clock auction (CCA) (Ausubel, Cramton, and Milgrom 2006) has been used to conduct more than 15 spectrum auctions and has generated more than $\$ 20$ Billion in total revenue (Ausubel and Baranov 2017). Another important application of ICAs are auctions for building offshore wind farms (Ausubel and Cramton 2011). Given the value of the resources allocated in these real-world ICAs, increasing their efficiency by $1-2 \%$ points already translates into welfare gains of millions of dollars. Therefore, improving the efficiency of ICAs is an important research challenge.

### 1.1 Machine Learning and Mechanism Design

Researchers have proposed various ways to further increase the efficiency of CAs by integrating machine learning (ML) methods into the mechanism. This research goes back to Blum et al. (2004) and Lahaie and Parkes (2004), who studied the relationship between computational learning theory and preference elicitation in CAs. More recently, Brero and Lahaie (2018) and Brero, Lahaie, and Seuken (2019) introduced a Bayesian CA where they integrated ML into a CA to achieve faster convergence. In a different strand of research, Dütting et al. (2015; 2017), Narasimhan, Agarwal, and Parkes (2016) and Golowich, Narasimhan, and Parkes (2018) used ML to directly learn a new mechanism (following the automated mechanism design paradigm).

Most related to the present paper is the work by Brero, Lubin, and Seuken (2017; 2018; 2019), who proposed an MLpowered ICA. The core of their auction is an ML-powered preference elicitation algorithm. As part of their algorithm, they used kernelized support vector regressions (SVRs) to learn the nonlinear value functions of bidders. Recently, Brero, Lubin, and Seuken (2019) showed that their MLbased ICA achieves even higher efficiency than the CCA. However, because of runtime complexity issues, Brero, Lubin, and Seuken (2018; 2019) focused on SVRs with linear
and quadratic kernels. This leaves room for improvement, since bidders' valuations can have more complex structures than can be captured by linear or quadratic kernels.

### 1.2 Our Approach Using Deep Neural Networks

In this paper, we propose using DNNs instead of SVRs in ML-powered ICAs. In each round of the auction, we approximate bidders' value functions by DNNs and subsequently solve an optimization problem, a $D N N$-based winner determination problem (WDP), to determine which query to ask every bidder in the next round. Since our design involves doing this in each round of the auction, a central requirement for the practical implementation of the auction mechanism is to efficiently solve these DNN-based WDPs. Therefore, we show how to reformulate the WDP based on DNNs with rectified linear units (ReLUs) as activation functions into a (linear) mixed integer program (MIP) (Section 4).

Our approach is related to a recent line of research that uses MIP formulations to study specific properties of DNNs. For example, Cheng, Nührenberg, and Ruess (2017) studied resilience properties of DNNs. Similarly, Fischetti and Jo (2018) used a MIP formulation for finding adversarial examples in image recognition.

To experimentally evaluate the performance of our DNNbased approach, we use the Spectrum Auction Test Suite (SATS) (Weiss, Lubin, and Seuken 2017) to generate synthetic CA instances (Section 5). We first compare the prediction performance of DNNs against SVRs in the two medium-sized domains GSVM and LSVM. Then we compare the economic efficiency of our DNN-powered ICA against the SVR-powered ICA. In GSVM (a domain perfectly suited for the quadratic kernel), our DNN-powered ICA matches the efficiency of the SVR-powered ICA, while in the more complex LSVM domain, our DNN-powered ICA outperforms the SVR-powered ICA by $1.74 \%$ points. Finally, we also demonstrate that our DNN-based approach scales well to a very large domain, by evaluating it in the MRVM domain (with 98 items and 10 bidders). Overall, our results show that, perhaps surprisingly, even small-sized neural networks can be advantageous for the design of ICAs.

## 2 Preliminaries

We now present our formal model and review the MLpowered ICA by Brero, Lubin, and Seuken (2018). ${ }^{1}$

### 2.1 Iterative Combinatorial Auction

We consider a CA setting with $n$ bidders and $m$ indivisible items. Let $N:=\{1, \ldots, n\}$ and $M:=\{1, \ldots, m\}$ denote the set of bidders and items, respectively. We denote by $x \in$ $\mathcal{X}:=\{0,1\}^{m}$ 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 $v_{i}:\{0,1\}^{m} \rightarrow \mathbb{R}_{+}, i \in$ $N$, i.e., $v_{i}(x)$ represents bidder $i$ 's true value for bundle $x$.

[^1]Let $v:=\left(v_{1}, \ldots, v_{n}\right)$ denote the vector of bidders' value functions. The (possibly untruthful) reported valuations are denoted by $\hat{v}_{i}$ and $\hat{v}$, respectively.

By $a:=\left(a_{1}, \ldots, a_{n}\right) \in \mathcal{X}^{n}$ we denote an allocation of bundles to bidders, where $a_{i} \in \mathcal{X}$ is the bundle bidder $i$ obtains. An allocation $a$ is feasible if each item is allocated to at most one bidder, i.e., $\forall j \in M$ : $\sum_{i \in N} a_{i j} \leq 1$. We denote the set of feasible allocations by $\mathcal{F}:=\left\{a \in \mathcal{X}^{n}: \sum_{i \in N} a_{i j} \leq 1, \forall j \in M\right\}$. Payments are denoted by $p=\left(p_{1}, \ldots, p_{n}\right) \in \mathbb{R}^{n}$, where $p_{i}$ is bidder $i$ 's payment. Furthermore, we assume that bidders have quasilinear utility functions $u_{i}(a):=v_{i}\left(a_{i}\right)-p_{i}$. The (true) social welfare of an allocation $a$ is defined as $V(a):=$ $\sum_{i \in N} v_{i}\left(a_{i}\right)$. Let $a^{*} \in \arg \max _{a \in \mathcal{F}} V(a)$ be a feasible, social-welfare maximizing, i.e., efficient, allocation given true value functions $v$. Then the efficiency of any feasible allocation $a \in \mathcal{F}$ is measured in terms of $a^{*}$ by $\frac{V(a)}{V\left(a^{*}\right)}$.

An ICA mechanism defines how the bidders interact with the auctioneer, how the final allocation is determined, and how payments are computed. In this paper, we only consider ICAs that ask bidders to iteratively report their valuations $\hat{v}_{i}(x)$ for particular bundles $x$ selected by the mechanism. A finite set of such reported bundle-value pairs of bidder $i$ is denoted as $B_{i}:=\left\{\left(x^{(k)}, \hat{v}_{i}\left(x^{(k)}\right)\right)\right\}_{k \in\left\{1, \ldots, n_{i}\right\}}, n_{i} \in \mathbb{N}, x^{(k)} \in \mathcal{X}$, where $n_{i}$ is the total number of bundle-value pairs reported by bidder $i$. We let $B:=\left(B_{1}, \ldots, B_{n}\right)$ denote the tuple of reported bundle-value pairs obtained from all bidders. We define the reported social welfare of an allocation $a$ given $B$ as

$$
\begin{equation*}
\widehat{V}(a \mid B):=\sum_{i \in N:\left(a_{i}, \hat{v}_{i}\left(a_{i}\right)\right) \in B_{i}} \hat{v}_{i}\left(a_{i}\right), \tag{1}
\end{equation*}
$$

where the condition $\left(a_{i}, \hat{v}_{i}\left(a_{i}\right)\right) \in B_{i}$ ensures that only values for reported bundles contribute to the sum. Finally, the optimal feasible allocation $a_{B}^{*}$ given $B$ is defined as

$$
\begin{equation*}
a_{B}^{*} \in \underset{a \in \mathcal{F}}{\arg \max } \widehat{V}(a \mid B) . \tag{2}
\end{equation*}
$$

In the ICA mechanisms we consider in this paper, the final outcome is only computed based on the reported values $B$ at termination. Specifically, the mechanism determines a feasible allocation $a_{B}^{*} \in \mathcal{F}$ and charges payments $p$.

As the auctioneer can generally only ask each bidder $i$ for a limited number of bundle-value pairs $B_{i}$, the ICA mechanism needs a sophisticated preference elicitation algorithm. This leads to the following preference elicitation problem, where the goal is to find an (approximately) efficient allocation with a limited number of value queries. More formally:
Problem 1 (Preference Elicitation in ICA). Given a cap $c_{e}$ on the number of value queries in an ICA, elicit from each bidder $i \in N$ a set of reported bundle-value pairs $B_{i}$ with $\left|B_{i}\right| \leq c_{e}$ such that the resulting efficiency of $a_{B}^{*}$ is maximized, i.e.,

$$
\begin{equation*}
B \in \underset{B:\left|B_{i}\right| \leq c_{e}}{\arg \max } \frac{V\left(a_{B}^{*}\right)}{V\left(a^{*}\right)} \tag{3}
\end{equation*}
$$

In practice, a small domain-dependent cap on the number of queries is chosen, e.g., $c_{e} \leq 500$.

### 2.2 SVR-powered ICA

We now present a brief review of the ML-based ICA introduced by Brero, Lubin, and Seuken (2018). At the core of their auction is an ML-based preference elicitation algorithm which we reprint here as Algorithm 1.

```
Algorithm 1: ML-BASED Elicitation (Brero et al.
2018)
    Parameter: Machine learning algorithm \(\mathcal{A}\)
    \(B^{0}=\) initial tuple of reported bundle-value pairs at
    \(t=0\)
    do
        \(t \leftarrow t+1\)
        Estimation step: \(\tilde{V}^{t}:=\mathcal{A}\left(B^{t-1}\right)\)
        Optimization step: \(a^{(t)} \in \underset{a \in \mathcal{F}}{\arg \max } \tilde{V}^{t}(a)\)
        for each bidder \(i\) do
            if \(a_{i}^{(t)} \notin B_{i}^{t-1}\) then
                Query value \(\hat{v}_{i}\left(a_{i}^{(t)}\right)\)
                \(B_{i}^{t}=B_{i}^{t-1} \cup\left\{\left(a_{i}^{(t)}, \hat{v}_{i}\left(a_{i}^{(t)}\right)\right)\right\}\)
            else
                \(B_{i}^{t}=B_{i}^{t-1}\)
            end
        end
    while \(\exists i \in N: a_{i}^{(t)} \notin B_{i}^{t-1}\)
    Output tuple of reported bundle-value pairs \(B^{t}\)
```

This algorithm is a procedure to determine $B$, i.e., for each bidder $i$ a set of reported bundle-value pairs $B_{i}$. Note that the algorithm is described in terms of a generic ML algorithm $\mathcal{A}$ which is used in the estimation step (Line 4) to obtain the estimated social welfare function $\tilde{V}^{t}$ in iteration $t$. In the optimization step (Line 5), an ML-based winner determination problem is then solved to find an allocation $a^{(t)}$ that maximizes $\tilde{V}^{t}$. Finally, given the allocation $a^{(t)}$ from iteration $t$, each bidder $i$ is asked to report his value for the bundle $a_{i}^{(t)}$. The algorithm stops when it reaches an allocation $a^{(t)}$ for which all bidders have already reported their values for the corresponding bundles $a_{i}^{(t)}$.

As the ML-algorithm $\mathcal{A}$, Brero, Lubin, and Seuken (2018) used a sum of kernelized SVRs, i.e,

$$
\begin{equation*}
\mathcal{A}\left(B^{t-1}\right):=\sum_{i \in N} \mathrm{SVR}_{i} \tag{4}
\end{equation*}
$$

Given a bundle $x$, each $\mathrm{SVR}_{i}$ computes the predicted value as $\operatorname{SVR}_{i}(x)=w_{i} \cdot \phi(x)$, where the weights $w_{i}$ are determined through training on the reported bundle-value pairs $B_{i}^{t-1}$. Kernelized SVRs are a popular non-linear regression technique, where a linear model is fitted on transformed data. The transformation of bundles $x$ is implicitly conducted by setting a kernel $k\left(x, x^{\prime}\right):=\phi(x)^{T} \phi\left(x^{\prime}\right)$ in the dual optimization problem (Smola and Schölkopf 2004).

Brero, Lubin, and Seuken (2018) called their entire auction mechanism the Pseudo Vickrey-Clarke-Groves mechanism (PVM). We reprint it here as Algorithm 2. PVM calls

## Algorithm 2: PVM (Brero et al. 2018)

1 Run Algorithm $1 n+1$ times:
$B^{(-\emptyset)}, B^{(-1)}, \ldots, B^{(-n)}$.
Determine allocations: $a^{(-\emptyset)}, a^{(-1)}, \ldots, a^{(-n)}$, where
$a^{(-i)} \in \arg \max _{a \in \mathcal{F}} \widehat{V}\left(a \mid B^{(-i)}\right)$.
Pick $a^{p v m} \in\left\{a^{(-\emptyset)}, a^{(-1)}, \ldots, a^{(-n)}\right\}$ with maximal $\widehat{V}$.
4 Charge each bidder $i$ according to:

$$
\begin{equation*}
p_{i}^{p v m}:=\sum_{j \neq i} \hat{v}_{j}\left(a_{j}^{(-i)}\right)-\sum_{j \neq i} \hat{v}_{j}\left(a_{j}^{p v m}\right) . \tag{5}
\end{equation*}
$$

the preference elicitation algorithm (Algorithm 1) $n+1$ times: once including all bidders (called the main economy) and $n$ times excluding a different bidder in each run (called the marginal economies). The motivation for this design, which is inspired by the VCG mechanism, is to obtain payments such that the auction aligns bidders' incentives with allocative efficiency. Here, $B^{(-i)}$ denotes the output of Algorithm 1 by excluding bidder $i$ from the set of bidders. For each of the reported bundle-value pairs $B^{(-i)}$ obtained from the $n+1$ runs, PVM calculates a corresponding allocation that maximizes the reported social welfare (Line 2). The final allocation $a^{p v m}$ is determined as the allocation of the $n+1$ runs with the largest reported social welfare (Line 3). Finally, VCG-style payments are calculated (Line 4).

## 3 Deep Neural Network-powered ICA

In this section, we present the high level design of our DNNpowered ICA and discuss its advantages compared to the SVR-based design by Brero, Lubin, and Seuken (2018).

Observe that the choice of the ML algorithm $\mathcal{A}$ affects Algorithm 1 in two ways: first, in the estimation step (Line 4), $\mathcal{A}$ determines how well we can predict bidders' valuations; second, in the optimization step (Line 5), it determines the complexity of the ML-based WDP. Thus, our situation is different from standard supervised learning because of the added optimization step. In particular, when choosing $\mathcal{A}$, we must also ensure that we obtain a practically feasible MLbased WDP. Given that we have to solve the optimization step hundreds of times throughout an auction, in practice, we must impose a time limit on this step. In our experiments (Section 5), we follow Brero, Lubin, and Seuken (2018) and impose a 1 hour time limit on this step.

To make the optimization step feasible, Brero, Lubin, and Seuken (2018) used SVRs with quadratic kernels, for which the ML-based WDP is a quadratic integer program (QIP) and still practically solvable within a 1 hour time limit for most settings. However, note that a quadratic kernel, while more expressive than a linear kernel, can still at most model two-way interactions between the items. To this end, Brero, Lubin, and Seuken $(2017 ; 2019)$ also evaluated more expressive kernels (gaussian and exponential). Even though these kernels had good prediction performance, the corresponding ML-based WDPs were too complex such that they always
timed out and had a large optimization gap, thus leading to worse economic efficiency than the quadratic kernel. However, using SVRs with quadratic kernels leaves room for improvement, since bidders' valuations can be more complex than can be captured by quadratic kernels.

In this work, we show how these shortcomings can be addressed by using DNNs instead of SVRs in the estimation and optimization steps of Algorithm 1. DNNs are a concatenation of affine and non-linear mappings (see Figure 1). They consist of several layers, which are themselves composed of multiple nodes. Between each of the layers an affine transformation is applied, which is followed by a nonlinear mapping called the activation function.

One advantage of DNNs compared to (nonlinear) kernelized SVRs is that, for any number of layers and nodes, we always obtain a (linear) MIP for the DNN-based WDP whose size only grows linearly in the number of bidders and items (as we will show in Section 4). The key insight for this is to use rectified linear units (ReLUs) as activation functions. Furthermore, in contrast to SVRs, DNNs do not use predefined feature transformations. While with SVRs, the choice of a good kernel usually relies on prior domain knowledge, DNNs automatically learn features in the process of training.

Following Brero, Lubin, and Seuken (2018), we decompose the estimated social welfare function in Line 4 of Algorithm 1, $\tilde{V}^{t}=\mathcal{A}\left(B^{t-1}\right)$, as follows:

$$
\begin{equation*}
\tilde{V}^{t}=\sum_{i \in N} \tilde{v}_{i}^{t} \tag{6}
\end{equation*}
$$

where $\tilde{v}_{i}^{t}$ is an estimate of bidder $i$ 's true value function $v_{i}$ and is trained on the data set $B_{i}^{t-1}$, i.e., the values queried up to round $t-1$. In this work, for every $i \in N$, we model $\tilde{v}_{i}^{t}$ using a fully connected feed-forward $D N N \mathcal{N}_{i}:\{0,1\}^{m} \rightarrow \mathbb{R}_{+}$. Consequently, the estimated social welfare function $\tilde{V}^{t}$ is given as a sum of DNNs, i.e.,

$$
\begin{equation*}
\tilde{V}^{t}:=\sum_{i \in N} \mathcal{N}_{i} . \tag{7}
\end{equation*}
$$

Note that each bidder's value function is modeled as a distinct DNN (with different architectures and parameters) because bidders' preferences are usually highly idiosyncratic. ${ }^{2}$

## 4 MIP Formulation of the DNN-based Winner Determination Problem

We now present the parameterization of each DNN and show how to reformulate the DNN-based WDP into a MIP. Thus, we focus on the optimization step (Line 5) of Algorithm 1.

### 4.1 Setting up the DNN-based WDP

Figure 1 shows a schematic representation of a DNN $\mathcal{N}_{i}$. We now define the parameters of the DNNs. To simplify the exposition, we consider a fixed iteration step $t$ and no longer highlight the dependency of all variables on $t$.

[^2]

Figure 1: Schematic representation of a DNN $\mathcal{N}_{i}$.

Each DNN $\mathcal{N}_{i}$ consists of $K_{i}-1$ hidden layers for $K_{i} \in \mathbb{N}$, with the $k^{\text {th }}$ hidden layer containing $d_{k}^{i}$ hidden nodes, where $k \in\left\{1, \ldots, K_{i}-1\right\}$. As $\mathcal{N}_{i}$ maps bundles $x \in\{0,1\}^{m}$ to values, the dimension of the input layer $d_{0}^{i}$ is equal to the number of items $m$ (i.e., $d_{0}^{i}:=m$ ) and the dimension of the output layer is equal to 1 (i.e., $d_{K_{i}}^{i}:=1$ ). Hence, in total, a single DNN consists of $K_{i}+1$ layers. Furthermore, let $\varphi: \mathbb{R} \rightarrow \mathbb{R}_{+}, \varphi(s):=\max (0, s)$ denote the ReLU activation function. ${ }^{3}$ The affine mappings between the $k^{\text {th }}$ and the $k+1^{\text {st }}$ layer are parameterized by a matrix $W^{i, k} \in$ $\mathbb{R}^{d_{k+1}^{i} \times d_{k}^{i}}$ and a bias $b^{i, k} \in \mathbb{R}^{d_{k+1}^{i}}, k \in\left\{0, \ldots, K_{i}-1\right\}$.

To estimate the parameters $W^{i, k}$ and $b^{i, k}$ from data (i.e., from the bundle-value pairs $B_{i}$ ) we use the $A D A M$ algorithm, which is a popular gradient-based optimization technique (Kingma and Ba 2015 ). ${ }^{4}$ This is done in the estimation step in Line 4 of Algorithm 1. Thus, after the estimation step, $W^{i, k}$ and $b^{i, k}$ are constants. In summary, given estimated parameters $W^{i}:=\left\{W^{i, k}\right\}_{0 \leq k \leq K_{i}-1}$ and $b^{i}:=$ $\left\{b^{i, k}\right\}_{0 \leq k \leq K_{i}-1}$, each DNN $\mathcal{N}_{i}\left(W^{i}, b^{i}\right):\{0,1\}^{m} \rightarrow \mathbb{R}_{+}$ represents the following nested function:

$$
\begin{align*}
& \mathcal{N}_{i}\left(W^{i}, b^{i}\right)(x)=  \tag{8}\\
& =\varphi\left(W^{i, K_{i}-1} \varphi\left(\ldots \varphi\left(W^{i, 0} x+b^{i, 0}\right) \ldots\right)+b^{i, K_{i}-1}\right)
\end{align*}
$$

The $D N N$-based $W D P$ in the optimization step in Line 5 of Algorithm 1 can now be formulated as follows:

$$
\begin{array}{ll} 
& \max _{a \in \mathcal{X}^{n}}\left\{\sum_{i \in N} \mathcal{N}_{i}\left(W^{i}, b^{i}\right)\left(a_{i}\right)\right\}  \tag{OP1}\\
\text { s.t. } & \sum_{i \in N} a_{i j} \leq 1, \quad \forall j \in M \\
& a_{i j} \in\{0,1\}, \quad \forall j \in M, \forall i \in N
\end{array}
$$

### 4.2 The MIP Formulation

In its general form, (OP1) is a nonlinear, non-convex optimization problem and there do not exist practically feasible algorithms that are guaranteed to find a globally optimal solution. Therefore, we now reformulate (OP1) into a MIP.

[^3]Consider bidder $i \in N$. For every layer $k \in\left\{1, \ldots, K_{i}\right\}$ let $o^{i, k} \in \mathbb{R}_{+}^{d_{k}^{i}}$ denote the output of the $k^{\text {th }}$ layer, which can be recursively calculated as

$$
\begin{align*}
o^{i, k} & =\varphi\left(W^{i, k-1} o^{i, k-1}+b^{i, k-1}\right)= \\
& =\max \left(0, W^{i, k-1} o^{i, k-1}+b^{i, k-1}\right) \tag{9}
\end{align*}
$$

For $k \in\left\{1, \ldots, K_{i}\right\}$, we introduce $d_{k}^{i}$ binary decision variables that determine which node in the corresponding layer is active, represented as a vector $y^{i, k} \in\{0,1\}^{d_{k}^{i}}$. We also introduce $2 d_{k}^{i}$ continuous variables, represented as vectors $z^{i, k}, s^{i, k} \in \mathbb{R}^{d_{k}^{i}}$. Each $z^{i, k}$ corresponds to the positive components of the output value $o^{i, k}$ of each layer and each $s^{i, k}$ is used as a slack variable representing the absolute value of the negative components of $o^{i, k}$.

In our final MIP formulation, we will make use of "big$M^{\prime \prime}$ constraints. For our theoretical results to hold, we need to make the following standard assumption.
Assumption 1. (Big-M CONSTRAINT) For all $i \in N$ and $k \in\left\{1, \ldots, K_{i}\right\}$ there exists a large enough constant $L \in \mathbb{R}_{+}$, such that $\left|\left(W^{i, k-1} o^{i, k-1}+b^{i, k-1}\right)_{j}\right| \leq L$ for $1 \leq j \leq d_{k}^{i}$.

In the following lemma, we show how to recursively encode a layer of $\mathcal{N}_{i}$ given the output value of the previous layer as multiple linear constraints. ${ }^{5}$
Lemma 1. Let $W^{i, k-1} o^{i, k-1}+b^{i, k-1} \in \mathbb{R}^{d_{k}^{i}}$ be fixed for a $k \in\left\{1, \ldots, K_{i}\right\}$. Furthermore, let $z^{i, k}, s^{i, k} \in \mathbb{R}^{d_{k}^{i}}$, $y^{i, k} \in\{0,1\}^{d_{k}^{i}}$. Consider the following linear constraints:

$$
\begin{align*}
& z^{i, k}-s^{i, k}=W^{i, k-1} o^{i, k-1}+b^{i, k-1}  \tag{10}\\
& 0 \leq z^{i, k} \leq y^{i, k} \cdot L  \tag{11}\\
& 0 \leq s^{i, k} \leq\left(1-y^{i, k}\right) \cdot L \tag{12}
\end{align*}
$$

The polytope defined by (10)-(12) is not empty and every element $\left(z^{i, k}, s^{i, k}, y^{i, k}\right)$ of this polytope satisfies $z^{i, k}=o^{i, k}$.

Proof. For notational convenience, let $c:=W^{i, k-1} o^{i, k-1}+$ $b^{i, k-1}$. Non-emptiness follows since $\left|c_{j}\right| \leq L, 1 \leq j \leq d_{k}^{i}$, by Assumption 1. From Constraints (11) and (12) it follows, for $1 \leq j \leq d_{k}^{i}$, that if $z_{j}^{i, k}>0$ then $s_{j}^{i, k}=0$, and if $s_{j}^{i, k}>0$ then $z_{j}^{i, k}=0$. We now have to distinguish the following three cases for each component of $c$ :

$$
\begin{aligned}
& c_{j}<0 \Longrightarrow y_{j}^{i, k}=0, s_{j}^{i, k}=-c_{j}, z_{j}^{i, k}=0=\varphi\left(c_{j}\right) \\
& c_{j}>0 \Longrightarrow y_{j}^{i, k}=1, s_{j}^{i, k}=0, z_{j}^{i, k}=c_{j}=\varphi\left(c_{j}\right) \\
& c_{j}=0 \Longrightarrow\left(z_{j}^{i, k}, s_{j}^{i, k}, y_{j}^{i, k}\right) \in\{(0,0,0),(0,0,1)\}
\end{aligned}
$$

Combining all cases yields that $z^{i, k}=\varphi\left(c_{j}\right)=o^{i, k}$.
Given Lemma 1, we can now reformulate the DNN-based WDP as a MIP. For this, we let $W^{i}, b^{i}$ denote the estimated parameters corresponding to $\mathcal{N}_{i}\left(W^{i}, b^{i}\right)$. Furthermore, let $a \in \mathcal{X}^{n}, y^{i, k} \in\{0,1\}^{d_{k}^{i}}$ and $z^{i, k}, s^{i, k} \in \mathbb{R}^{d_{k}^{i}}$, for $1 \leq k \leq K_{i}$ and $L$ be a constant satisfying Assumption 1.

[^4]\[

$$
\begin{equation*}
\max _{a \in \mathcal{X}^{n}, z^{i, k}, s^{i, k}, y^{i, k}}\left\{\sum_{i \in N} z^{i, K_{i}}\right\} \tag{OP2}
\end{equation*}
$$

\]

s.t.

$$
\left.\begin{array}{l}
z^{i, 0}=a_{i} \\
z^{i, k}-s^{i, k}=W^{i, k-1} z^{i, k-1}+b^{i, k-1} \\
0 \leq z^{i, k} \leq y^{i, k} \cdot L \\
0 \leq s^{i, k} \leq\left(1-y^{i, k}\right) \cdot L \\
y^{i, k} \in\{0,1\}^{d_{k}^{i}} \\
a_{i j} \in\{0,1\}, \quad \forall j \in M, \forall i \in N \\
\sum_{i \in N} a_{i j} \leq 1, \quad \forall j \in M
\end{array}\right\} \begin{aligned}
& \\
& \\
& \\
& \\
& \\
& \\
& \\
&
\end{aligned} \quad \forall k \in\left\{1, \ldots, K_{i}\right\}
$$

We are now ready to state our main theorem.
Theorem 1. (MIP FORMULATION) The DNN-based WDP as defined in (OP1) is equivalent to the MIP defined in (OP2).

Proof. Consider (OP1). For each bidder $i \in N$, we first set $z^{i, 0}$ equal to the input bundle $a_{i}$. Then we proceed by using Lemma 1 for $k=1$, i.e., we reformulate the output of the $1^{\text {st }}$ layer as the linear constraints (10), (11) and (12). We iterate this procedure until we have reformulated the final layer, i.e, $k=K_{i}$. Doing so for each bidder $i \in N$ and adding the feasibility constraints yields (OP2).

Using the MIP formulation (OP2), we can solve the DNN-based WDP using standard optimization packages like Cplex. ${ }^{6}$
We now provide a simple worked example for how to reformulate (OP1) into (OP2) which illustrates Theorem 1.
Example 1. Consider a setting with one bidder $(n=1)$, $m$ items $\left(d_{0}^{1}=m\right)$ and one hidden layer $\left(K_{1}=2\right)$. Given $W^{1,0} \in \mathbb{R}^{d_{1}^{1} \times m}, W^{1,1} \in \mathbb{R}^{1 \times d_{1}^{1}}, b^{1,0} \in \mathbb{R}^{d_{1}^{1}}$, and $b^{1,1} \in \mathbb{R}$, (OP1) can be written as

$$
\begin{aligned}
& \max _{a_{1} \in \mathcal{X}^{1}}\left\{\max \left(0, W^{1,1} \max \left(0, W^{1,0} a_{1}+b^{1,0}\right)+b^{1,1}\right)\right\} \\
& \text { s.t. } \quad a_{1 j} \in\{0,1\}, \quad \forall j \in M,
\end{aligned}
$$

where the constraint $\sum_{i \in N} a_{i j} \leq 1, \forall j \in M$ is redundant in this case, since we only consider one bidder. First, we replace the inner maximum using $y^{1,1} \in\{0,1\}^{d_{1}^{1}}, z^{1,1}, s^{1,1} \in$ $\mathbb{R}^{d_{1}^{1}}$ and arrive at the equivalent optimization problem

$$
\begin{array}{ll} 
& \max _{\substack{a_{1} \in \mathcal{X}^{1} \\
z^{1,1}, s^{1,1}, y^{1,1}}}\left\{\max \left(0, W^{1,1} z^{1,1}+b^{1,1}\right)\right\} \\
\text { s.t. } \quad z^{1,1}-s^{1,1}=W^{1,0} a_{1}+b^{1,0} \\
& 0 \leq z^{1,1} \leq y^{1,1} \cdot L \\
& 0 \leq s^{1,1} \leq\left(1-y^{1,1}\right) \cdot L \\
& a_{1 j} \in\{0,1\}, \quad \forall j \in M
\end{array}
$$

[^5]Applying Lemma 1 again by using $y^{1,2} \in\{0,1\}$ and $z^{1,2}, s^{1,2} \in \mathbb{R}$ yields the final MIP formulation

$$
\begin{array}{ll} 
& \max _{\substack{a_{1} \in \mathcal{X}^{1}, 1^{1, k}}}\left\{z^{1,2}\right\} \\
y^{1, k}, s^{1, k}, k \in\{1,2\} \\
\text { s.t. } & z^{1,1}-s^{1,1}=W^{1,0} a_{1}+b^{1,0} \\
& 0 \leq z^{1,1} \leq y^{1,1} \cdot L \\
& 0 \leq s^{1,1} \leq\left(1-y^{1,1}\right) \cdot L \\
& z^{1,2}-s^{1,2}=W^{1,1} z^{1,1}+b^{1,1} \\
& 0 \leq z^{1,2} \leq y^{1,2} \cdot L \\
& 0 \leq s^{1,2} \leq\left(1-y^{1,2}\right) \cdot L \\
& a_{1 j} \in\{0,1\}, \quad \forall j \in M .
\end{array}
$$

Remark 1. The number of decision variables in the MIP defined in (OP2) is given by

$$
\sum_{\text {\#bidders }}^{\sum_{i \in N}}(\underbrace{m}_{\text {\#items }}+\underbrace{3}_{y^{i, k}, s^{i, k}, z^{i, k}} \cdot(\underbrace{\text { \#hidden layers }}_{\sum_{k=1}^{\sum_{i}-1}} \underbrace{d_{k}^{i}}_{\text {\#nodes per layer }}+\overbrace{1}^{\text {output }}))
$$

## 5 Experimental Evaluation

In this section, we evaluate the performance of our deep learning-powered ICA and compare it against the SVRbased approach using quadratic kernels by Brero, Lubin, and Seuken (2018). We release our code under an open-source license at: https://github.com/marketdesignresearch/DL-ICA.

### 5.1 Experiment setup

Spectrum auctions are one of the most prominent applications of CAs, which is why we choose them for our experiments. Specifically, we use the spectrum auction test suite (SATS) version 0.6.4 (Weiss, Lubin, and Seuken 2017). ${ }^{7}$ SATS enables us to generate 1000s of CA instances in different domains. Furthermore, we have access to each bidder's true value $v_{i}(x)$ for all $2^{m}$ possible bundles $x \in \mathcal{X}$ as well as the efficient allocation $a^{*} \in \mathcal{F}$, which we can use to measure the efficiency of any other allocation $a$ by $V(a) / V\left(a^{*}\right)$. We evaluate our approach in the following three domains:

The Global Synergy Value Model (GSVM) (Goeree and Holt 2010) consists of 6 regional bidders, 1 national bidder, and 18 items. In GSVM the value of a package increases by a certain percentage with every additional item of interest. Thus, the value of a bundle only depends on the total number of items contained in a bundle which makes it one of the simplest models in SATS. In fact, bidders' valuations can be exactly learned by SVRs with quadratic kernels (Brero, Lubin, and Seuken (2019)) which implies that the valuations exhibit at most two-way interactions between items.

The Local Synergy Value Model (LSVM) (Scheffel, Ziegler, and Bichler 2012) consists of 5 regional bidders, 1 national bidder and 18 items. The items are arranged on a rectangle of size $3 \times 6$. The national bidder is interested in all

[^6]items, while the regional bidders are only interested in certain subsets of items. Complementarities arise from spatial proximity of items and are modeled via a logistic function, which makes it more complex than GSVM.
The Multi-Region Value Model (MRVM) (Weiss, Lubin, and Seuken 2017) consists of 98 items and 10 bidders. It models large US and Canadian spectrum auctions and captures both geography (different regions) as well as frequency dimensions (different bands). A bidder is categorized as national, regional or local, depending on the magnitude of the synergies between different regions.
In Sections 5.2 and 5.3 , we first evaluate our approach in detail using the two medium-sized domains GSVM and LSVM. Then, in Section 5.4, we use MRVM to evaluate how well our approach scales to very large domains.

### 5.2 Prediction Performance

We first compare the prediction performance of DNNs to SVRs. Using SATS, we generate a data set of bundle-value pairs $\left\{\left(x^{(k)}, v_{i}\left(x^{(k)}\right)\right)\right\}$ for all bidders $i \in N$. For each auction domain, we draw 100 auction instances uniformly at random. For each such instance, we sample, for each bidder type, a training set $T$ of equal size and a disjoint test set $V$ consisting of all remaining bundles, i.e., $|V|:=2^{|M|}-|T|$. For each bidder type, we train the ML algorithm on $T$ and test it on $V$. We report the mean absolute error (MAE) for both bidder types averaged over the 100 instances. ${ }^{8}$
We denote by $\left[d_{1}, d_{2}, d_{3}\right]$ a 3 -hidden-layer DNN with $d_{1}, d_{2}$ and $d_{3}$ hidden nodes, respectively. For both, SVRs and DNNs, we performed a hyperparameter optimization for each bidder type. For the DNNs, we optimized the architecture ${ }^{9}$, the L2-regularization parameter for the affine mappings, the dropout rate per layer, and the learning rate of ADAM. For SVRs with a quadratic kernel $k(x, y):=$ $x^{T} y+\gamma\left(x^{T} y\right)^{2}$, we optimized $\gamma$ (i.e., the influence of the quadratic term), the regularization parameter $C$, and the loss function parameter $\epsilon$. In what follows, we present the winner models resulting from this hyperparameter optimization.

In Table 1, we present prediction performance results for the GSVM domain. Consider the last column of the table, which shows the MAE on the test set. We observe the very good prediction performance of the SVRs and in particular that the test error converges to 0 when increasing $|T|$. This is due to the fact that in GSVM, bidders' value functions can be perfectly captured by quadratic kernels. In this sense, GSVM represents a "worst case" auction domain w.r.t our comparison of DNNs against quadratically-kernelized SVRs. Looking at the performance of the DNNs, we observe that the test error also decreases with $|T|$, but, not surprisingly, is always larger than for the SVRs with quadratic kernels. Furthermore, we observe that the optimal architectures are always a 1 -hidden layer network.

In Table 2, we present the results for the more complex LSVM domain. We observe that for $|T|=50$, the DNNs and

[^7]| ML Algorithm | $\|T\|$ | Bidder Type | DNN <br> Architecture | $\mathrm{MAE}_{\text {train }}$ | $\mathrm{MAE}_{\text {test }}$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| DNNs | 50 | National | [100] | 1.99 | 5.25 (0.11) |
|  |  | Regional | [100] | 2.06 | 6.20 (0.19) |
|  | 100 | National | [100] | 2.10 | 3.66 (0.07) |
|  |  | Regional | [100] | 2.71 | 4.64 (0.13) |
|  | 200 | National | [100] | 1.61 | 2.22 (0.05) |
|  |  | Regional | [100] | 2.11 | 2.89 (0.09) |
|  |  |  | Kernel |  |  |
| SVRs | 50 | National | quadratic | 0.03 | 4.38 (0.11) |
|  |  | Regional | quadratic | 0.03 | 4.98 (0.20) |
|  | 100 | National | quadratic | 0.03 | 1.71 (0.04) |
|  |  | Regional | quadratic | 0.03 | 2.07 (0.07) |
|  | 200 | National | quadratic | 0.03 | 0.12 (0.00) |
|  |  | Regional | quadratic | 0.03 | 0.13 (0.00) |

Table 1: Prediction performance in GSVM. All results are averaged over 100 auction instances. For $\mathrm{MAE}_{\text {test }}$, standard errors are shown in parentheses.

| ML Algorithm | $\|T\|$ | Bidder Type | DNN <br> Architecture | $\mathrm{MAE}_{\text {train }}$ | $\mathrm{MAE}_{\text {test }}$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| DNNs | 50 | National | [10] | 24.68 | 29.90 (0.23) |
|  |  | Regional | [100] | 4.22 | 16.58 (0.39) |
|  | 100 | National | [10, 10, 10] | 9.01 | 25.62 (0.36) |
|  |  | Regional | [100] | 5.01 | 13.74 (0.26) |
|  | 200 | National | [10, 10] | 10.52 | 20.58 (0.21) |
|  |  | Regional | [100, 100, 100] | 3.64 | 11.27 (0.23) |
| Kernel |  |  |  |  |  |
| SVRs | 50 | National | quadratic | 18.51 | 32.61 (0.59) |
|  |  | Regional | quadratic | 3.11 | 15.30 (0.34) |
|  | 100 | National | quadratic | 20.03 | 27.86 (0.28) |
|  |  | Regional | quadratic | 3.21 | 14.21 (0.28) |
|  | 200 | National | quadratic | 20.03 | 25.44 (0.16) |
|  |  | Regional | quadratic | 8.23 | 12.67 (0.26) |

Table 2: Prediction performance in LSVM. All results are averaged over 100 auction instances. For $\mathrm{MAE}_{\text {test }}$, standard errors are shown in parentheses.

SVRs with quadratic kernels have similar test error. But for $|T|=100$ and $|T|=200$, the DNNs significantly outperform the SVRs with quadratic kernels. Specifically, DNNs better capture the national bidder in LSVM, which is important, since this bidder is interested in all items and usually gets a large portion of the items in the final allocation, which matters a lot for efficiency. In contrast to GSVM, we observe that for $|T| \geq 100$, multi-hidden-layer networks were found to be best. This suggests that DNNs may indeed be advantageous for capturing more complex preference structures. ${ }^{10}$

### 5.3 Efficiency Results

Finally, we compare the economic efficiency of our DNNpowered ICA against the SVR-powered ICA. When conducting the efficiency experiments, we follow Brero, Lubin, and Seuken (2018) and assume that bidders answer all value queries truthfully. Furthermore, we also use their experiment

[^8]| DNN Architectures ${ }^{11}$ | $c_{0}$ | Efficiency $\%$ | Revenue $\%^{12}$ |
| :--- | :---: | :---: | :---: |
| $\mathrm{R}:[16,16] \mid \mathrm{N}:[10,10]$ | 40 | $98.53 \%$ | $69.26 \%$ |
| $\mathrm{R}:[16,16] \mid \mathrm{N}:[10,10]$ | 30 | $98.41 \%$ | $68.29 \%$ |
| $\mathrm{R}:[16,16] \mid \mathrm{N}:[10,10,10]$ | 40 | $98.51 \%$ | $68.91 \%$ |
| $\mathrm{R}:[16,16] \mid \mathrm{N}:[10,10,10]$ | 30 | $98.32 \%$ | $68.59 \%$ |
| $\mathrm{R}:[32,32] \mid \mathrm{N}:[10,10]$ | 40 | $98.75 \%$ | $71.14 \%$ |
| R:[32,32] $\mid \mathrm{N}:[10,10]$ | $\mathbf{3 0}$ | $\mathbf{9 8 . 9 4 \%}$ | $\mathbf{6 8 . 4 7 \%}$ |
| $\mathrm{R}:[32,32] \mid \mathrm{N}:[10,10,10]$ | 40 | $98.69 \%$ | $71.63 \%$ |
| $\mathrm{R}:[32,32] \mid \mathrm{N}:[10,10,10]$ | 30 | $98.92 \%$ | $68.88 \%$ |
| $\mathrm{R}:[100] \mid \mathrm{N}:[100]$ | 30 | $98.27 \%$ | $66.73 \%$ |

Table 3: Efficiency results for 9 configurations of a DNNpowered ICA on a training set of 100 GSVM auction instances. The selected winner model is marked in bold. All results are averaged over the 100 auction instances.

| Auction <br> Mechanism | \#Queries | Max <br> \#Queries | $\%$ <br> Efficiency | $\%$ <br> Revenue | t-test on <br> Efficiency ${ }^{13}$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| VCG | $2^{18}$ | $2^{18}$ | $100.00(0.00)$ | 80.4 | - |
| SVR-ICA | 41.9 | 42.8 | $\mathbf{9 8 . 8 5}(0.13)$ | 77.80 | 0.337 |
| DNN-ICA | 53 | 78 | $\mathbf{9 8 . 6 3}(0.18)$ | 67.81 |  |

Table 4: A comparison of the DNN-powered ICA against the SVR-powered ICA and VCG (as reported in Brero, Lubin, and Seuken (2018)) on a test set of 100 GSVM instances. All results are averaged over the 100 instances. For efficiency, standard errors are shown in parentheses.
setup and define a cap $c_{e}$ on the total number of queries in Algorithm 1 and set $c_{e}:=50$. The initial set of reported bundle-value pairs $B_{i}^{0}$ per bidder $i$ is drawn uniformly at random and set to be equal across bidders. We denote the number of initial reports by $c_{0}:=\left|B_{i}^{0}\right|, \forall i \in N$, resulting in a maximum of $c_{0}+n \cdot\left(c_{e}-c_{0}\right)$ queries per bidder.

GSVM. In Table 3, we first present the results from comparing nine different network architectures on a training set of 100 GSVM instances. As we can see, the winning model is among the largest multi-layer networks we tested. It is noteworthy that the one-hidden-layer network (R:[100]|N:[100]), which performed best in terms of prediction performance, did not perform best in terms of efficiency.

In Table 4, we compare the performance of the winner model from Table 3 against the SVR-based approach on a test set of 100 GSVM instances. Even though GSVM can be perfectly captured by quadratic kernels, our DNN-based approach achieves a similar result w.r.t to efficiency, where the difference is not statistically significant ( $p=0.337$ ).

LSVM. We now turn to the more complex LSVM domain. We first select a winner model based on a training set of 100 LSVM instances (Table 5). As in GSVM, the best model

[^9]| DNN Architectures | $c_{0}$ | Efficiency $\%$ | Revenue \% |
| :--- | :---: | :---: | :---: |
| $\mathrm{R}:[16,16] \mid \mathrm{N}:[10,10]$ | 40 | 97.40 | 60.51 |
| $\mathrm{R}:[16,16] \mid \mathrm{N}:[10,10]$ | 30 | 96.87 | 56.85 |
| $\mathrm{R}:[16,16] \mid \mathrm{N}:[10,10,10]$ | 40 | 97.45 | 61.15 |
| $\mathrm{R}:[16,16] \mid \mathrm{N}:[10,10,10]$ | 30 | 97.12 | 59.31 |
| $\mathrm{R}:[32,32] \mid \mathrm{N}:[10,10]$ | 40 | 97.40 | 62.01 |
| $\mathrm{R}:[32,32] \mid \mathrm{N}:[10,10]$ | 30 | 96.83 | 59.07 |
| R:[32,32] \| N:[10,10,10] | $\mathbf{4 0}$ | $\mathbf{9 7 . 7 4}$ | $\mathbf{6 1 . 9 5}$ |
| $\mathrm{R}:[32,32] \mid \mathrm{N}:[10,10,10]$ | 30 | 97.12 | 59.56 |
| $\mathrm{R}:[100] \mid \mathrm{N}:[10]$ | 40 | 96.78 | 58.71 |

Table 5: Efficiency results for 9 configurations of a DNNpowered ICA on a training set of 100 LSVM auction instances. The selected winner model is marked in bold. All results are averaged over the 100 auction instances.

| Auction <br> Mechanism | \#Queries | Max <br> \#Queries | $\%$ <br> Efficiency | $\%$ <br> Revenue | t-test on <br> Efficiency |
| :---: | :---: | :---: | :---: | :---: | :---: |
| VCG | $2^{18}$ | $2^{18}$ | $100.00(0.00)$ | 83.4 | - |
| SVR-ICA | 48.2 | 52.8 | $\mathbf{9 6 . 0 3}(0.33)$ | 65.60 | 0.000038 |
| DNN-ICA | 65 | 77 | $\mathbf{9 7 . 7 4}(0.24)$ | 62.45 |  |

Table 6: A comparison of the DNN-powered ICA against the SVR-powered ICA and VCG (as reported in Brero, Lubin, and Seuken (2018)) on a test set of 100 LSVM auction instances. All results are averaged over the 100 instances. For efficiency, standard errors are shown in parentheses.
is among the largest architectures and the best model w.r.t prediction performance does not yield the highest efficiency.

In Table 6, we compare the performance of the selected winner model from Table 5 against the SVR-based approach on a test set of 100 new auction instances. Here, we see that our DNN-powered ICA substantially outperforms the SVR-powered ICA by $1.74 \%$ points, and this effect is highly statistically significant ( $p=0.000038$ ). This demonstrates the advantage of DNNs over SVRs with quadratic kernels in complex domains like LSVM.

In Figure 2, we present a histogram of the efficiency obtained by the selected winner model on the test set. We see that for 29 auction instances, our approach (impressively) obtains an economic efficiency of $100 \%$. However, for two instances, the efficiency is less than $90 \%$. Thus, it is a promising avenue for future work to investigate these outliers to further increase the average efficiency.


Figure 2: Histogram of efficiency results in LSVM of the selected DNN winner model on the test set.

Remark 2. In Tables 4 and 6, we see that our DNN-based approach achieves lower revenue than the SVR-based approach. This may be explained as follows. A bidder's payment in PVM, depends on the difference between the social welfare in the marginal and the main economy. However, PVM has one oddity: a bidder's payment may be negative. This happens more frequently with DNNs than with SVRs: consider an auction where bidder $i$ is not allocated in the main economy. Then, ideally, the allocation (and thus the welfare) should be exactly the same in the marginal economy where bidder $i$ is excluded, resulting in a zero payment. When using SVRs, this is guaranteed if the same set of bundle-value pairs is used in the main and marginal economy, because SVRs use a deterministic algorithm for estimation. In contrast, DNNs use a non-deterministic algorithm, sometimes resulting in different allocations in the main and marginal economies. However, this is more a limitation of PVM itself. In practice, one should lower bound the payments as also suggested by Brero, Lubin, and Seuken (2019). Lower-bounding all payments by zero increases the revenue in GSVM by $7.9 \%$ points and in LSVM by $8.3 \%$ points.

### 5.4 Scaling to Larger Domains

We now present results for the MRVM domain (with 98 items and 10 bidders) to show that our DNN-powered ICA also scales well to larger domains (we present detailed runtime results in Section 5.5). We use the experiment setup of Brero, Lubin, and Seuken (2018) and set $c_{e}:=100$.

In Table 7, we present the results for different DNN architectures and different values of $c_{0}$, evaluated on a training set of MRVM instances. First, we observe that the efficiency increases as we decrease $c_{0}$. This can be explained by the fact

| DNN Architectures | $c_{0}$ | Efficiency $\%$ | Revenue \% |
| :--- | :---: | :---: | :---: |
| $\mathrm{L}:[10,10]\|\mathrm{R}:[32,32]\| \mathrm{N}:[32,32]$ | 70 | 93.54 | 31.02 |
| $\mathrm{~L}:[10,10]\|\mathrm{R}:[32,32]\| \mathrm{N}:[32,32]$ | 50 | 94.07 | 33.51 |
| $\mathrm{~L}:[16,16]\|\mathrm{R}:[16,16]\| \mathrm{N}:[16,16]$ | 30 | 94.46 | 31.39 |
| $\mathrm{~L}:[10,10]\|\mathrm{R}:[32,32]\| \mathrm{N}:[32,32]$ | 30 | 94.73 | 31.88 |
| $\mathrm{~L}:[16,16]\|\mathrm{R}:[16,16]\| \mathrm{N}:[16,16]$ | 20 | 94.88 | 30.31 |
| $\mathrm{~L}:[10,10]\|\mathrm{R}:[32,32]\| \mathrm{N}:[32,32]$ | 20 | 94.42 | 34.23 |
| $\mathrm{~L}:[16,16]\|\mathrm{R}:[16,16]\| \mathrm{N}:[16,16]$ | $\mathbf{1 0}$ | $\mathbf{9 5 . 0 0}$ | $\mathbf{3 1 . 9 7}$ |
| $\mathrm{L}:[10,10]\|\mathrm{R}:[32,32]\| \mathrm{N}:[32,32]$ | 10 | 94.54 | 34.78 |
| $\mathrm{~L}:[10,10]\|\mathrm{R}:[16,16,16]\| \mathrm{N}:[16,16,16]$ | 10 | 94.74 | 31.12 |

Table 7: Efficiency results for 9 configurations of a DNNpowered ICA on a training set of 19 MRVM auction instances. The selected winner model is marked in bold. All results are averaged over the 19 auction instances.

| Auction <br> Mechanism | \#Queries | Max <br> \#Queries | $\%$ <br> Efficiency | $\%$ <br> Revenue | t-test on <br> Efficiency |
| :---: | :---: | :---: | :---: | :---: | :---: |
| VCG | $2^{98}$ | $2^{98}$ | $100.00(0.00)$ | 44.3 | - |
| SVR-ICA | 265 | 630 | $\mathbf{9 4 . 5 8}(0.14)$ | 35.20 | 0.0268 |
| DNN-ICA | 334 | 908 | $\mathbf{9 5 . 0 4}(0.14)$ | 30.59 |  |

Table 8: A comparison of the DNN-powered ICA against the SVR-powered ICA and VCG (as reported in Brero, Lubin, and Seuken (2018)) on a test set of 50 MRVM auction instances. All results are averaged over the 50 instances. For efficiency, standard errors are shown in parentheses.
that a smaller $c_{0}$ tends to lead to more iterations of the preference elicitation algorithm, resulting in a larger number of elicited bundle-value pairs. In terms of which DNN architectures performed better or worse, no clear pattern emerged.

In Table 8, we compare the performance of the selected winner model from Table 7 against the SVR-based approach on a test set of 50 MRVM instances. We see that our DNN-powered ICA outperforms the SVR-powered ICA by $0.46 \%$ points. While this is a relatively modest increase in efficiency, this effect is statistically significant ( $p=$ 0.0268 ). We also observe that our DNN-based approach (while obeying all caps) asks a larger number of queries than the SVR-based approach. It is unclear how much of the efficiency increase is driven by the DNN or by the larger number of queries. Future work should compare the two approaches by holding the total number of queries constant.

### 5.5 Runtime Analysis

| Domain | $\varnothing$ MIP Runtime | $\varnothing$ Iteration Runtime | $\varnothing$ Auction Runtime |
| :---: | :---: | :---: | :---: |
| GSVM | 15.90 sec | 30.51 sec | 44 min |
| LSVM | 39.75 sec | 51.69 sec | 65 min |
| MRVM | 3.67 sec | 26.75 sec | 457 min |

Table 9: Average runtime results of the selected DNN winner models in different domains. All values are averaged over 100 (GSVM and LSVM) and 50 (MRVM) auction instances.

In Table 9, we present runtime results for our DNNpowered ICA for the winner models from Tables 4, 6 and 8. Specifically, we show average runtime results of the MIP (OP2), of an iteration of Algorithm 1, and of a whole auction (PVM). We observe that the average runtime of a whole auction takes approximately 1 hour in GSVM and LSVM and 7 hours in the larger MRVM domain. The increase in total runtime in MRVM can be explained by the fact that we use a smaller number of initial queries $\left(c_{0}:=10\right)$ and a larger total query cap ( $c_{e}:=100$ ) compared to LSVM and GSVM. This results in a larger number of iterations of Algorithm 1. Additionally, MRVM consists of 10 bidders resulting in 11 calls of Algorithm 1 in contrast to 7 calls in LSVM and GSVM. Even though in MRVM the average MIP runtime is smaller, the larger number of iterations and bidders lead to this increase in total runtime. Overall, these results show that our DNN-based approach is practically feasible and scales well to the larger MRVM domain. Brero, Lubin, and Seuken (2018) do not provide runtime information such that we cannot provide a runtime comparison with SVRs. ${ }^{14}$

In Figure 3, we present additional MIP runtime results for selected DNN architectures in LSVM (results in GSVM and MRVM are qualitatively similar). We observe two effects: First, increasing the number of nodes per layer slightly increases the average runtime. Second, adding an additional

[^10]

Figure 3: MIP runtimes defined in (OP2) based on 50 different LSVM instances. Results are shown for a selection of various DNN architectures and for $c_{0}:=40, c_{e}:=50$.
layer (for the national bidder) significantly increases the average runtime. Not surprisingly, the largest DNN architectures lead to the highest runtime.

The runtime of our MIPs heavily depends on the size of the "big-M" variable $L$. In practice, $L$ should be chosen as small as possible to obtain a MIP formulation that is as tight as possible. We initialized $L:=5000$ and tightened this bound further by using interval arithmetic (see, e.g., Tjeng, Xiao, and Tedrake (2019)). Recently, Singh et al. (2018) proposed a novel technique for tightening such MIP formulations. Evaluating this in more detail is subject to future work.

## 6 Conclusion

In this paper, we have designed a deep learning-powered ICA. We have compared our approach against prior work using SVRs with quadratic kernels. Our experimental results have shown that our DNN-based approach leads to significantly higher economic efficiency in complex auction domains and scales well to large domains.

On a technical level, our main contribution was to reformulate the DNN-based WDP into a MIP. The main insight to achieve this was to use ReLU activation functions, which can be re-written as multiple linear constraints. From an experimental point of view, we were pleasantly surprised to see that even DNNs with a small number of layers and nodes and with a small number of training samples (i.e., bids) were able to achieve high prediction performance and ultimately high economic efficiency in the overall auction mechanism.

Future work should investigate the trade-off between larger DNN architectures and the resulting MIP runtime, to further increase efficiency.

## Acknowledgments

We thank Gianluca Brero, Nils Olberg, and Stefania Ionescu for insightful discussions and 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).

## References

Ausubel, L. M., and Baranov, O. 2017. A practical guide to the combinatorial clock auction. Economic Journal 127(605):F334F350.

Ausubel, L., and Cramton, P. 2011. Auction design for wind rights. Report to Bureau of Ocean Energy Management, Regulation and Enforcement.
Ausubel, L. M.; Cramton, P.; and Milgrom, P. 2006. Combinatorial Auctions. MIT Press. chapter The clock-proxy auction: A practical combinatorial auction design, 115-138.
Bichler, M.; Davenport, A.; Hohner, G.; and Kalagnanam, J. 2006. Combinatorial Auctions. MIT Press. chapter Industrial procurement auctions, 593-612.

Blum, A.; Jackson, J.; Sandholm, T.; and Zinkevich, M. 2004. Preference elicitation and query learning. Journal of Machine Learning Research 5:649-667.
Brero, G., and Lahaie, S. 2018. A bayesian clearing mechanism for combinatorial auctions. In Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence.
Brero, G.; Lahaie, S.; and Seuken, S. 2019. Fast iterative combinatorial auctions via bayesian learning. In Proceedings of the Thirty-third AAAI Conference of Artificial Intelligence.
Brero, G.; Lubin, B.; and Seuken, S. 2017. Probably approximately efficient combinatorial auctions via machine learning. In ThirtyFirst AAAI Conference on Artificial Intelligence.
Brero, G.; Lubin, B.; and Seuken, S. 2018. Combinatorial auctions via machine learning-based preference elicitation. In Proceedings of the 27th International Joint Conference on Artificial Intelligence and the 23 rd European Conference on Artificial Intelligence.
Brero, G.; Lubin, B.; and Seuken, S. 2019. Machine learning-powered iterative combinatorial auctions. arXiv preprint https://arxiv.org/abs/1911.08042.
Cheng, C.-H.; Nührenberg, G.; and Ruess, H. 2017. Automated Technology for Verification and Analysis. ATVA 2017. Lecture Notes in Computer Science, volume 10482. Springer, Cham. chapter Maximum resilience of artificial neural networks.
Cramton, P. 2013. Spectrum auction design. Review of Industrial Organization 42(2):161-190.
Dütting, P.; Fischer, F.; Jirapinyo, P.; Lai, J. K.; Lubin, B.; and Parkes, D. C. 2015. Payment rules through discriminant-based classifiers. ACM Transactions on Economics and Computation 3(1):5.
Dütting, P.; Feng, Z.; Narasimhan, H.; Parkes, D. C.; and Ravindranath, S. S. 2017. Optimal auctions through deep learning. In Proceedings of the 36th International Conference on Machine Learning.
Fischetti, M., and Jo, J. 2018. Deep neural networks and mixed integer linear optimization. Constraints 23(3):296-309.
Goeree, J. K., and Holt, C. A. 2010. Hierarchical package bidding: A paper \& pencil combinatorial auction. Games and Economic Behavior 70(1):146-169.
Golowich, N.; Narasimhan, H.; and Parkes, D. C. 2018. Deep learning for multi-facility location mechanism design. In Proceedings of the Twenty-seventh International Joint Conference on Artificial Intelligence and the Twenty-third European Conference on Artificial Intelligence, 261-267.
Kingma, D., and Ba, J. 2015. Adam: A method for stochastic optimization. In Proceedings of the 3rd International Conference on Learning Representations.
Lahaie, S. M., and Parkes, D. C. 2004. Applying learning algorithms to preference elicitation. In Proceedings of the 5th ACM Conference on Electronic Commerce.
Narasimhan, H.; Agarwal, S. B.; and Parkes, D. C. 2016. Automated mechanism design without money via machine learning. In

Proceedings of the 25th International Joint Conference on Artificial Intelligence.
Nisan, N., and Segal, I. 2006. The communication requirements of efficient allocations and supporting prices. Journal of Economic Theory 129(1):192-224.
Rassenti, S. J.; Smith, V. L.; and Bulfin, R. L. 1982. A combinatorial auction mechanism for airport time slot allocation. The Bell Journal of Economics 402-417.
Scheffel, T.; Ziegler, G.; and Bichler, M. 2012. On the impact of package selection in combinatorial auctions: an experimental study in the context of spectrum auction design. Experimental Economics 15(4):667-692.
Singh, G.; Gehr, T.; Mirman, M.; Püschel, M.; and Vechev, M. 2018. Fast and effective robustness certification. In Advances in Neural Information Processing Systems, 10802-10813.
Smola, A. J., and Schölkopf, B. 2004. A tutorial on support vector regression. Statistics and computing 14(3):199-222.
Tjeng, V.; Xiao, K. Y.; and Tedrake, R. 2019. Evaluating robustness of neural networks with mixed integer programming. In International Conference on Learning Representations.
Weiss, M.; Lubin, B.; and Seuken, S. 2017. Sats: A universal spectrum auction test suite. In Proceedings of the 16th Conference on Autonomous Agents and Multi-Agent Systems.


[^0]:    Copyright © $C^{2020}$, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[^1]:    ${ }^{1}$ We compare our DNN-powered ICA against the mechanism described in (Brero, Lubin, and Seuken 2018) because, when we wrote this paper, (Brero, Lubin, and Seuken 2019) was not available yet. We slightly adopt the notation and use $B_{i}$ instead of $\widehat{\vartheta}_{i}$.

[^2]:    ${ }^{2}$ A second reason for this construction is that this prevents bidders from influencing each others' ML-models. Please see (Brero, Lubin, and Seuken 2019) for a detailed incentive analysis of PVM.

[^3]:    ${ }^{3} \varphi$ acts on vectors componentwise, i.e., for $s \in \mathbb{R}^{k}$ let $\varphi(s):=\left(\varphi\left(s_{1}\right), \ldots, \varphi\left(s_{k}\right)\right)$.
    ${ }^{4}$ For fitting the DNNs in all of our experiments we use PYTHON 3.5.3, KERAS 2.2.4 and TENSORFLOW 1.13.1.

[^4]:    ${ }^{5}$ In the following, all constraints containing vectors are defined componentwise.

[^5]:    ${ }^{6}$ For solving MIPs of the form (OP2) in our experiments we use Cplex 12.8.0.0 with the python library DOcplex 2.4.61.

[^6]:    ${ }^{7}$ Experiments were conducted on machines with Intel Xeon E52650 v 42.20 GHz processors with 40 logical cores.

[^7]:    ${ }^{8}$ For training the DNNs, we use the MAE as the loss function.
    ${ }^{9}$ We considered the following architectures: for the national bidders: [10], [100], [10,10], ..., [100,100,100,100], and for the regional bidders: [32], [100], [32,32], . . ., [100,100,100,100].

[^8]:    ${ }^{10}$ We also evaluated the prediction performance of other kernels (linear, gaussian and exponential) and observed that DNNs were as good or better for almost all combinations of bidder types and $|T|$.

[^9]:    ${ }^{11}$ We denote by R and N the architectures used for the regionaland national bidders, respectively.
    ${ }^{12}$ Revenue is calculated as $\left(\sum_{i \in N} p_{i}^{p v m}\right) / V\left(a^{*}\right)$.
    ${ }^{13}$ We performed a two-sided unpaired Welch Two Sample t-test with $\mathcal{H}_{0}: \mu_{1}=\mu_{2}$ against $\mathcal{H}_{A}: \mu_{1} \neq \mu_{2}$. We thank Brero, Lubin, and Seuken (2018) for providing us with the detailed results of their experiments to enable all t-tests reported in this paper.

[^10]:    ${ }^{14}$ In conversations with the authors, they told us that for gaussian and exponential kernels, their MIPs always timed out (1h cap) in GSVM, LSVM and MRVM. The average MIP runtime for the quadratic kernel was a few seconds in GSVM and LSVM. In MRVM, the quadratic kernel also regularly timed out resulting in an average MIP runtime of 30 min and of 36 h for a whole auction.

