



Optimal Economic Design Through Deep Learning (Short Paper)

Citation

Dütting, Paul, Zhe Feng, Harikrishna Narasimhan, and David C. Parkes. 2017. Optimal Economic Design through Deep Learning (Short Paper). 31st Conference on Neural Information Processing Systems (NIPS), December 4-9, 2017, Long Beach, CA, USA.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:41718766>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Optimal Economic Design through Deep Learning (Short paper)*

Paul Dütting[†] Zhe Feng[‡] Harikrishna Narasimhan[‡] David C. Parkes[‡]

Abstract

Designing an auction that maximizes expected revenue is an intricate task. Despite major efforts, only the single-item case is fully understood. We explore the use of tools from deep learning on this topic. The design objective is revenue optimal, dominant-strategy incentive compatible auctions. For a baseline, we show that multi-layer neural networks can learn almost-optimal auctions for a variety of settings for which there are analytical solutions, and even without encoding characterization results into the design of the network. Our research also demonstrates the potential that deep nets have for deriving auctions with high revenue for poorly understood problems.

A fundamental result in auction theory is the characterization of revenue optimal auctions as *virtual value maximizers* [21]. We know, for example, that second price auctions with a suitably chosen reserve price are optimal when selling to bidders with i.i.d. values, and how to prioritize one bidder over another in settings with bidder asymmetry. Myerson’s theory is as rare as it is beautiful. In a single item auction, a bidder’s type is a single number (her value for the item), making this a single-dimensional mechanism design problem. The design of optimal auctions for multiple items has proved much more difficult, and defied a thorough theoretical understanding.

Tracing the contours of analytical results reveals the difficulty of this problem of multi-dimensional mechanism design. Decades after Myerson’s result, we do not have precise descriptions of optimal auctions with two or more bidders and more than two items. Even the design of the optimal auction for selling two items to a single buyer is not fully understood.¹ For a single additive buyer with item values i.i.d. $U(0, 1)$, Manelli and Vincent [20] handle two items, and Giannakopoulos and Koutsoupias [14] up to six items. Yao [27] provides the optimal design for any number of additive bidders and two items, as long as item values can take on one of two possible values.

A promising alternative is to use computers to solve problems of optimal economic design. The framework of *automated mechanism design* [8] suggests to use algorithms for the design of optimal mechanisms. Early approaches required an explicit representation of all possible type profiles, which is exponential in the number of agents and does not scale. Others have proposed to search through a parametric subfamily of mechanisms, and are not fully general [17, 18, 25, 22]. In recent years, efficient algorithms have been developed for the design of optimal, *Bayesian incentive compatible (BIC)* auctions in multi-bidder, multi-item settings [2, 5, 1, 3, 6, 4, 9]. But despite this, many

*Full version of the paper: Dütting et al. (2017): “Optimal auctions through deep learning,” *CoRR*, abs/1706.03459.

[†]Department of Mathematics, London School of Economics, Houghton Street, London WC2A 2AE, UK. Email: p.d.duetting@lse.ac.uk.

[‡]John A. Paulson School of Engineering and Applied Sciences, Harvard University, 33 Oxford Street, Cambridge, MA 02138, USA. Email: {zhe_feng, hnarasimhan, parkes}@g.harvard.edu.

¹Results are known for additive i.i.d. $U(0, 1)$ values on items [20], additive, independent and asymmetric distributions on item values [9, 15, 26], additive, i.i.d. exponentially distributed item values [9] and extended to multiple items [13], additive, i.i.d. Pareto distributions on item values [19], and unit-demand valuations with item values i.i.d. $U(c, c + 1)$, $c > 0$ [23].

important questions remain unsolved. For instance, while there is a characterization of optimal mechanisms as virtual-value maximizers [4, 7], relatively little is known about the exact structure of these mechanisms. Similar progress has not been made on the design of optimal, *dominant-strategy incentive compatible* (DSIC) mechanisms (rather, there has been emphasis to the design of approximate, DSIC mechanisms; e.g. Hart and Nisan [19]).

The disruptive developments in machine learning suggest an opportunity to use machine learning for the design of optimal economic mechanisms. The use of machine learning for mechanism design was earlier pioneered by Dütting et al. [11], who use support vector machines to design payment rules for a given allocation rule (which can be designed to be scalable). But their framework need not provide incentive compatibility when the rule is not implementable and does not support design objectives stated on payments.²

We have initiated our research into the use of deep learning for optimal design on the problem of multi-item, optimal auction design [10]. Subsequently, we have also investigated problems with private budgets [12], as well as problems of mechanism design without money [16] (with N. Golowich). We give here only a brief overview of the methodology and results from Dütting et al. [10]. For *type profile* $v = (v_1, \dots, v_n)$ (for $N = \{1, \dots, n\}$ agents), parametrized *allocation rule* g^w and *payment rule* p^w (mapping reported types to an allocation and payments, respectively), with weights w , and with *loss function* $\mathcal{L}(v; g^w, p^w) = -\sum_{i \in N} p_i^w(v)$, the machine learning problem of interest for optimal auction design can be stated as:

$$\begin{aligned} & \min_w \mathbf{E}_{v \sim F_V} [\mathcal{L}(v; g^w, p^w)] & (1) \\ \text{s.t. [IC]} & \quad rgt_i(w) = 0, \quad \forall i \in N \\ \text{[IR]} & \quad irp_i(w) = 0. \quad \forall i \in N \end{aligned}$$

The type profile is sampled from n agents for some value distribution F_V . The *expected ex post regret* for agent i , given parameters w , is

$$rgt_i(w) = \mathbf{E}_{v \sim F_V} \left[\max_{v'_i \in V_i} u_i(v'_i, v_{-i}; v_i, g^w, p^w) - u_i(v_i, v_{-i}; v_i, g^w, p^w) \right], \quad (2)$$

where V_i is the set of possible valuations for agent i , and $u_i(v'_i, v_{-i}; v_i, g^w, p^w)$ is the utility (value minus price) to agent i with valuation v_i when reporting v'_i , when others report v_{-i} , and with allocation and payment rule g^w, p^w , respectively. Zero expected ex post regret corresponds to a mechanism that is, except with measure zero, dominant-strategy incentive compatible (or strategy-proof). The *expected violation of individual rationality* for agent i , given parameters w , is

$$irp_i(w) = \mathbf{E}_{v \sim F_V} \left[\max\{0, -u_i(v; v_i, g^w, p^w)\} \right]. \quad (3)$$

Zero expected violation of individual rationality corresponds to a mechanism that ensures, except with measure zero, that the utility from participation is non-negative.

We use multi-layer, feed-forward neural networks to represent the parametrized economic mechanism. These networks provide differentiable, non-linear function approximations, where the training problem is optimized through stochastic gradient descent together with augmented Lagrangian optimization. Our *fully agnostic* approach proceeds without the use of characterization results and, because of this, holds the most promise in discovering new economic designs. The input layer of the REGRETNET architecture represents bids, and the network has two logically distinct components: the *allocation network* and *payment network* (see Figure 1). Each network is a fully-connected, feed-forward network with multiple hidden layers (denoted $h^{(r)}$ and $c^{(t)}$) and an output layer. In our experiments these networks make use of two hidden layers, each with 100 units. Each hidden unit has a sigmoidal activation function applied to a weighted sum of outputs from the previous layer. These weights form the parameters of the network.³

²Procaccia et al. [24] studied the learnability of voting rules, but without considering incentives.

³For a given bid profile b , illustrated here as providing a number for each agent for each of m items, the allocation network outputs a vector of allocation probabilities $z_{1j}(b), \dots, z_{nj}(b)$, for each item $j \in [m]$, through a *softmax activation function*, with $\sum_{i=1}^n z_{ij}(b) \leq 1$ for each item $j \in [m]$. Bundling of items is possible because the value on output units corresponding to allocating each of two different items to the same agent can be correlated. In another variation, we handle unit-demand valuations by using an additional set of softmax activation functions, one per agent, and taking the minimum of these item-wise and agent-wise softmax components in defining the output layer. The output layer of the payment network defines the payment for each agent for a given type profile, and makes use of *ReLU activation units* ($relu(s) = \max\{s, 0\}$).




Figure 1: REGRETNET: The allocation and payment networks for a setting with multiple agents $(1, \dots, n)$ and multiple items $(1, \dots, m)$ [10].

In practice, the loss, regret and IR penalty involved in formulating (1) are estimated from samples of valuation profiles (with regret, for example, estimated for a given profile as the maximum utility gain over a set of additional samples of possible deviations). Augmented Lagrangian optimization handles the IC and IR constraints through a weighted penalty in the objective, with stochastic gradient descent on this unconstrained problem interleaved with updates to Lagrangian multipliers. Through this approach, almost optimal auctions with almost zero expected *ex post* regret and almost zero expected IR penalty can be obtained across a number of different economic environments.




Figure 2: Two items, a single unit-demand agent, item values $x_1, x_2 \sim U[2, 3]$. Comparing the alloc. rules of REGRETNET and the optimal design [23], for item 1 and 2 respectively, varying the agent's value. The design of the optimal rule is superimposed, with different alloc. regions separated by dashed lines and numbers in each region indicating the prob. with which the item is allocated.

See Figure 2 for just one illustrative comparison between an allocation rule in the learned mechanism and that of an optimal rule (in the case where an analytical result exists). In this case, the network attains the optimal expected revenue of 0.387 (at truthful bidding) and with per-agent expected *ex post* regret of 0.01. We have also used this approach to design essentially incentive-aligned auctions with high revenue for economic environments out of reach of theoretical analysis [10]. These include the single, additive bidder environment with ten items (there is no analytical solution with more than six items), as well as a setting with two items and two additive bidders, where item values are i.i.d. uniform on $\{0.5, 1.0, 1.5\}$ (analytical results are known only for supports of size two). In this case, we learn an auction with effectively zero regret and IR penalty, and expected revenue of 1.868 (compared with 1.818 for an optimal item-wise Myerson auction, and 1.697 for an optimal bundle-wise Myerson auction).

References

- [1] ALAEI, S., H. FU, N. HAGHPANAH, AND J. D. HARTLINE (2013): “The Simple Economics of Approximately Optimal Auctions,” in *54th Ann. IEEE Symp. on Found. of Computer Science, FOCS*, 628–637.
- [2] ALAEI, S., H. FU, N. HAGHPANAH, J. D. HARTLINE, AND A. MALEKIAN (2012): “Bayesian optimal auctions via multi- to single-agent reduction,” in *ACM Conference on Electronic Commerce, EC’12*, 17.
- [3] BHALGAT, A., S. GOLLAPUDI, AND K. MUNAGALA (2013): “Optimal auctions via the multiplicative weight method,” in *Proceedings of the 14th ACM Conference on Economics and Computation*, 73–90.
- [4] CAI, Y., C. DASKALAKIS, AND M. S. WEINBERG (2012): “Optimal Multi-dimensional Mechanism Design: Reducing Revenue to Welfare Maximization,” in *Proceedings of the 53rd IEEE Symposium on Foundations of Computer Science*, 130–139.
- [5] CAI, Y., C. DASKALAKIS, AND S. M. WEINBERG (2012): “An algorithmic characterization of multi-dimensional mechanisms,” in *Proc. 44th Annual ACM Symp. on Theory of computing*, 459–478.
- [6] ——— (2013): “Reducing revenue to welfare maximization: Approximation algorithms and other generalizations,” in *Proc. 24th Ann ACM-SIAM Symp. on Discrete Algs.*, 578–595.
- [7] ——— (2013): “Understanding Incentives: Mechanism Design Becomes Algorithm Design,” in *Proceedings of the 54th IEEE Symposium on Foundations of Computer Science*, 618–627.
- [8] CONITZER, V. AND T. SANDHOLM (2002): “Complexity of Mechanism Design,” in *Proceedings of the 18th Conference in Uncertainty in Artificial Intelligence*, 103–110.
- [9] DASKALAKIS, C., A. DECKELBAUM, AND C. TZAMOS (2016): “Strong Duality for a Multiple-Good Monopolist,” *Econometrica*.
- [10] DÜTTING, P., Z. FENG, H. NARASIMHAN, AND D. C. PARKES (2017): “Optimal Auctions through Deep Learning,” *CoRR*, abs/1706.03459.
- [11] DÜTTING, P., F. A. FISCHER, P. JIRAPINYO, J. K. LAI, B. LUBIN, AND D. C. PARKES (2015): “Payment Rules through Discriminant-Based Classifiers,” *ACM Trans. Economics and Comput.*, 3, 5:1–5:41.
- [12] FENG, Z., H. NARASIMHAN, AND D. C. PARKES (2017): “Deep Learning for Revenue-Optimal Auctions with Budgets,” Tech. rep., Harvard SEAS.
- [13] GIANNAKOPOULOS, Y. (2015): “Bounding the optimal revenue of selling multiple goods,” *Theor. Comput. Sci.*, 581, 83–96.
- [14] GIANNAKOPOULOS, Y. AND E. KOUTSOPIAS (2014): “Duality and optimality of auctions for uniform distributions,” in *Proceedings of the 15th ACM Conference on Economics and Computation*, 259–276.
- [15] ——— (2015): “Selling Two Goods Optimally,” in *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming*, 650–662.
- [16] GOLOWICH, N., H. NARASIMHAN, AND D. C. PARKES (2017): “Deep Learning for Multi-Facility Location Mechanism Design,” Tech. rep., Harvard SEAS.
- [17] GUO, M. AND V. CONITZER (2009): “Worst-case optimal redistribution of VCG payments in multi-unit auctions,” *Games and Economic Behavior*, 67, 69–98.
- [18] ——— (2010): “Computationally Feasible Automated Mechanism Design: General Approach and Case Studies,” in *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*.
- [19] HART, S. AND N. NISAN (2012): “Approximate revenue maximization with multiple items,” in *Proceedings of the 13th ACM Conference on Economics and Computation*, 656.
- [20] MANELLI, A. AND D. VINCENT (2006): “Bundling as an optimal selling mechanism for a multiple-good monopolist,” *Journal of Economic Theory*, 127, 1–35.
- [21] MYERSON, R. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- [22] NARASIMHAN, H., S. AGARWAL, AND D. C. PARKES (2016): “Automated Mechanism Design without Money via Machine Learning,” in *Proc. 25th Int. Joint Conf. on Art. Intell.*, 433–439.
- [23] PAVLOV, G. (2011): “Optimal Mechanism for Selling Two Goods,” *B.E. Journ. of Theor. Econ.*, 11, 1–35.
- [24] PROCACCIA, A., A. ZOHAR, Y. PELEG, AND J. ROSENSCHEIN (2009): “The Learnability of Voting Rules,” *Artificial Intelligence*, 173, 1133–1149.
- [25] SANDHOLM, T. AND A. LIKHODEDOV (2015): “Automated Design of Revenue-Maximizing Combinatorial Auctions,” *Operations Research*, 63, 1000–1025.
- [26] THIRUMULANATHAN, D., R. SUNDARESAN, AND Y. NARAHARI (2016): “Optimal Mechanism for Selling Two Items to a Single Buyer Having Uniformly Distributed Valuations,” in *12th International Conference on Web and Internet Economics (WINE)*, 174–187.
- [27] YAO, A. C.-C. (2017): “Dominant-Strategy Versus Bayesian Multi-item Auctions: Maximum Revenue Determination and Comparison,” in *Proceedings of the 2017 ACM Conference on Economics and Computation, EC ’17*, 3–20.