

TOWARDS PRIOR-FREE APPROXIMATELY TRUTHFUL ONE-SHOT AUCTION LEARNING VIA DIFFERENTIAL PRIVACY*

Daniel Reusche
research@degreat.net

Nicolás Della Penna
nikete@mit.edu

ABSTRACT

Designing truthful, revenue maximizing auctions is a core problem of auction design. Multi-item settings have long been elusive. Recent work of Dütting et al. (2020) introduces effective deep learning techniques to find such auctions for the prior-dependent setting, in which distributions about bidder preferences are known. One remaining problem is to obtain priors in a way that excludes the possibility of manipulating the resulting auctions. Using techniques from differential privacy for the construction of approximately truthful mechanisms, we modify the RegretNet approach to be applicable to the prior-free setting. In this more general setting, no distributional information is assumed, but we trade this property for worse performance. We present preliminary empirical results and qualitative analysis for this work in progress.

1 INTRODUCTION

Auction design is a core problem of economic theory. The resulting auctions find practical applications for example in spectrum auctions for allocation of frequency bands to wireless carriers, commodity auctions, as well as online auctions platforms like eBay.

In the standard model of *independent private valuations*, bidders have valuations over items and utility dependent on the items allocated to them. The auctioneer does not know the realized valuations of the bidders, but has access to aggregate information in the form of distributions over valuations. Since valuations are private, incentivizing bidders to report them truthfully is important for finding revenue maximizing auctions.

Myerson (1981) presents an optimal (truthful and revenue maximizing) auction for the single-item multi-bidder setting, but the multi-item setting has long been elusive for reasons of computational intractability; For a survey on intractability, see the introduction of Rahme et al. (2020). The last 10 years have seen advances in partial characterizations, algorithmic results, albeit satisfying weaker notions than truthfulness, as well as applications of tools from machine learning and computational learning theory; A survey on these developments can be found in the introduction of Dütting et al. (2020).

Recently a line of work by Dütting et al. (2020) introduces deep learning techniques to find revenue maximizing, truthful (or: dominant strategy incentive compatible) auctions. It develops the RegretNet approach, where regret, a measure for incentive compatibility, is used to constrain the learning problem of finding revenue maximizing auctions by way of the augmented Lagrangian method. It is able to recover known solutions and finds low-regret solutions for multi-item settings. This approach is limited to the prior-dependent setting though, in which knowledge about the distributions of valuations is assumed. These distributions need to be aggregated in an incentive compatible manner, to prevent influence on the outcome of the mechanism by this avenue. Most recently Rahme et al. (2020) builds on RegretNet, but increases efficiency and applies reductions to get truthful auctions from low regret ones.

*We want to thank Olivia Röhrig for her invaluable contributions to increasing the clarity of this document, the reviewers of the DPML workshop at ICLR 2021 for their instructive feedback, as well as NLnet for its manifold organizational support. Code for the experiments is available at: <https://github.com/degreat/one-shot-approx-auctions/>

There also exist interesting connections between robust mechanism design and differential privacy: McSherry & Talwar (2007) introduce its use as a solution technique for mechanism design, by using the guarantees it provides to bound the influence agents can have on the outcome of mechanisms. These mechanisms are approximately truthful (and approximately optimal), since the bounded influence results in bounded incentives to misreport. Nissim et al. (2011) further analyze the potential for optimality approximation of mechanisms utilizing differential privacy.

1.1 CONTRIBUTIONS

To work towards an approach for solving the more general prior-free setting, in which no knowledge about valuation distributions is assumed, we integrate the above techniques.

By using the work of Abadi et al. (2016), we make the training of RegretNet differentially private, and thus robust to changes in distribution. This robustness allows us to remove the distribution requirement and enables us to use RegretNet on single bids profiles to perform one-shot learning, giving us one auction per bid profile.

Preliminary empirical analysis leads us to believe that the resulting auctions are approximately truthful, prior-free approximations of optimal mechanisms.

We present a qualitative analysis of computational experiments with the modified codebase and formulate theoretical problems, the solution of which we deem necessary for a thorough characterization.

2 OUR APPROACH

In this section we will give an explanation of the implementation details and the current state of results. We believe our approach to lead to an approximately truthful prior-free one-shot learner, giving agents only bounded incentive to misreport without depending on knowledge of valuation distributions. Thus every bid should be an ε -dominant strategy. Evidence is not yet conclusive, but preliminary results look promising. For background on the required theory see App. B.

2.1 PRIOR-FREE ONE-SHOT LEARNING

To work towards learning prior-free approximations of optimal auctions, we attempt to make the RegretNet approach approximately truthful by making the regret computation differentially private.

As we do not assume the existence of any valuation distribution, we can only elicit bids and proceed to do one-shot learning on them, to produce one auction per set of bids. We view the whole learner L as the mechanism to be analyzed, since we learn an auction A from a bid profile b and then apply it to the same bid profile to receive the outcome o : $L(b) = A$, then $A(b) = o$. From the perspective of the learner, the bids are treated as valuations.

Concretely, we make two modifications to RegretNet Training (Alg. (1) from Dütting et al. (2020)):

1. We take one set of reports and use it in every iteration of training. (By turning a bid profile into a δ -distribution and repeatedly sampling from it.)
2. We make the computation of the parameters w for the regret gradient differentially private, by using a differentially private optimizer (Alg. (1) from Abadi et al. (2016)).

If we would only do the first, we would always learn an auction that overfits the reports, meaning misreports could have a large influence. By introducing differential privacy, overfitting is reduced, resulting in bounded influence of each bidders reports. *Note:* We use uniform distributions $\mathcal{U}(0, 1)$

to generate misreports v'_i . Since we only use the samples to test a set of auction parameters w , this is permissible without assuming valuation distributions.

ALGORITHM 1: One-Shot RegretNet Training

Input: one bid profile b

Output: one auction $A = (g^w, p^w)$ with g^w and p^w being neural nets parametrized by w

Parameters: auction learning rate $\eta > 0$, misreport learning rate $\gamma > 0$, Lagrange update rates

$\forall t, \rho_t > 0$, noise scale σ , gradient norm bound C , $(\eta, \gamma, \rho_t, \sigma, C \in \mathbb{R})$, training steps T ,

misreport computation steps Γ , Lagrange update frequency Q , $(T, \Gamma, Q \in \mathbb{N})$, set of bidders N

Initialize: auction parameters $w^0 \in \mathbb{R}^d$, Lagrange multipliers $\lambda^0 \in \mathbb{R}^n$

for $t = 0, \dots, T$ **do**

Initialize misreports: $v'_i \sim \mathcal{U}(0, 1), i \in N$

for $r = 0, \dots, \Gamma$ **do**

forall $i \in N$ **do**

$v'_i \leftarrow v'_i + \gamma \nabla_{v'_i} u_i^w(b_i; (v'_i, b_{-i}))$

end

end

forall $i \in N$ **do**

Compute regret gradient:

$g_i^t = \nabla_w [u_i^w(b_i; (v'_i, b_{-i})) - u_i^w(b_i; b)] \Big|_{w=w^t}$

Clip regret gradient:

$\bar{g}_i^t \leftarrow g_i^t / \max\left(1, \frac{\|g_i^t\|_2}{C}\right)$

 /* Differential Privacy */

Add gaussian noise:

$\tilde{g}_i^t \leftarrow \bar{g}_i^t + \mathcal{N}(0, \sigma^2 C^2 \mathbf{I})$

end

Compute Lagrangian gradient using Eq. (1) and update w^t :

$w^{t+1} \leftarrow w^t - \eta \nabla_w C_{\rho_t}(w^t, \lambda^t)$

Update Lagrange multipliers once in Q iterations:

if t is a multiple of Q **then**

$\lambda_i^{t+1} \leftarrow \lambda_i^t + \rho_t \tilde{r} g_{t,i}(w^{t+1}), \forall i \in N$

else

$\lambda^{t+1} \leftarrow \lambda^t$

end

end

The Lagrangian function (Section 4 of Dütting et al. (2020))¹ is

$$\nabla_w C_{\rho}(w, \lambda^t) = - \sum_{i \in N} \nabla_w p_i^w(b) + \sum_{i \in N} \lambda_i^t g_i + \rho \sum_{i \in N} \tilde{r} g_{t,i}(w) g_i \quad (1)$$

$$\text{with } g_i = \nabla_w \left[\max_{v'_i \in V_i} u_i^w(b_i; (v'_i, b_{-i})) - u_i^w(b_i; b) \right].$$

The above algorithm returns one auction, consisting of an allocation g^w and payment function p^w , per set of reports.

2.2 REGRET COMPUTATION FOR ANALYSIS

To empirically evaluate whether this auction learner satisfies approximate truthfulness, and to measure the worst case approximation of an optimal auction, we need to compute a set of auctions with different sets of reports. Then we can calculate regret and worst case revenue or welfare over them.

Since we need to train one auction for each set of bids, any evaluation will be computationally expensive. We decided that the smallest meaningful test would be to measure the regret of a single

¹ The only difference to RegretNet is that $L = 1$.

misreporting agent for a sample of points from the valuation space. To calculate regret for one agent on a single point of the valuation space, we take a valuation sample, and compute the corresponding auction A_0 . We then construct a set of reports by enumerating all possible misreports $\forall v_1^{j'} \in V_1$ of the first agent, while keeping the valuation reports $v_{-1} = (v_2, \dots, v_n)$ of the other agents fixed. For each set of reports $(v_1^{j'}, v_{-1})$ we then compute an auction A_j , with $j \in J$ and $J = (1, \dots, |M \times \mathcal{V}|)$. Since we only handle additive valuations here, we model the valuation space per agent as $M \times \mathcal{V}$ with $\mathcal{V} = \{0, 1\}$ being the discrete valuations we permit, to make enumeration feasible.

The regret for the first agent, over all auctions $A_j \in \mathcal{A}$, with $A_j = (g^{w_j}, p^{w_j}), \forall j \in J$, utility $u_i^w(v_i; b) = v_i(g_i^w(b)) - p_i^w(b)$, and w_0 the initial parameters, then is

$$rgt_1(w) = \max_{j \in J} u_1^{w_j}(v_1; (v_1^{j'}, v_{-1})) - u_1^{w_0}(v_1; (v_1, v_{-1})) \tag{2}$$

ALGORITHM 2: Computing regret for One-Shot RegretNet

Input: Valuation space $\Theta = N \times M \times \mathcal{V}$ with $\mathcal{V} = \{0, 1\}$

Parameters: number of valuation samples S

for $s = 0, \dots, S$ **do**

Sample one set of valuations v **from** Θ

Compute one auction A_0 **on the valuations:**

$Alg. (1) \leftarrow (v)$

forall misreports $v_1^{j'} \in V_1$ **do**

Compute A_j **with one misreport** $v_1^{j'}$ **and valuations of the other agents** v_{-1} :

$Alg. (1) \leftarrow (v_1^{j'}, v_{-1})$

end

Calculate regret for $agent_1$ **using Eq. (2)**

end

2.3 QUALITATIVE ANALYSIS OF EMPIRICAL RESULTS

We can now do a qualitative analysis of the experimental results (figures in App. A): In the cases without privacy, misreports that outperform truthful reports exist. When applying differential privacy, we can observe the following stages with increasing noise, dependent on noise multiplier σ :

1. More and more misreports stop outperforming truthful reports. Outperforming misreports are visualized by the individual lines outside the main bundles of the misreporting agent in Fig. 1.
2. At some threshold regret bounds for truthfully reporting and misreporting agents align. This is marked by the bundles of the misreporting and truthfully reporting agent being of roughly the same width from $\sigma = 0.05$ onwards.
3. Regret bounds widen, as can be seen by the bundles of misreporting and truthfully reporting agents increasing in width. Revenue suffers, but its bounds tighten, which is marked by the revenue bundles reducing their width and flattening their slope in Fig. 2. This can also be seen in the steady decline of revenue in Tab. 2.

This is to be expected, since with differentially private training we exchange privacy for accuracy of the resulting model (see Figure 4, Abadi et al. (2016)). In the frameworks of McSherry & Talwar (2007); Nissim et al. (2011) privacy controls approximation of truthfulness, accuracy controls outcome quality.

In cases with very small valuation spaces, the technique does not provide reasonable tradeoffs, most likely since noise quickly outweighs available information. Once our pipeline supports finer valuation spaces, we will be able to analyze this in more depth.

All in all, these preliminary results are in line with our hypothesis of being able to train prior-free approximations of optimal mechanisms which are approximately truthful. As this is a work in progress, we have not yet achieved general quantifications for the degrees of approximation for optimality and truthfulness. Producing conclusive evidence will require further investigation.

3 FURTHER WORK

3.1 IMPROVING EMPIRICAL ANALYSIS

Since experiments are still comparatively expensive, we don't have good data yet for larger instances, finer valuation spaces or longer training. Application of the technique above to Rahme et al. (2020), if successful, could make these experiments feasible. An implementation in TensorFlow 2 would allow us to make use of the new SIMD features of TensorFlow Privacy (Subramani et al. (2020)), further reducing training time.

We will also evaluate whether it is possible to use an approximately truthful bid elicitation step for the online algorithm. This might lead to a prior-independent solution, further improving efficiency by being able to train auctions for the approximated distributions, instead of one auction per bid sample.

It might also lead to improved outcomes, if the online learning can be used to turn non-truthful reports into no longer approximately dominant strategies, while preserving approximate dominance for truthful reports (as described on p.3, Lecture 1, Roth (2014)).

3.2 APPROXIMATION BOUNDS

Describing bounds for approximation of incentive compatibility, as well as prior-free performance is an open problem for this work and related to the description of generalization bounds for the learning system.

Our current hypothesis is, that to describe meaningful approximation bounds, it is necessary to come up with generalization bounds for iterated, (ϵ, δ) -differentially private learning, which are also sensitive to the information capacity of the hypothesis space. To our knowledge, this is still an open problem.

Since the results of the above model depend on the auction setting (different networks are used to model different auctions), as well as on the parameters of the differentially private learning, we can not reuse the generalization bounds from the original model (Sections 2.4 and 3.3 of Dütting et al. (2020)). To get meaningful generalization bounds, one approach would be to extend covering number based techniques, as used in Section 2.4 Dütting et al. (2020) to also account for the learning algorithm that is being used, especially in regard to its privacy, as in He et al. (2020). The resulting technique should bind on the capacity of the hypothesis space, as well as on the sensitivity of the learning.

3.3 PRIVACY AWARE AGENTS

Another interesting topic of further investigation is the influence the technique has on incentive compatibility if the agents are privacy aware (Nissim et al., 2012).

ACKNOWLEDGEMENTS

Daniel Reusche was funded through the NGIO PET Fund, a fund established by NLnet with financial support from the European Commission's Next Generation Internet programme, under the aegis of DG Communications Networks, Content and Technology under grant agreement No 825310.

REFERENCES

- Martín Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. Oct 2016. doi: 10.1145/2976749.2978318. URL <http://arxiv.org/abs/1607.00133v2>. Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (ACM CCS), pp. 308-318, 2016.
- Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9:211–407, 2014. URL <https://www.cis.upenn.edu/~aaroht/Papers/privacybook.pdf>.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, 2006. doi: 10.1007/11681878_14. URL <https://iacr.org/archive/tcc2006/38760266/38760266.pdf>.
- Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David C. Parkes, and Sai Srivatsa Ravindranath. Optimal auctions through deep learning. Aug 2020. URL <http://arxiv.org/abs/1706.03459v5>.
- Jason Hartline. *Manuscript Mechanism Design and Approximation*. 2016. URL <http://jasonhartline.com/MDnA/>.
- Fengxiang He, Bohan Wang, and Dacheng Tao. Tighter generalization bounds for iterative differentially private learning algorithms. Aug 2020. URL <http://arxiv.org/abs/2007.09371v2>.
- Frank McSherry. Privacy integrated queries. *Communications of the ACM*, 53(9):89–97, September 2010. doi: 10.1145/1810891.1810916. URL <https://www.microsoft.com/en-us/research/wp-content/uploads/2010/09/pinq-CACM.pdf>.
- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pp. 94–103, 2007. doi: 10.1109/FOCS.2007.66. URL <http://kunalatalwar.org/papers/expmech.pdf>.
- Roger Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- Kobbi Nissim, Rann Smorodinsky, and Moshe Tennenholtz. Approximately optimal mechanism design via differential privacy. Mar 2011. URL <http://arxiv.org/abs/1004.2888v4>.
- Kobbi Nissim, Claudio Orlandi, and Rann Smorodinsky. Privacy-aware mechanism design. Feb 2012. URL <http://arxiv.org/abs/1111.3350v2>.
- Jad Rahme, Samy Jelassi, and S. Matthew Weinberg. Auction learning as a two-player game. Jun 2020. URL <http://arxiv.org/abs/2006.05684v2>.
- Aaron Roth. Lecture notes on differential privacy in game theory and mechanism design, Spring 2014. URL <https://www.cis.upenn.edu/~aaroht/courses/gametheoryprivacyS14.html>.
- Pranav Subramani, Nicholas Vadivelu, and Gautam Kamath. Enabling fast differentially private SGD via just-in-time compilation and vectorization. Oct 2020. URL <http://arxiv.org/abs/2010.09063v1>.

A FIGURES AND TABLES

In the following figures, each experiment is visualized by a single line. One experiment corresponds to the application of the auction learner to one valuation sample from the valuation space. For each sample we enumerate the misreports of one agent and train one auction each to calculate maximal regret as well as minimal revenue.

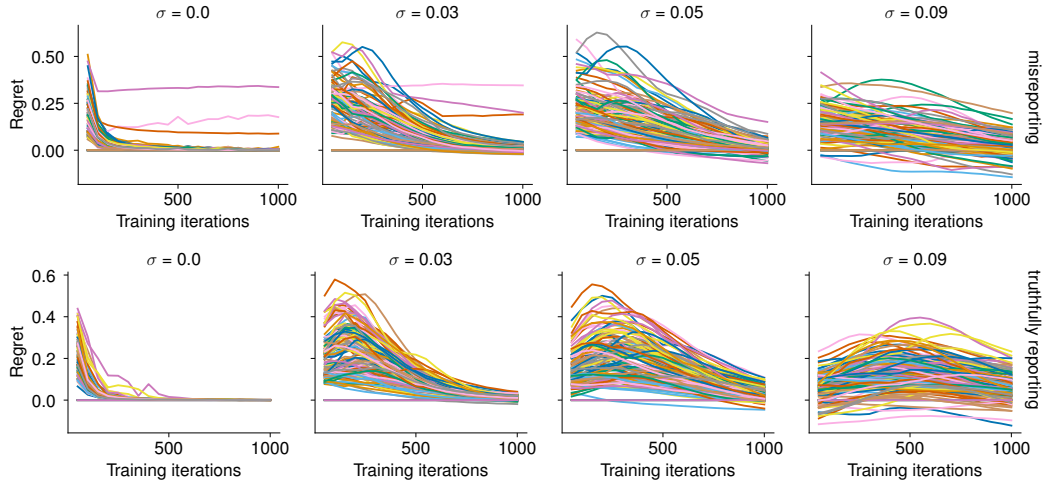


Figure 1: Max regret of a misreporting and a truthfully reporting bidder (5 bidders, 3 items)




Figure 2: Min revenue (5 bidders, 3 items)

Learner	Prior	Agents	Items	Regret	Revenue	Iterations	σ
RegretNet	uniform	5	3	0.0022	2.10	40000	N/A
one-shot RegretNet	none	5	3	up to 0.0330	1.71 - 2.57	1000	0.05

Table 1: Comparison of prior-free and prior-dependent results

σ	Regret $agent_1$	min Revenue overall
no DP	0.337	1.993
0.03	0.346	1.857
0.05	0.151	1.713
0.09	0.198	1.191

Table 2: Approximation of truthfulness and approximation of outcome optimality depending on σ

B BACKGROUND

As mentioned in the introduction, it is possible to learn multi-item auctions in the prior-dependent setting. What follows is a more in depth explanation of the techniques we have used to tackle the prior-free setting.

B.1 AUCTION DESIGN

The **multi-item auctions** we will consider consist of sets $N = \{1, \dots, n\}$ and $M = \{1, \dots, m\}$ of n bidders and m items. Each bidder i has valuations $v_i(\{j\})$ for all items j . We focus on bidders with **additive valuations**², where valuations for sets of items $S \subseteq M$ are $v_i(S) = \sum_{j \in S} v_i(\{j\})$.

In the **prior-dependent setting**, the auctioneer does not know the valuation profile $v = (v_1, \dots, v_n)$ of the agents in advance, but has prior knowledge about the distribution F from which v is drawn. $V = \mathbb{R}^{n \times m}$ is the space of possible valuations v and bids b , with $v_i, b_i \in V_i$ and $V_i = \mathbb{R}^m$. Bidders report their bids, $b = (b_1, \dots, b_n)$ with $b_i = v_i$ being a **truthful report** and $b_i \neq v_i$ being a **misreport** of agent i .

Each auction is defined by a pair of allocation and payment rules (g, p) with $g_i : V \rightarrow [0, 1]^m$ giving allocation probability of each item to agent i and $p_i : V \rightarrow \mathbb{R}_{\geq 0}$ giving the necessary payment. Any item is allocated at most once: $\sum_i g_i(b) \leq 1$ for all $b \in V$. Bidder i , with valuation v_i , receives **utility** $u_i(v_i; b) = v_i(g_i(b)) - p_i(b)$ for a set of bids b from all bidders. The **revenue** of an auction is $\sum_{i \in N} p_i(v)$.

Further, let v_{-i} be the valuation profile v without v_i . b_{-i} and V_{-i} are used analogously. An auction is **dominant strategy incentive compatible (DSIC)** if a bidder's utility is maximized when reporting truthfully, independent of the bids of others: $v'_i \neq v_i$ being a misreport, $u_i(v_i; (v_i, b_{-i})) \geq u_i(v_i; (v'_i, b_{-i}))$ holds for all $v_i, v'_i \in V$ and bids of others $b_{-i} \in V_{-i}$.

An auction is ex post **individually rational (IR)** if each bidder always receives non-negative utility: $u_i(v_i; (v_i, b_{-i})) \geq 0$ for all $v_i \in V_i, b_{-i} \in V_{-i}$.

Optimal auction design seeks a DSIC auction that maximizes revenue and is IR.

B.2 OPTIMAL AUCTION DESIGN WITH DEEP LEARNING

To translate optimal auction design into a **learning problem** (Section 2.2.2 of Dütting et al. (2020)) we take a parametrized class of auctions (g^w, p^w) , with parameters $w \in \mathbb{R}^d, d \in \mathbb{N}$.

With expected ex post **regret** being the utility gain for optimal misreports, and $u_i^w(v_i; b) = v_i(g_i^w(b)) - p_i^w(b)$, we can measure the deviation from DSIC:

$$rgt_i(w) = \mathbf{E}[\max_{v'_i \in V_i} u_i^w(v_i; (v'_i, v_{-i})) - u_i^w(v_i; (v_i, v_{-i}))]. \quad (3)$$

Thus an auction is DSIC iff $rgt_i(w) = 0, \forall i \in N$, except for measure zero events. We then minimize expected negated revenue, subject to a regret constraint for each bidder and F being the valuation distribution:

$$\min_{w \in \mathbb{R}^d} \mathbf{E}_{v \sim F} \left[- \sum_{i \in N} p_i^w(v) \right] \text{ s.t. } rgt_i(w) = 0 \forall i \in N, v \in V \quad (4)$$

For implementation with a deep learning pipeline, we formulate the **empirical regret** for a sample of L valuation profiles as

$$\widehat{rgt}_i(w) = \frac{1}{L} \sum_{l=1}^L \left[\max_{v'_i \in V_i} u_i^w(v_i^{(l)}; (v'_i, v_{-i}^{(l)})) - u_i^w(v_i^{(l)}; v^{(l)}) \right] \quad (5)$$

as well as the empirical loss, which we want to minimize, subject to an empirical regret constraint:

$$\min_{w \in \mathbb{R}^d} - \frac{1}{L} \sum_{l=1}^L \sum_{i=1}^n p_i^w(v^{(l)}) \text{ s.t. } \widehat{rgt}_i(w) = 0 \forall i \in N, v \in V \quad (6)$$

² For a more general exposition see Section 2 of Dütting et al. (2020). Section 3 of Rahme et al. (2020) shows another description of the additive setting.

The allocation and payment rules are modeled as neural networks (Section 3.2 of Dütting et al. (2020)), which ensure IR by restricting to auctions which don't charge any bidder more than their valuations for any allocation. This implementation is called RegretNet.

The regret constraint can be incorporated into the objective by using the technique of Lagrange multipliers (Section 4 of Dütting et al. (2020)).

B.3 PRIOR-FREE APPROXIMATIONS TO OPTIMAL MECHANISMS

A **prior-free mechanism** (Section 7 of Hartline (2016)) is a mechanism (Section 9.4 of Nisan et al. (2007)) which does not make assumptions about the valuation distribution of agents. In the case of auction design, this would be the distribution F of the valuations v of the bidders. Since it has a weaker informational requirement, this type of mechanism can be applied in a larger variety of settings.

A mechanism \mathcal{M} is a **prior-free β -approximation** to a benchmark (Def. 7.1. of Hartline (2016)) \mathcal{B} if for all valuation profiles v its performance is at least a β fraction of the benchmark: $\mathcal{M}(v) \geq \frac{1}{\beta}\mathcal{B}(v)$.

In Dütting et al. (2020) auctions are the mechanisms of interest, in our setting though, we seek to analyze the whole one-shot learner.

B.4 DIFFERENTIAL PRIVACY

Differential privacy (Dwork et al., 2006)³ gives us strong guarantees on the distinguishability of the outcomes of a mechanism executed on adjacent datasets. A randomized mechanism $\mathcal{M}: \mathcal{D} \rightarrow \mathcal{R}$ with domain \mathcal{D} and range \mathcal{R} satisfies **(ϵ, δ) -differential privacy** if for any two adjacent inputs $d, d' \in \mathcal{D}$ and for any subset of outputs $S \subseteq \mathcal{R}$ it holds that

$$\Pr[\mathcal{M}(d) \in S] \leq e^\epsilon \Pr[\mathcal{M}(d') \in S] + \delta \quad (7)$$

For the auction design context, we take adjacent datasets to be bid profiles which differ in one entry, i.e. it would be contained in one and missing from the other.

A single application of a **gaussian noise mechanism**

$$\mathcal{M}(d) \triangleq f(d) + \mathcal{N}(0, S_f^2 \cdot \sigma^2) \quad (8)$$

to any function f , with sensitivity S_f being the amount any single argument can change its output, and σ the noise multiplier, satisfies (ϵ, δ) -differential privacy if $\delta \geq \frac{4}{5}e^{-\frac{\epsilon^2}{2}}$ and $\epsilon \geq 1$. Since the analysis can be applied post hoc, there are infinite (ϵ, δ) pairs s.t. this is fulfilled.

B.5 DEEP LEARNING WITH DIFFERENTIAL PRIVACY

To introduce differential privacy to deep learning, we can use a differentially private version of SGD, which has two extra steps between the computation of per example gradients $g(x_i)$ and descent. The gradients get clipped to the gradient norm bound C (which bounds sensitivity S_f):

$$\bar{g}(x_i) = g(x_i) / \max\left(1, \frac{\|g(x_i)\|_2}{C}\right) \quad (9)$$

Then noise is added to the gradient of each iteration, with σ being the noise multiplier, L being the group size of the sample from the whole dataset:

$$\tilde{g} = \frac{1}{L} \left(\sum_i \bar{g}(x_i) + \mathcal{N}(0, \sigma^2 C^2 \mathbf{I}) \right) \quad (10)$$

This way, training of neural networks can be modelled as a composition of single applications of the gaussian mechanism Eq. (8). For details see Section 3.1. of Abadi et al. (2016).

³ The definitions used here are lifted from Abadi et al. (2016). For a more in depth exposition of (ϵ, δ) -differential privacy see Def. 2.4 of Dwork & Roth (2014) and onwards.

The bookkeeping of the privacy budget resulting from the repeated application of gaussian mechanisms is handled with a moments accountant, which is introduced in McSherry (2010), Section 3.2. of Abadi et al. (2016) describes its application to deep learning.

B.6 MECHANISM DESIGN VIA DIFFERENTIAL PRIVACY

McSherry & Talwar (2007) introduce differential privacy as a solution concept for mechanism design problems.

For any mechanism \mathcal{M} , truthful reporting is an ϵ -**approximately dominant strategy** (Definition 10.2 of Dwork & Roth (2014)) for player i if for every pair of types, $t_i, t'_i \in T$, t_i being the private information player i holds, and for every vector of types t_{-i} from the other players, and utility function u the following holds (in the auction setting, types are valuations):

$$u(t_i, \mathcal{M}(t_i, t_{-i})) \geq u(t_i, \mathcal{M}(t'_i, t_{-i})) - \epsilon \quad (11)$$

Using Eq. (7), at the expected utility for any $(\epsilon, 0)$ -differentially private mechanism \mathcal{M} and any non-negative function g of its range, with $d, d' \in \mathcal{D}$ differing only in one data point

$$\mathbf{E}[g(\mathcal{M}(d))] \leq e^\epsilon \mathbf{E}[g(\mathcal{M}(d'))] \quad (12)$$

we can derive that $(\epsilon, 0)$ -differentially private mechanisms being $(\epsilon, 0)$ -**approximately dominant strategy truthful** (Section 2.1 of McSherry & Talwar (2007)), for $\epsilon \leq 1$ and utilities bounded in $[0, 1]$ (Full proof: Lecture 1, Claim 6, Roth (2014)).

This means, given the above constraints, when using an $(\epsilon, 0)$ -differentially private mechanism, no user can cause a change of more than ϵ in their utility.