



CIRANO
Knowledge into action

IMPLEMENTING A HIERARCHICAL DEEP LEARNING APPROACH FOR SIMULATING MULTILEVEL AUCTION DATA

IGOR SADOUNE
MARCELIN JOANIS
ANDREA LODI

2023s-23
WORKING PAPER

CS

The purpose of the **Working Papers** is to disseminate the results of research conducted by CIRANO research members in order to solicit exchanges and comments. These reports are written in the style of scientific publications. The ideas and opinions expressed in these documents are solely those of the authors.

Les cahiers de la série scientifique visent à rendre accessibles les résultats des recherches effectuées par des chercheurs membres du CIRANO afin de susciter échanges et commentaires. Ces cahiers sont rédigés dans le style des publications scientifiques et n'engagent que leurs auteurs.

CIRANO is a private non-profit organization incorporated under the Quebec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the government of Quebec, and grants and research mandates obtained by its research teams.

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du gouvernement du Québec, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO Partners – Les partenaires du CIRANO

Corporate Partners – Partenaires corporatifs

Autorité des marchés financiers
Bank of Canada
Bell Canada
BMO Financial Group
Business Development Bank of Canada
Caisse de dépôt et placement du Québec
Desjardins Group
Énergir
Hydro-Québec
Innovation, Science and Economic Development Canada
Intact Financial Corporation
Manulife Canada
Ministère de l'Économie, de l'Innovation et de l'Énergie
Ministère des finances du Québec
National Bank of Canada
Power Corporation of Canada
PSP Investments
Ville de Montréal

Academic Partners – Partenaires universitaires

Concordia University
École de technologie supérieure
École nationale d'administration publique
HEC Montréal
McGill University
National Institute for Scientific Research
Polytechnique Montréal
Université de Montréal
Université de Sherbrooke
Université du Québec
Université du Québec à Montréal
Université Laval

CIRANO collaborates with many centers and university research chairs; list available on its website. *Le CIRANO collabore avec de nombreux centres et chaires de recherche universitaires dont on peut consulter la liste sur son site web.*

© September 2023. Igor Sadoune, Marcelin Joanis and Andrea Lodi. All rights reserved. *Tous droits réservés.* Short sections may be quoted without explicit permission, if full credit, including © notice, is given to the source. *Reproduction partielle permise avec citation du document source, incluant la notice ©.*

The observations and viewpoints expressed in this publication are the sole responsibility of the authors; they do not represent the positions of CIRANO or its partners. *Les idées et les opinions émises dans cette publication sont sous l'unique responsabilité des auteurs et ne représentent pas les positions du CIRANO ou de ses partenaires.*

Implementing a Hierarchical Deep Learning Approach for Simulating multilevel Auction Data

Igor Sadoune^{}, Marcelin Joanis[†] and Andrea Lodi[‡]*

Abstract/Résumé

We present a deep learning solution to address the challenges of simulating realistic synthetic first-price sealed-bid auction data. The complexities encountered in this type of auction data include high-cardinality discrete feature spaces and a multilevel structure arising from multiple bids associated with a single auction instance. Our methodology combines deep generative modeling (DGM) with an artificial learner that predicts the conditional bid distribution based on auction characteristics, contributing to advancements in simulation-based research. This approach lays the groundwork for creating realistic auction environments suitable for agent-based learning and modeling applications. Our contribution is twofold: we introduce a comprehensive methodology for simulating multilevel discrete auction data, and we underscore the potential of DGM as a powerful instrument for refining simulation techniques and fostering the development of economic models grounded in generative AI.

Nous proposons une solution basée sur l'apprentissage profond pour simuler de manière réaliste des données d'enchères scellées. Les enjeux liés à ce type de données résident dans la gestion des variables discrètes de grande dimension et de la structure multiniveau liée à la présence de multiples offres pour une seule et même enchère. Notre approche intègre une modélisation générative profonde avec un système d'apprentissage artificiel, capable de prévoir la distribution des offres en fonction des propriétés de l'enchère. Cette stratégie constitue une base solide pour l'élaboration d'environnements d'enchères artificiels mais réalistes, adaptés à l'apprentissage et à la modélisation basés sur les agents. Notre contribution est double: nous introduisons une méthodologie complète pour simuler des données d'enchères discrètes à plusieurs niveaux, et nous mettons en lumière le potentiel de la modélisation générative profonde pour améliorer les techniques de simulation et promouvoir le développement de modèles économiques s'appuyant sur l'intelligence artificielle générative.

Keywords/Mots-clés: simulation crafting, discrete deep generative modeling, multilevel discrete data, auction data / simulation, modélisation générative discrète et profonde, données discrètes multiniveaux, données d'enchères

To quote this document

Sadoune, I., Joanis, M., & Lodi, A. (2023). Implementing a Hierarchical Deep Learning Approach for Simulating multilevel Auction Data (2023s-23, Working Papers, CIRANO.) <https://doi.org/10.54932/LQOG8430>

* Department of Mathematics and Industrial Engineering, Polytechnique Montreal, Montreal, QC, Canada, CIRANO
Corresponding author: igor.sadoune@polymtl.ca

† Department of Mathematics and Industrial Engineering, Polytechnique Montreal, Montreal, QC, Canada, CIRANO

‡ Jacobs Technion-Cornell Institute, Cornell Tech and Technion - IIT, New York, NY, USA

These authors contributed equally to this work: marcelin.joanis@polymtl.ca, andrea.lodi@cornell.edu

1 Introduction

In this paper, we propose a hierarchical deep learning method for generating synthetic yet realistic first-price auction data. Our approach is designed to address the challenges associated with high-cardinality discrete feature spaces and multilevel structures inherent to auction data. By leveraging the capabilities of generative deep learning, our primary contribution is to provide the right methodology for crafting realistic auction simulations, which in turn can facilitate the application of agent-based modeling (ABM) and complex system approaches to the study of auction markets.

Auction markets, as a prominent example of complex systems, have gained significant attention from researchers and practitioners due to their inherent complexity and potential applications across various fields, such as finance, economics, and operations research [1]. The study of first-price auctions, in particular, is crucial for understanding the mechanisms and dynamics that govern their functioning, which are prevalent in numerous industries and applications, including online advertising [2], electricity markets [3], and public procurement in general [4].

We argue that simulations, ABM, and the complex system approach, in general, may offer advantages over static statistical analysis by providing a more detailed representation of individual interactions and capturing the emergent properties of complex systems [5, 6, 7]. The integration of deep learning techniques, such as deep generative modeling, with the crafting of realistic auction simulations is expected to promote more robust and resilient economic systems, ultimately benefiting various fields, including computational economics [8, 9, 10].

The method developed in this paper utilizes generative adversarial networks (GANs) [11] and variational autoencoders (VAEs) [12] for the replication of the auction feature space, while employing a neural network, referred to as Bidnet, to predict the conditional bid distribution for each instance of auction from the generated feature space. Recent advancements in deep generative modeling (DGM), such as the development of CTGAN [13], have enabled the efficient handling of high-cardinality discrete distributions, which is a critical aspect for the generation of realistic auction data. Additionally, our Bidnet architecture is specifically designed to address the multilevel issue inherent to auction data, by capturing the relationships between auctions and their associated bids. This combination of DGM and Bidnet provides a comprehensive and effective approach for simulating first-price auction data, while accounting for the complexities and challenges associated with high-cardinality discrete feature spaces and multilevel structures.

1.1 Paper Organization

Section 2 delves into the mathematical background of GANs and VAEs, discussing their advantages over traditional methods for data generation. This section also explores the application of deep generative modeling in social sciences and its use in high-dimensional discrete spaces.

In Section 3, we detail the process of generating synthetic multilevel auction data. This section first presents the multilevel problem and the problem formulation, followed by the solution framework. We then discuss approximating auction features joint density using GAN-based approaches and tabular variational encoding. The section continues with the training of a generator for continuous bids and concludes with the sampling of synthetic auction instances.

Section 4 focuses on the validation of the synthetic data generated by our proposed methods. This section evaluates the faithfulness of the synthetic auction features and the performance of the BidNet in generating synthetic bids.

In Section 5, we present a discussion that summarizes the study's findings, emphasizing the credibility and usefulness of our contribution. This section also discusses the limitations of our approach and highlights the superiority of CTGANs in our specific context.

Finally, the algorithms used and developed in the paper are detailed in the Appendix.

2 Deep Generative Modeling

Deep generative modeling includes a variety of techniques that aim to perform density estimation using artificial neural networks as function approximators. Certain models can obtain an explicit representation of the target distribution, while others provide only an implicit, or "black box," representation of the data structure to be replicated. In this study, we focus on the latter category of models, specifically GANs [11] and VAEs [12]. The strength of these models lies in their ability to eliminate the necessity for detailed knowledge of the underlying data structure being replicated. In essence, GANs and VAEs possess the capability to generate an extensive array of synthetic data points from a limited number of empirical observations.


Adversarial learning is the core concept in GANs, which consist of two neural networks, a generator (G) and a discriminator (D), that compete in a two-player minimax game. The generator's objective is to generate synthetic samples $G(z)$, where z is a random noise vector, that are similar to the true data distribution $p_{data}(x)$. The discriminator's goal is to differentiate between real samples (x) from the true data distribution

and synthetic samples generated by the generator. The discriminator assigns a probability value $D(x)$ to each input x . The learning process can be described by the following objective function:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{data}(x)}[\log D(x)] + \mathbb{E}_{z \sim p_z(z)}[\log(1 - D(G(z)))] \quad (1)$$

During training, the generator and discriminator optimize this objective function in a sequential way. The generator seeks to minimize the function while the discriminator attempts to maximize it. This adversarial process leads to a convergence where the generator produces samples that the discriminator can no longer differentiate from real data samples. As the training progresses, the generator becomes increasingly skilled at producing realistic samples, while the discriminator improves its ability to distinguish between real and synthetic samples. When the equilibrium is reached, the generator generates synthetic samples that resemble the true data distribution, and the discriminator is unable to distinguish between real and generated samples, assigning a probability of $\frac{1}{2}$ to each input.

Figure 2.1: Chart of adversarial learning (left) and variational autoencoding (right).



VAEs are a type of generative model that leverages variational inference to learn a probabilistic mapping between the data and a latent space. The key idea behind VAEs is to introduce a probabilistic encoder $q_\phi(z|x)$, that approximates the true posterior distribution $p_\theta(z|x)$, and a probabilistic decoder $p_\theta(x|z)$, that models the data distribution conditioned on the latent variable z . Here, ϕ and θ represent the parameters of the encoder and decoder neural networks, respectively. Variational inference is used to optimize the evidence lower bound (ELBO) on the log-likelihood of the data, i.e.,

$$\log p_\theta(x) \geq \mathbb{E} q_\phi(z|x)[\log p_\theta(x|z)] - D_{KL}(q_\phi(z|x)||p_\theta(z)) = \mathcal{L}(\theta, \phi; x). \quad (2)$$

The ELBO consists of two terms: the reconstruction term, $\mathbb{E} q_\phi(z|x)[\log p_\theta(x|z)]$, which measures the ability of the model to reconstruct the data given the latent variable, and the regularization term, $D_{KL}(q_\phi(z|x)||p_\theta(z))$, which measures the divergence between the approximate posterior and the prior distribution over the latent space. During training, the VAE jointly optimizes the encoder and decoder parameters to maximize the ELBO with respect to the model parameters θ and ϕ . This process encourages the learned latent space to have a meaningful structure, enabling efficient generation of new samples by sampling from the prior $p_\theta(z)$ and decoding them using the decoder $p_\theta(x|z)$.

2.1 Traditional Methods for Data Generation

Market simulations and simulation crafting have been widely used across various disciplines to study complex systems, replicate real-world dynamics, and develop policies and strategies. The use of market simulations can be traced back to the early works of economists and computer scientists [14]. Agent-based modeling has emerged as a prominent approach to studying markets and economic systems, allowing researchers to capture the interactions between heterogeneous agents and their influence on market outcomes [5]. A large body of

literature has explored different aspects of market simulations, from designing auction mechanisms [15] to investigating the dynamics of financial markets [16].

Historically, traditional methods for simulation crafting, such as Monte Carlo sampling, bootstrapping, cellular automata, and system dynamics modeling, have been widely used in various applications [17, 5, 18]. Monte Carlo sampling generates random samples from a given probability distribution, cellular automata simulates spatial and temporal dynamics through discrete cells evolving based on predefined rules, and system dynamics modeling uses differential equations to describe relationships among system components.

However, these traditional methods exhibit limitations when applied to complex systems with high-cardinality discrete spaces or multilevel structures, such as auction markets. These limitations include difficulties in accurately representing high-dimensional probability distributions, modeling intricate dependencies between variables, accommodating nonlinear dynamics and non-stationary behaviors, and integrating domain-specific knowledge or constraints [19]. Furthermore, traditional simulation methods often require substantial computational resources and are constrained by their assumptions and rules [6].

The advent of machine learning and artificial intelligence has significantly impacted simulation crafting, paving the way for advanced techniques such as DGM. Deep generative models including GANs and VAEs, address the limitations of traditional techniques by learning and representing high-dimensional probability distributions, modeling complex dependencies between variables, and accommodating nonlinearities and non-stationarities in complex systems [20]. These innovative methods have been applied to various domains, such as finance [21]. Although GANs have been recognized as having potential for creating more realistic market simulations and fostering the development of more effective policies and strategies [22], the use of GANs in economics remains limited, with only a small number of applications currently found in the field.

Leveraging unsupervised learning and deep neural networks, DGM methods identify and represent the underlying data structure in a flexible and scalable manner. Additionally, DGM techniques can model complex dependencies and correlations between variables, providing a more accurate representation of the relationships present in the data. DGM techniques suitable for handling the nonlinearities and non-stationarities that are inherent to complex systems. Due to their deep architectures and non-linear activation functions, deep generative models can capture complex, multi-scale structures in data, which may be challenging for traditional methods to accurately represent. Moreover, DGM approaches can be adapted to incorporate domain-specific knowledge and constraints, further enhancing the realism and validity of the generated synthetic data.

2.2 DGM in Social Sciences

Although deep generative models are still underused in social sciences, they can potentially provide substantial improvements in any application relying on qualitative data.

They have emerged as a powerful tool for addressing low degrees of freedom. Specifically, DGMs amplify observations and enhance weak signals within categorical configurations, enabling the generation of synthetic data that closely approximates the characteristics of real-world data. This technology has found applications in a diverse range of domains, such as credit card datasets [23] and medical data recovery [24].

In causal effect research, DGMs play a pivotal role by enabling counterfactual analysis using propensity scores and facilitating the estimation of unseen or partially unseen distributions through GANs [25] and VAEs [26, 27]. Additionally, DGMs have been shown to benefit privacy-sensitive applications involving medical or financial data, leading to specialized literature in this area [28].

Moreover, DGMs have contributed to the reinvigoration of agent-based modeling in economics by generating synthetic data for creating realistic artificial environments and maintaining simulation realism [21]. For example, DGM-based travel behavior simulations have employed restricted Boltzmann machines [29], while GAN-based financial correlation matrices and time-series sampling have been used for simulating

financial systems [16, 30].

2.3 GANs for High-Dimensional Discrete Spaces

Despite the progress in market simulations and simulation crafting, several challenges remain. In our case, generating high-cardinality discrete data and multilevel data structures is still a complex task, as highlighted by [13]. Furthermore, incorporating the dynamics of real-world markets and their ever-evolving nature into simulation models requires continuous research and development. Various strategies have been developed to enhance the stability and performance of GANs, particularly focusing on innovations that address the challenges of handling discrete inputs.

For instance, the Wasserstein GAN (WGAN) [31, 32] and the least-square GAN (LSGAN) [33] employ a critic rather than a traditional discriminator. The critic predicts the distance between a given point and the decision boundary separating real and fake samples, thereby improving the signal quality for the generator during training and leading to better convergence properties. WGAN's popularity stems from its Earth Mover (EM) or Wasserstein-I loss, which measures the distance between real and fake sample distributions, thus promoting stable and robust training [31].

The boundary-seeking GAN (BGAN) [34] introduces a modified training process, where the generator uses a policy gradient that accommodates both discrete and continuous inputs. This approach results in smoother training and mitigates issues related to mode collapse, a common problem in GANs.

Meanwhile, the conditional tabular GAN (CTGAN) [13] is a framework specifically designed to address data imbalances and discrete inputs in tabular data. Functioning as a meta-algorithm, CTGAN is compatible with various loss functions, network topologies, or training processes. The core concept of CTGAN, training by sampling, augments the generator's input space with a conditional vector that encodes the selection, thereby enhancing the model's ability to capture the underlying structure of discrete data and generate realistic synthetic samples.

3 Generating Synthetic multilevel Auction Data

In this study, we aim to create realistic synthetic auctions using a novel method that combines two distinct functions, each handling a specific aspect of the auction data. The first function generates artificial auction characteristics, capturing the features of the contracts being offered. The second function approximates the distribution of bids given these auction features, providing an aggregated representation of the bidding firms. By breaking down the process into these two sequential functions, our method addresses technical challenges while maintaining the overall structure of the auction. We validate our approach by training a predictor using the synthetic data and evaluating its performance on real-world data. This method effectively separates the generation of discrete and continuous data types and simplifies the complex, multilevel auction structure, making it suitable for generating realistic market simulations.

We have chosen the context of public procurement to demonstrate our ability to simulate realistic first-price auctions effectively. To this end, we rely on the data provided by the "Système électronique d'appel d'offre" (SEAO), which encompasses 119,029 contracts offered in public market auctions within the Canadian province of Quebec over a ten-year period (2010-2020). In this scenario, the high-cardinality discrete feature space represents various auction attributes, such as the type of contract or the designated delivery region. The multilevel structure of the auction data emerges from the multiple firms participating in the bidding process for each auction. Utilizing this comprehensive dataset allows us to develop and refine our simulation methodology, thereby showcasing its potential for generating accurate and realistic representations of first-price auctions in the domain of public procurement.

3.1 The multilevel Problem

The multilevel structure of first-price sealed-bid auctions presents a challenge due to the varying number of bids associated with each observation. Consequently, it is necessary to approximate both the conditional distribution of the random variable representing the number of bids and the conditional distribution of the bids themselves. Employing a single function approximator for both the discrete feature space and the bids may compromise accuracy, as the mapping involved in the task is surjective in nature, meaning that multiple inputs can map to the same output.

An alternative approach could involve using an autoencoding structure, such as the one implemented in the MEDGAN framework [35, 24], which learns the correspondence between the original and a latent space. However, MEDGAN is not well-suited to this context, as it is designed to learn mappings between categorical sets and sub-categorical combinatorial sets. Although MEDGAN could potentially be a viable option if the goal were to replicate the combinatorial space represented by the firms associated with each auction, the cardinality of the set of firms is excessively high, making it challenging for any algorithm to accurately handle such combinatorial structures.

Given these limitations, we choose to focus on generating bids independently from the labels of the firms behind each bid. This approach circumvents the complexities associated with the high cardinality of the set of firms and enables more effective generation of realistic auction simulations. Table 1 summarizes the structure of the data to be replicated and highlights the combinatorial issue associated to the multilevel structure of first-price auctions. In Table 1, we differentiate between multiclass and multilabel variables. A multiclass variable is a categorical one, with each instance belonging to exactly one out of a finite set of classes. On the other hand, a multilabel variable represents cases where an observation can simultaneously have multiple values. For instance, in our scenario, an auction could be either municipal or provincial, making the variable *municipality* a binary multiclass variable. It encodes whether an auction is issued at the municipal level (1) or not (0). Conversely, the variable *firms* can take varying numbers of values per auction, as multiple firms may participate in an auction. Therefore, the variable *firms* is multilabel. This implies that the cardinality (the number of possible states) of such a variable increases exponentially with the number of potential values.

Table 1: Structure of the data.

Variables	Type	Cardinality	Count
Discrete	multiclass	6,749	9
	multilabel	40,659	1 (<i>firms</i>)
Continuous	-	-	1 (<i>bids</i>)

3.2 Problem Formulation and Solution Framework

An instance of public procurement auction is labeled $a \in \{1, \dots, N\}$, where N is the number of examples. Let C be a collection of multiclass variables representing the features of the contracts being offered, and let $nb \in C$ be the number of bidders in an auction. We adopt a sparse one-hot encoded form of our discrete space, meaning that each $c \in C$ is represented by a binary vector such that $\sum_i c(i) = 1$, and \mathbf{c} is the aggregated vector of multiclass variables such that $\sum_j \mathbf{c}(j) = |C|$. Introducing the variable *firms*, we represent it by the binary variable \mathbf{f} , which encodes the presence or absence of firms bidding in each auction, such that $\sum_i \mathbf{f}(i) = nb^a$ for each instance a . Thus, an auction a is given by the set $\mathbf{x} = \{\mathbf{c}^a, \mathbf{f}^a, \mathbf{b}^a\}$, where \mathbf{b} is an array of continuous bids. It follows that the number of elements in \mathbf{b}^a equals nb^a . Let a^* be a hypothetical auction with only one feature ($|C| = 1$), and two bidders among four firms composing the market. Then, $nb^{a^*} = 2$, and \mathbf{x}^{a^*} could be (depending on which firms are actually bidding) $\{[1, 0], [1, 1, 0, 0], [b_1, b_2]\}$, where $b_1, b_2 > 0$. The

problem of density estimation in our case is to approximate the resulting joint distribution $p(\mathbf{x})$. Since we wish to generate new samples, our problem is to optimize the set of parameters ψ for the mapping $M : \mathbf{x} \xrightarrow{p(\mathbf{x};\psi)} \tilde{\mathbf{x}}$, where $\tilde{\mathbf{x}}$ is a fake but realistic synthetic auction, i.e., it is impossible to say if $\tilde{\mathbf{x}} \sim p(\mathbf{x})$ or $\tilde{\mathbf{x}} \sim p(\cdot)$, where $p(\cdot)$ is an arbitrary probability function.

We can decompose the mapping M into two independent and sequential functions, specifically by considering the functions $A : \mathbf{z} \xrightarrow{p(\mathbf{c}|\mathbf{z};\alpha)} \tilde{\mathbf{c}}$ and $B : \mathbf{c} \xrightarrow{p(\mathbf{b}|\mathbf{c};\beta)} \hat{\theta}$. The function A generates synthetic samples of auction characteristics, $\tilde{\mathbf{c}}$, from the noise input \mathbf{z} , by approximating $p(\mathbf{c}|\mathbf{z}; \alpha)$ using the set of parameters α . Utilizing the latent signal \mathbf{z} is advantageous because synthetic samples can be generated from scratch, with real data required only for training the approximator.

The function B approximates the conditional distributions of the bids given auction features, thus providing an aggregated representation of the firms. The vector $\hat{\theta}$ consists of the estimated parameters for the conditionals and depends on the probability function employed to describe the distribution of bids. Consequently, the bid generator B produces estimates that serve as arguments for a random generator, from which nb bids are sampled, i.e., $\hat{\mathbf{b}} \sim \mathcal{P}(\hat{\theta})$.

While the specific functional forms for A and B have not yet been determined, we can already observe how the general structure coherently articulates. Once both models are trained with respect to α and β , we can generate synthetic but realistic vectors of auction features, $\tilde{\mathbf{c}}$, using function A . These vectors are then utilized as input for function B , reconstructing the space \mathbf{c} , \mathbf{b} . Consequently, the generation of each data type is separated between A (discrete) and B (continuous), effectively flattening the multilevel auction structure, as B represents all firms simultaneously. At this stage, the problem initially formalized by the mapping M has been divided into functions A and B , or equivalently, we have defined $\psi = (\alpha, \beta)$.

The proposed pipeline does not explicitly incorporate the multilabel variable *firms*, as its cardinality is too large to feasibly preserve its original form and interpretability. Furthermore, we find it counterproductive to use a smaller continuous code representing *firms* if the original space cannot be retrieved subsequently. It is also worth noting that retaining only a subset of the data to reduce cardinality would undermine the objective.

For validation, we employ the inception scoring method, which involves training a predictor F using the newly sampled synthetic data and evaluating its performance on the real data. The bid generator is subsequently assessed separately with real examples that were not seen during training. Figure 3.1 offers a visual representation of the comprehensive system we have outlined.

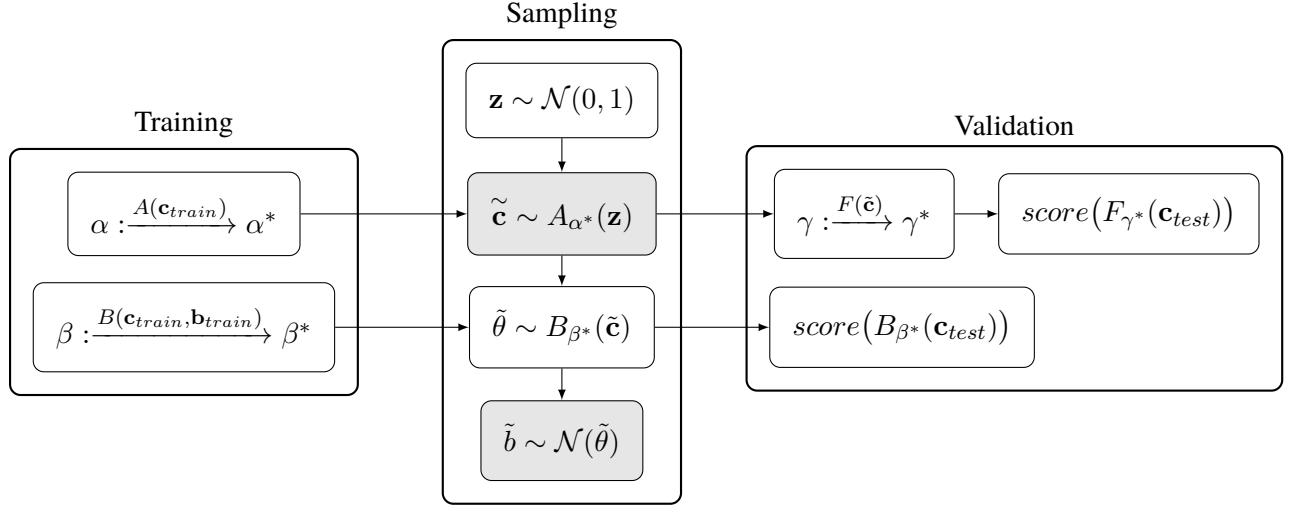
3.3 Approximating Auction Features Joint Density

This section discusses the details regarding the training of two synthesizers based on GANs and VAEs for the mapping $A : \mathbf{z} \xrightarrow{p(\mathbf{c}|\mathbf{z};\alpha)} \tilde{\mathbf{c}}$. It is essential to note that they are not complementary but competitors. It is a priori not possible to predict which one will perform best, so both are tried. Furthermore, they both deserve to be discussed in the context of our endeavor.

3.3.1 GAN-based Approach

Our GAN-based algorithm consists of a generator G and a critic C that we optimize with respect to the Wasserstein loss. The "training-by-sampling" principle is applied. In this context, G approximates $p(\mathbf{c}|\mathbf{z}, cond)$, where $cond$ is a binary conditional vector of size $|C|$ that sums to one. Consequently, the loss function must be augmented with a cross-entropy penalty term that enforces the sampling of a synthetic data point admitting $c_{i^*} = 1$, where i^* is the selected state of the chosen variable c . The resulting objective function is expressed as follows:


Figure 3.1: The graph illustrates the sequential structure of our meta-algorithm. The approximators A and B are trained, with regard to their respective set of parameters α and β , on real examples. Then, the solidified forms A^* and B^* are used to generate the synthetic features $\tilde{\mathbf{c}}$ and bids $\tilde{\mathbf{b}}$.




$$\min_G \max_C E \left[C(G(\mathbf{z})) - C(\mathbf{c}) \right] + CE(\tilde{\mathbf{c}}_i^*, cond). \quad (3)$$

The training procedure surrounding Equation (3) is detailed by Algorithm 1 in the Appendix while Figure 3.2 provides a graphical account of the algorithmic architecture of the model.

Figure 3.2: Neural representation of the generator and the critic used to synthesize auction characteristics. Note that W and s are respectively weight matrices and biases, while h are outputs from hidden blocks. The dimensions of h , and therefore of W and s , depend on the width parameters for a given network. The critic C outputs a scalar while G outputs a continuous array of size N_x (the number of continuous variables) and N_c one-hot arrays (one for each discrete multiclass variable).



(a) Generator



(b) Critic

We use the PacGAN configuration [36] in order to prevent mode collapse, a common problem in GANs

that refers to when the model fails to capture the diversity of the underlying data distribution and instead produces a limited set of similar or identical output samples [37]. The PacGAN configuration expands the input space of the critic with multiple stacks of the original input space.

Additionally, we implement a gradient penalty technique, also known as WGAN-GP, instead of the traditional "weight clipping" method of imposing a Lipschitz constraint on the critic. This approach stabilizes the training process and encourages the critic to learn the correct gradients by imposing a penalty on the norm of the critic's gradient with respect to its inputs. [32].

In the original GAN framework, the training process alternates between steps for the discriminator and the generator. The parameter k , a positive integer, determines how many times the discriminator is trained for each time the generator is trained. However, in the context of Wasserstein GANs, it is typically recommended to train the generator as frequently as the critic, essentially setting $k = 1$. This is to prevent the critic from overpowering the generator too quickly. We recall that in the context of Wasserstein GANs, the critic replaces the discriminator of the original framework, and predicts the distance between a given point and the decision boundary separating real and fake samples (instead of predicting the probability that the input is real).

Finally, we use the adaptive moment estimation (Adam) optimizer, an extension of the stochastic gradient descent that improves performance by adapting the learning rate of each weight parameter based on its first and second moments [38].

3.3.2 Tabular Variational Autoencoding

In the variational autoencoding framework, the encoder models $p(\mathbf{z}|\mathbf{c}, \mathbf{x})$, while the decoder replicates the original space by approximating $p(\mathbf{c}, \mathbf{x}|\mathbf{z})$. "Training-by-sampling" is thus neither necessary nor possible since the conditional vector would be encoded in a continuous layer. It means that only the decoder needs to be modified to generate tabular data. In addition, the reconstruction loss needs to be augmented with a cross-entropy term to ensure the integrity of the discrete structure. Algorithm 2 in the Appendix details the training procedure of the tabular variational autoencoder (TVAE), and Figure 3.3 provides its pictorial representation.


3.4 Training a Generator of Continuous Bids

Recall that B must represent all the firms at once by learning a bidding function based on auction features, which are synthesized with A . Ultimately, the best model from the two trained for A will be selected. However, for now, this information is irrelevant, as only inputs sampled from the original data need to be considered to train an approximator $B : \mathbf{c} \xrightarrow{p(b|\mathbf{c};\beta)} \hat{\theta}$. As shown by Figure 3.4, the standardized logarithmic bid distribution is Gaussian. Accordingly, the assumption of log-normality for the conditionals can be made, meaning that $\theta = (\mu, \sigma^2)$, where μ and σ^2 are the first two moments of the Gaussian density.

In accordance with the pipeline delineated earlier, the input space of the bid generator B must correspond to the output space of the auction generator A . The challenge arises due to the necessity of generating bids separately yet not independently from their conditional vector. Given that θ is a bi-dimensional vector, the task is defined as a multi-output regression. Since a multi-output linear model only represents the fitting of independent models and disregards the statistical dependence between each model's parameter set, a nonlinear approximator is required to model B .

A neural network, or more specifically a multi-layer perceptron (MLP), is the most comprehensive approximator capable of performing multi-target regression while maintaining the statistical dependence among its parameters [39]. Taking an input vector x and a target y , an MLP can represent the function $f(x; \gamma)$ by producing parameters for a distribution over y instead of directly predicting y [40, 41, 42]. This

Figure 3.3: Neural representation of the encoder and decoder composing the tabular VAE used to synthesize auction characteristics. Similarly to Figure 3.2, W and s are respectively weight matrices and biases, while h are outputs from hidden blocks. The σ_j are parameters of the decoder.



is applicable when γ is optimized with respect to $-\log P(y|x)$. Considering an MLP for B , its output is interpreted as a bi-dimensional vector containing the predicted first two moments of a conditional distribution $p(b|c)$. This model will be referred to as the *BidNet*.


To ensure the stability of the BidNet (i.e., systematic convergence), it was trained on different folds using a cross-validation procedure. The methodology involves dividing the training set into K folds (in this case, $K = 5$) and utilizing $K - 1$ folds to optimize the parameters with respect to the negative log-likelihood (NLL). Subsequently, the aggregated loss is computed on the remaining fold. After each pass over the training set, the validation set (the remaining fold) is employed to assess the objective NLL, which forms the basis for an early stopping rule. This strategy serves as a regularization method to prevent overfitting. The process is repeated until all K folds have been used as validation sets. Although various early stopping designs have been identified by [43], a customized one was employed to best suit the requirements of this study. The pseudocode related to the training of the BidNet with cross-validation is presented in Algorithm 3 in the Appendix.

3.5 Sampling Synthetic Auction Instances

The GAN-based model relies on its generator to sample synthetic features from noise. However, following the principle of training-by-sampling, the user must also create a conditional vector or a batch of conditional vectors alongside the latent space \mathbf{z} . To achieve this, one must apply steps 3 to 6 described in Algorithm 1 and then feed the trained generator with the noise and the conditional vector, as illustrated in Figure 3.2. A key advantage of this approach is that one can manually specify a conditional vector by disregarding steps 3 to 5 in Algorithm 1, enabling the replication of the signal of interest. The tabular VAE also samples through latent noise that can be created from a standard Gaussian distribution.

A synthetic instance of public procurement comprises a vector of auction features $\tilde{\mathbf{c}}$ and its associated array of bids. It is important to note that the process of bid generation occurs in two steps because the number of bids to be generated per auction varies. The variable "number of bidders" encodes this variability and has

Figure 3.4: Comparison of normal quantile-to-quantile plots relating to several numerical representations of the logarithmic bids. Left to right: mode specific normalization, minmax normalization, and standardization. Here, mode-specific normalization was applied to the bids in hope that it would provide a precise representation of the target, concerning the incumbent regression task.



been included in the joint distribution of auction characteristics; consequently, it is included in the output of A . Since the BidNet predicts the first moments of a Gaussian distribution, it is straightforward to sample nb bids from a random generator. In fact, when given synthetic inputs originating from A , the BidNet provides $\tilde{\theta} = (\tilde{\mu}, \tilde{\sigma}^2)$, which are the parameters for the conditionals $p(b|c; \theta)$. Synthetic bids are then drawn according to $\tilde{b} \sim \mathcal{N}(\tilde{\theta})$.

To maintain consistency with the notation introduced in Figure 3.1, we distinguish \tilde{b} from \hat{b} , the latter representing the predicted bids originating from $\mathcal{N}(\hat{\theta})$, where $\hat{\theta}$ is the output of the BidNet when provided with an original sample from the test set \mathbf{c}_{test} . It is important to note that the predicted bids \hat{b} should be used for validation purposes only.

4 Validation

We assess the faithfulness of the synthetic auction features, and by extension the performances of our synthesizers, by employing an inception score. The concept of inception scoring is based on the principle that a classifier or a regressor that has been trained successfully using synthetic data should perform well when tested on real-world instances. This principle underlies the expectation that the characteristics of a predictor’s output are primarily determined by its input, given a fixed set of parameters.

Inception scoring is a well-established method for validating the output of GANs and has been widely used for this purpose in the literature. The use of inception scoring as a validation metric for GANs was first proposed by [37] and has since been adopted and expanded upon by numerous researchers upon investigation on its usefulness as a metric for evaluating the quality of generated outputs [44].

In our case, we perform inception scoring by utilizing the binary variable *municipality*—one of the auction characteristics—as a target in the following binary classification problem:

$$f(\mathbf{c}_{-mun}) = p(mun),$$

where \mathbf{c}_{-mun} is the one-hot encoded set of auction features that excludes *municipality*. In other terms, we define the classifier $f : \frac{p(mun|\mathbf{c}_{-mun})}{\sum} \rightarrow [0, 1]$. If the synthetic data is realistic, that is, if the synthesizer

being evaluated managed to efficiently approximate the targeted joint distribution, then we expect the overall accuracy of $f(\mathbf{c}_{-mun})$ and $f(\tilde{\mathbf{c}}_{-mun})$ to be similar; provided f has been successfully trained.

Table 2: The classification accuracy for each class (Recall), as well as the average F1-score are reported. The classifiers have been trained two times each, using 100,000 training examples generated with the GANs and VAE-based method. Both synthesizers have been previously trained over 200 epochs. The numbers in the parenthesis indicate the performance gap (in percentage) between scores achieved on synthetic and real test-beds.

(a) Evaluation metrics of models trained on synthetic data generated by the GAN-based method.

Test-bed	Model	Recall (0)	Recall (1)	F1-score
Synthetic	Decision Tree	0.95	0.91	0.93
	k-NN	0.90	0.84	0.87
	CMLP	0.88	0.80	0.84
Real	Decision Tree	0.94 (-0.01)	0.57 (-0.15)	0.78 (-0.15)
	k-NN	0.81 (-0.09)	0.74 (-0.10)	0.78 (-0.09)
	CMLP	0.88 (0.00)	0.76 (-0.04)	0.83 (-0.01)

(b) Evaluation metrics of models trained on synthetic data generated by the VAE-based method.

Test-bed	Model	Recall (0)	Recall (1)	F1-score
Synthetic	Decision Tree	0.99	1.00	0.99
	k-NN	0.98	0.96	0.97
	CMLP	0.98	0.97	0.97
Real	Decision Tree	1.00 (+0.01)	0.00 (-1.00)	0.44 (-0.55)
	k-NN	0.00 (-0.98)	1.00 (+0.04)	0.24 (-0.73)
	CMLP	0.66 (-0.32)	0.89 (-0.08)	0.75 (-0.22)

Three models were utilized to represent the classifier f : a *decision tree*, *k-nearest neighbors* (k-NN) and a *multi-layer perceptron*. Here, we will name our MLP classifier as CMLP, where C stands for classification, in order to distinguish this classifier from the MLP we used earlier to predict the bids. Evaluation metrics for the three classifiers when trained on synthetic examples generated by the GAN-based and VAE-based methods are reported by Table 2. Based on the results, we can draw a clear conclusion: the GAN-based model succeeded in synthesizing reliable and realistic data, while the VAE-based model did not. Indeed, the CMLP achieved an overall classification F1-score on the real test-bed only 1% below what was attained on the synthetic test-bed. We recall that, having been trained with synthetic data, it is natural that the classifiers perform best on a synthetic test-bed. Despite the decision tree model performing the worst with an F1-score 15% below its reference point, it still achieved a 78%, which is acceptable according to classification standards. Furthermore, a specific pattern emerges from this experiment, indicating that instances belonging to class 0 are more accurately classified than those of class 1 (Recall (0) versus Recall (1)). All three classifiers have preserved this structure when tested with real data, providing additional evidence supporting our conclusion.

In contrast, when examining the VAE-based model, classifiers display relatively poor F1-scores, ranging from 44% (Decision Tree) to 75% (MLP). An F1-score below or around 50% suggests that the classifier was not better than a random algorithm. In addition, the decision tree and the k-NN both assigned all instances of the real test-bed to the same class (class 0 for the decision tree and class 1 for the k-NN). The MLP performed decently in recognizing some structure in the data generated by the VAE-based model, but the scores achieved on the synthetic test-bed are unrealistically high. This gap in performance between the GAN and VAE-based methods explains why we indulged in the specification of two methods in the first place. It should also be noted that the results in this kind of experiment may vary according to the hyperparameter settings, and in

theory, the tabular VAE should be able to reach a similar level of effectiveness with some fine tuning. The point to be made here is that the GAN-based model provides a reliable way to perform the task at hand without having to worry too much about hyperparameter tuning. Details about our hyperparameter tuning are to be found in the Appendix.

Now, we need to evaluate the efficiency of the BidNet, and by extension the reliability of the synthetic bids. An account of the BidNet’s relative efficiency is given by comparing it to a *regression tree* model and a *multi-target support vector regressor* (MSVR). The three regressors have been evaluated using the same metric—negative log-likelihood (NLL)—and their performances over five folds are reported in Table 3. The BidNet is, on average, the best and fastest model.

Nevertheless, since the NLL is a relative measure with values that can range from $-\infty$ to $+\infty$, a thorough investigation of the BidNet’s outputs is needed in order to assert the reliability of the resulting synthetic bid distribution. To that end, we propose a procedure for evaluating the distance between the distributions of fake and real bids, $Dist(p(b)||p(\tilde{b}))$, using the distance between the real and predicted distributions $Dist(p(b)||p(\hat{b}))$ as an identity, and the distance between the predicted and fake distributions $Dist(p(\hat{b})||p(\tilde{b}))$ as a control.

Table 3: The averages and standard deviations of the negative log-likelihood (NLL) over the five cross-validation folds are given for the BidNet, a tree-based regressor and a multi-target support vector regressor (MSVR). The column *best* reports the NLL to the best model identified with each method.

Model	NLL			Number of Iteration	
	mean	std	best	mean	std
BidNet	0.81	0.61	0.10	1,184	6,445
Regression Tree	2.11	0.03	2.06	92,500	-
MSVR	1.61	0.07	1.50	92,500	-

This procedure is also called "double validation" because it provides another way to validate the output of the synthesizers. Indeed, the opportunity to use the BidNet in order to construct an inception score naturally occurs since it has been trained on the real data and generates synthetic bids from fake auction features. The results displayed in Table 4 are coherent with those of the previous validation as the tabular VAE is at the origin of a noisy synthetic bid distribution. Meanwhile, $Dist(p(b)||p(\tilde{b}))$ and $Dist(p(\hat{b})||p(\tilde{b}))$ are very close to each other in both cases and also close to $Dist(p(b)||p(\hat{b}))$ in the GAN-generated data case. The BidNet is hence effective in preserving the statistical dependence between bids and their auction features and is a powerful approximator of $p(b|c)$.

5 Discussion

In this study, we have investigated the application of GANs and VAEs for crafting realistic auction simulations. We employed these data-driven methods to generate synthetic auction features and bids, aiming to provide a credible and useful tool for researchers and practitioners in the field of auction design and analysis. To ensure the reliability of the generated data, we evaluated the performances of the synthesizers by employing inception scores and assessing the faithfulness of the synthetic auction features. Additionally, we compared the BidNet’s efficiency to regression tree and multi-target support vector regressor models to further validate the newly generated synthetic data.

Our results demonstrated that while both GANs and VAEs are capable of generating realistic artificial data, the CTGAN outperformed the tabular VAE in our context. The CTGAN’s ability to generate data according to hand-crafted conditional vectors made it particularly suitable for simulating auction scenarios in

Table 4: The statistical distances have been measured with the *Wasserstein* or *Earth-Mover Distance* (EMD), and the *quantile-to-quantile root mean squared error* (QQ-RMSE). A score of 0 indicates identical distributions.

(a) Synthetic bids emanate from the real data.

	QQ-RMSE	EMD
$Dist(p(b) p(\hat{b}))$	1.17	0.02

(b) Synthetic bids emanate from the data generated by the GAN-based method.

	QQ-RMSE	EMD
$Dist(p(b) p(\tilde{b}))$	1.37	0.01
$Dist(p(\hat{b}) p(\tilde{b}))$	1.30	0.00

(c) Synthetic bids emanate from the data generated by the VAE-based method.

	QQ-RMSE	EMD
$Dist(p(b) p(\tilde{b}))$	400.55	0.28
$Dist(p(\hat{b}) p(\tilde{b}))$	400.55	0.30

our study. This finding highlights the importance of selecting the most appropriate method for the specific context and requirements of a given task.

The proposed method has broader implications and potential applications beyond simulating first-price auction data. For instance, the method could be extended to simulate other types of auctions and market mechanisms, as explored by [45]. Furthermore, our method could be integrated with multi-agent systems, such as those employed by [46], to simulate more complex interactions and behaviors in auction markets. Combining our method with other machine learning techniques, like those used by [47] to forecast GDP growth, could enable researchers to predict various economic indicators or outcomes based on synthetic data.

However, it is important to acknowledge the limitations of our approach. A primary limitation is that GANs and VAEs can only capture the dynamics of data structures through time based on past data. Consequently, simulations generated using data-driven methods like ours cannot incorporate novel shock scenarios. This limitation is not unique to our approach, as no model can accurately predict the future, and anticipating future shocks that substantially modify the data generation process of a joint distribution remains an ongoing challenge in economics in general.

In conclusion, this study offers a valuable contribution to the field of auction simulations by demonstrating the potential of GANs and VAEs for generating realistic and reliable synthetic data. While our approach is not without its limitations, the findings provide a strong foundation for future research in this area, particularly in exploring the potential of CTGANs and other generative models for simulating complex auction scenarios and their applications in various economic contexts.

Statements and Declarations

Funding

This work was supported by Meta Research following the authors application to the request for proposals on [agent-based user interaction simulation to find and fix integrity and privacy issues](#).

Competing Interests

The authors declare no competing interests.

Author Contributions

All authors contributed to the study conception and design. Material preparation, data collection, analysis were performed by Igor Sadoune. The first draft of the manuscript was written by Igor Sadoune and all authors commented on previous versions of the manuscript. All authors read and approved the final manuscript.

Appendix A: Methodology Overview

Initially, the SEAO dataset underwent a thorough cleaning process. This involved handling missing values, removing irrelevant columns, and reformatting specific columns to ensure their consistency and reliability.

Following the data cleaning, preprocessing was conducted to transform the dataset and make it suitable for machine learning applications. Discrete variables were transformed using one-hot encoding techniques, while continuous bid values were standardized.

To generate synthetic data, two primary generative models were utilized: CTGAN and TVAE. The CTGAN model was trained using an array of hyperparameters, including distinct embedding dimensions, generator and discriminator dimensions, learning rates, and specific decay rates. Likewise, the TVAE, a variant that incorporates a variational autoencoder structure, was trained with specific parameters, including hidden and latent dimensions.

Once trained, these models were then employed to sample synthetic datasets, replicating the patterns and distributions seen in the original SEAO dataset.

Subsequent to the synthetic data generation, we introduced the BidNet neural network model. This model was designed to predict bid values using both discrete and continuous inputs. For training efficiency, the model utilized cross-validation and early stopping methodologies. Several hyperparameters, including learning rate, batch size, and number of epochs, were tuned to optimize the model's performance.

To assess the quality of the synthetic data produced by CTGAN and TVAE, a series of classifiers, including Decision Trees, k-Nearest Neighbors, and Neural Networks, were trained on both the real and synthetic datasets. Performance metrics from these classifiers provided insights into the fidelity and utility of the synthetic data.

Finally, the BidNet model's performance was critically evaluated using various metrics. These metrics, namely the Root Mean Square Error (RMSE), Jensen-Shannon distance (JS), and Wasserstein distance (WS), compared the synthetic and real bids, giving a comprehensive understanding of the model's accuracy and effectiveness.

Throughout the entire process, special attention was given to reproducibility. Foundational functionalities ensured consistent random states, allowing for deterministic behavior across runs. Additionally, capabilities were established to save intermediate results, trained models, and to manage computation across various devices, whether CPU or GPU.

Appendix B: Algorithms

Algorithm 1 Training GANs-based auction features generator

```

1:  $D_{train} \leftarrow$  Initialize training set
2: while  $C(fake) > threshold$  do ▷ The critic can be optimized until  $C(fake)$  is near 0.
3:   Randomly select a discrete variable  $c$  with equal probability
4:   Compute the probability mass function (PMF) of  $c$ 
5:   Randomly select a state  $i^*$  inherent to  $c$  according its PMF
6:   Create the conditional vector  $cond$  so that  $\sum_i cond(i) = 1$  and  $cond(i^*) = 1$ 
7:   for  $batch \in \{1, \dots, N_{batches}\}$  do ▷ Gradient descent with mini-batch
8:      $real \leftarrow d(c_{i^*} = 1) \sim D_{train}$  ▷ Sample batch of real examples respecting the constraint
9:      $z \sim \mathcal{N}(0, 1)$  ▷ Sample noise
10:     $fake \leftarrow \tilde{d} \sim G(z)$  ▷ Sample fake examples
11:     $real \leftarrow [real] \times 10$  ▷ Stack input 10 times for Pac configuration
12:     $fake \leftarrow [fake] \times 10$ 
13:     $L^j \leftarrow (C(fake_j) - C(real_j)) + CE(\tilde{c}, cond)$ 
14:     $L^{batch} \leftarrow L^{batch} + \lambda(\|\nabla L^{batch}\|_2 - 1)^2$  ▷ Apply gradient penalty
15:     $w_{crit} \leftarrow w_{crit} + Adam(\nabla_{w_{crit}} \frac{1}{m} \sum_i L^{batch}(i))$  ▷ Updating  $C$  with Adam
16:    if  $batch \bmod k = 0$  then ▷ Synchronicity, depends on  $k$ 
17:       $w_{gen} \leftarrow w_{gen} + Adam(\nabla_{w_{gen}} \frac{1}{m} \sum_i -C(G(z)))$  ▷ Updating  $G$  with Adam

```

Algorithm 2 Training tabular VAE for auction features

```

1:  $D_{train} \leftarrow$  Initialize training set
2: for  $epoch \in N_{steps}$  do
3:   for  $batch \in \{1, \dots, N_{batches}\}$  do ▷ Gradient descent with mini-batch
4:      $real \sim D_{train}$  ▷ Sample batch of real examples
5:      $(\mu, \sigma^2) \sim Enc(real)$ 
6:      $z \sim \mathcal{N}(\mu, \sigma^2)$  ▷ Sample latent input
7:      $fake \sim Dec(z)$  ▷ Sample fake examples
8:      $L^j \leftarrow CE(\tilde{c}_j - \arg \max(c)) + (2\sigma^2)^{-1}(x_j - \tanh(\tilde{x}_j))^2 + KL(\mathcal{N}(\mu_j, \sigma_j^2), \mathcal{N}(0, 1))$ 
9:      $w \leftarrow w + Adam(\nabla_w \frac{1}{m} \sum_i L^{batch}(i))$  ▷ Updating parameters with Adam

```

Algorithm 3 K-folds cross-validation BidNet training procedure

```
1:  $D \leftarrow \{D_1, \dots, D_K\}$  ▷ Initialize K-folds
2:  $loss^* \leftarrow \infty$  ▷ Initialize best model
3: for  $fold \in D$  do
4:    $reset(w_{BidNet})$  ▷ Reset parameters before entering each new fold
5:    $D_{val} \leftarrow D(fold), D_{train} \leftarrow D(-fold)$ 
6:   while has not converged do
7:     for  $batch \in \{1, \dots, N_{batches}\}$  do ▷ Gradient descent with mini-batch
8:        $d \sim D_{train}$  ▷ Sample batch of real examples
9:        $\hat{\theta} \leftarrow BidNet(d)$ 
10:       $L^{train} \leftarrow m^{-1} \sum_i NLL(\hat{\theta})_i$  ▷ compute NLL on training batch
11:       $w \leftarrow w + Adam(\nabla L^{train})$  ▷ Update BidNet
12:       $converged \leftarrow ES(L^{val})$  ▷ Early stopping
13:       $L^{val} \leftarrow n^{-1} \sum_j NLL(BidNet(D_{val}))_j$  ▷ compute NLL on validation fold
14:      if  $L^{val} < loss^*$  then
15:         $loss^* \leftarrow L^{val}$ 
16:        save model
```

Algorithm 4 Synthetic bid validation / Double validation

```
1:  $\beta \leftarrow \beta^*$  ▷ load best set of parameters for BidNet
2:  $\alpha \leftarrow \alpha^*$  ▷ load optimized set of parameters for synthesizer
3:  $\tilde{\mathbf{c}} \sim A_{\alpha^*}(\mathbf{z})$  ▷ sample synthetic examples from the trained synthesizer
4:  $\mathbf{c} \sim D_{test}$  ▷ sample a test-set of real instances
5:  $\hat{\mathbf{b}} \sim B_{\beta^*}(\mathbf{c})$  ▷ sample predicted bids from the test-set of real instances using the BidNet
6:  $\tilde{\mathbf{b}} \sim B_{\beta^*}(\tilde{\mathbf{c}})$  ▷ sample fake bids with the synthetic data emanating from the synthesizer
7:  $Dist(p(\mathbf{b})||p(\tilde{\mathbf{b}}))$  ▷ compute the statistical distance between the fake and real distributions of bids
8:  $Dist(p(\mathbf{b})||p(\hat{\mathbf{b}}))$  ▷ compute the statistical distance between the predicted and real distributions of bids
9:  $Dist(p(\hat{\mathbf{b}})||p(\tilde{\mathbf{b}}))$  ▷ compute the statistical distance between the predicted and fake distributions of bids
```

References

- [1] P. Milgrom and E. Kwerel, *Putting Auction Theory to Work*. Cambridge University Press, 2003.
- [2] B. Edelman, M. Ostrovsky, and M. Schwarz, “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords,” *American Economic Review* **97** no. 1, (Mar, 2007) 242–259.
- [3] A. Hortaçsu and S. L. Puller, “Understanding Strategic Bidding in Multi-Unit Auctions: A Case Study of the Texas Electricity Spot Market,” *The RAND Journal of Economics* **39** no. 1, (2008) 86–114.
- [4] E. Uyerra, J. M. Zabala-Iturriagoitia, K. Flanagan, and E. Magro, “Public procurement, innovation and industrial policy: Rationales, roles, capabilities and implementation,” *Research Policy* **49** no. 1, (Feb, 2020) .
- [5] Tesfatsion, *Handbook of Computational Economics: Volume 2, Agent-based Computational Economics*. Elsevier, handbooks ed., 2006.
- [6] E. Bonabeau, “Agent-based modeling: Methods and techniques for simulating human systems,” *Proceedings of the National Academy of Sciences of the United States of America* **99** no. SUPPL. 3, (May, 2002) 7280–7287. <https://www.pnas.org/doi/abs/10.1073/pnas.082080899>.
- [7] H. Dawid and A. Pyka, “Introduction: Special Issue on Evolutionary Dynamics and Agent-Based Modeling in Economics,” *Computational Economics* **52** no. 3, (Oct, 2018) 707–710. <https://link.springer.com/article/10.1007/s10614-018-9831-8>.
- [8] S. Athey, “The Impact of Machine Learning on Economics,” *The Economics of Artificial Intelligence: An Agenda* (2019) 507–552. <https://doi.org/10.7208/9780226613475-023>.
- [9] J. A. Calpin, M. R. Salisbury, J. A. Vitkevich, and D. R. Woodward, “Extending the high level architecture paradigm to economic simulation,” *Computational Economics* **17** no. 2-3, (2001) 141–154. <https://link.springer.com/article/10.1023/A:1011619907538>.
- [10] D. Xie, N. Zhang, and D. A. Edwards, “Simulation Solution to a Two-Dimensional Mortgage Refinancing Problem,” *Computational Economics* **52** no. 2, (Aug, 2018) 479–492. <https://link.springer.com/article/10.1007/s10614-017-9689-1>.
- [11] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” *Advances in Neural Information Processing Systems* **3** no. January, (2014) 2672–2680, [arXiv:1406.2661](https://arxiv.org/abs/1406.2661). <http://www.github.com/goodfeli/adversarialhttp://arxiv.org/abs/1406.2661>.
- [12] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” in *2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings*. International Conference on Learning Representations, ICLR, 2014. [arXiv:1312.6114](https://arxiv.org/abs/1312.6114).
- [13] L. Xu, M. Skoularidou, A. Cuesta-Infante, and K. Veeramachaneni, “Modeling Tabular data using Conditional GAN,” 2019. [http://arxiv.org/abs/1907.00503](https://arxiv.org/abs/1907.00503).
- [14] W. B. Arthur, “Complexity and the economy,” Apr, 1999.
- [15] F. Decarolis, “Comparing Public Procurement Auctions.” 2017.

- [16] G. Marti, “CORRGAN: Sampling Realistic Financial Correlation Matrices Using Generative Adversarial Networks,” in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8459–8463. Oct, 2020. [arXiv:1910.09504](https://arxiv.org/abs/1910.09504). <http://arxiv.org/abs/1910.09504>.
- [17] R. Rubinstein Y. and D. Kroese P., *SIMULATION AND THE MONTE CARLO METHOD*. 2008.
- [18] A. C. Dvison, D. V. Hinkley, and E. Schechtman, “Efficient bootstrap simulation,” *Biometrika* **73** no. 3, (Dec, 1986) 555–566. <https://academic.oup.com/biomet/article/73/3/555/250081>.
- [19] S. Elsayah, T. Filatova, A. J. Jakeman, A. J. Kettner, M. L. Zellner, I. N. Athanasiadis, S. H. Hamilton, R. L. Axtell, D. G. Brown, J. M. Gilligan, M. A. Janssen, D. T. Robinson, J. Rozenberg, I. I. T. Ullah, and S. J. Lade, “Eight grand challenges in socio-environmental systems modeling,” *Socio-Environmental Systems Modelling* **2** (Jan, 2020) 16226.
- [20] X. Chen, Y. Duan, R. Houthoof, J. Schulman, I. Sutskever, and P. Abbeel, “InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets,” in *Advances in Neural Information Processing Systems*, pp. 2180–2188. 2016. [arXiv:1606.03657](https://arxiv.org/abs/1606.03657). <https://arxiv.org/abs/1606.03657>.
- [21] R. L. Axtell and J. D. Farmer, “Agent-Based Modeling in Economics and Finance: Past, Present, and Future,” *Journal of Economic Literature* no. August 2018, (2021) to appear.
- [22] S. Athey and G. W. Imbens, “Machine Learning Methods That Economists Should Know About,” *Annual Review of Economics* **11** no. 1, (2019) 685–725, [arXiv:1903.10075](https://arxiv.org/abs/1903.10075). <https://www.gsb.stanford.edu/faculty-research/working-papers/machine-learning-methods-economists-should-know-about>.
- [23] H. Ba, “Improving Detection of Credit Card Fraudulent Transactions using Generative Adversarial Networks,” [arXiv:1907.03355](https://arxiv.org/abs/1907.03355). <http://arxiv.org/abs/1907.03355>.
- [24] P. Jackson and M. Lussetti, “Extending a Generative Adversarial Network to Produce Medical Records with Demographic Characteristics and Health System Use,” *2019 IEEE 10th Annual Information Technology, Electronics and Mobile Communication Conference, IEMCON 2019* (2019) 515–518.
- [25] J. Yoon, J. Jordon, and M. Van Der Schaar, “Ganite: Estimation of individualized treatment effects using generative adversarial nets,” *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings* no. 2010, (2018) 1–15. <https://openreview.net/pdf?id=ByKWUeWA->.
- [26] C. Louizos, U. Shalit, J. Mooij, D. Sontag, R. Zemel, and M. Welling, “Causal effect inference with deep latent-variable models,” *Advances in Neural Information Processing Systems 2017-Decem* no. Nips, (May, 2017) 6447–6457, [arXiv:1705.08821](https://arxiv.org/abs/1705.08821). <http://arxiv.org/abs/1705.08821>.
- [27] R. Cai, J. Qiao, K. Zhang, Z. Zhang, and Z. Hao, “Causal discovery with cascade nonlinear additive noise models,” *IJCAI International Joint Conference on Artificial Intelligence 2019-Augus* (2019) 1609–1615, [arXiv:1905.09442](https://arxiv.org/abs/1905.09442).
- [28] Z. Cai, Z. Xiong, H. Xu, P. Wang, W. Li, and Y. Pan, “Generative Adversarial Networks: A Survey Toward Private and Secure Applications,” *ACM Computing Surveys* **54** no. 6, (2021) , [arXiv:2106.03785](https://arxiv.org/abs/2106.03785).
- [29] M. Wong and B. Farooq, “A bi-partite generative model framework for analyzing and simulating large scale multiple discrete-continuous travel behaviour data,” *Transportation Research Part C: Emerging Technologies* **110** (Jan, 2020) 247–268, [arXiv:1901.06415](https://arxiv.org/abs/1901.06415).

- [30] S. Takahashi, Y. Chen, and K. Tanaka-Ishii, “Modeling financial time-series with generative adversarial networks,” *Physica A: Statistical Mechanics and its Applications* **527** (Aug, 2019) .
- [31] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein GAN,” *Proceedings of the 34th International Conference on Machine Learning* **70** (Jan, 2017) 214–223, [arXiv:1701.07875](https://arxiv.org/abs/1701.07875).
<https://proceedings.mlr.press/v70/arjovsky17a.html><http://arxiv.org/abs/1701.07875>.
- [32] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, “Improved training of wasserstein GANs,” *Advances in Neural Information Processing Systems* **2017-Decem** (2017) 5768–5778, [arXiv:1704.00028](https://arxiv.org/abs/1704.00028). <https://github.com/igul222/improved>.
- [33] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. P. Smolley, “Least Squares Generative Adversarial Networks,” *Proceedings of the IEEE International Conference on Computer Vision* **2017-Octob** (2017) 2813–2821, [arXiv:1611.04076](https://arxiv.org/abs/1611.04076).
- [34] R. D. Hjelm, A. P. Jacob, T. Che, A. Trischler, K. Cho, and Y. Bengio, “Boundary-seeking generative adversarial networks,” in *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings*. 2018. [arXiv:1702.08431](https://arxiv.org/abs/1702.08431).
- [35] E. Choi, S. Biswal, B. Malin, J. Duke, W. F. Stewart, and J. Sun, “Generating Multi-label Discrete Patient Records using Generative Adversarial Networks,” [arXiv:1703.06490](https://arxiv.org/abs/1703.06490).
<http://arxiv.org/abs/1703.06490>.
- [36] Z. Lin, G. Fanti, A. Khetan, and S. Oh, “PacGan: The power of two samples in generative adversarial networks,” *Advances in Neural Information Processing Systems* **2018-Decem** no. 8, (2018) 1498–1507, [arXiv:1712.04086](https://arxiv.org/abs/1712.04086).
- [37] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, and X. Chen, “Improved Techniques for Training GANs,” in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, eds., vol. 29, pp. 1–9. Curran Associates, Inc., Dec, 2016. [arXiv:1701.00160](https://arxiv.org/abs/1701.00160). <http://arxiv.org/abs/1701.00160><https://proceedings.neurips.cc/paper/2016/file/8a3363abe792db2d8761d6403605aeb7-Paper.pdf>.
- [38] D. P. Kingma and J. L. Ba, “Adam: A Method for Stochastic Optimization,” *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings* (Dec, 2014) , [arXiv:1412.6980](https://arxiv.org/abs/1412.6980). <https://arxiv.org/abs/1412.6980v9>.
- [39] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. The MIT Press, 2016.
- [40] P. M. Williams, “Using Neural Networks to Model Conditional Multivariate Densities,” *Neural Computation* **8** no. 4, (1996) 843–854.
- [41] R. Neuneier, F. Hergert, W. Finnoff, and D. Ormoneit, “Estimation of Conditional Densities: A Comparison of Neural Network Approaches,” *Icann '94* (1994) 689–692.
- [42] H. Schioler and P. Kulczycki, “Neural network for estimating conditional distributions,” *IEEE Transactions on Neural Networks* **8** no. 5, (1997) 1015–1025.
- [43] L. Prechelt, “Early stopping - But when?,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **7700 LECTU** (2012) 53–67.

- [44] M. Lucic, K. Kurach, M. Michalski, O. Bousquet, and S. Gelly, “Are Gans created equal? A large-scale study,” *Advances in Neural Information Processing Systems* **2018-Decem** no. Nips, (2018) 700–709, [arXiv:1711.10337](https://arxiv.org/abs/1711.10337).
- [45] X. Zhou and H. Li, “Buying on Margin and Short Selling in an Artificial Double Auction Market,” *Computational Economics* **54** no. 4, (Dec, 2019) 1473–1489. <https://link.springer.com/article/10.1007/s10614-017-9722-4>.
- [46] J. Lussange, I. Lazarevich, S. Bourgeois-Gironde, S. Palminteri, and B. Gutkin, “Modelling Stock Markets by Multi-agent Reinforcement Learning,” *Computational Economics* **57** no. 1, (Jan, 2021) 113–147. <https://link.springer.com/article/10.1007/s10614-020-10038-w>.
- [47] B. Chu and S. Qureshi, “Comparing Out-of-Sample Performance of Machine Learning Methods to Forecast U.S. GDP Growth,” *Computational Economics* (Sep, 2022) 1–43. <https://link.springer.com/article/10.1007/s10614-022-10312-z>.