

Objective Variables for Probabilistic Revenue Maximization in Second-Price Auctions with Reserve

Maja R. Rudolph
Columbia University
New York, NY
maja@cs.columbia.edu

Joseph G. Ellis
Columbia University
New York, NY
jge2105@columbia.edu

David M. Blei
Columbia University
New York, NY
david.blei@columbia.edu

ABSTRACT

Many online companies sell advertisement space in second-price auctions with reserve. In this paper, we develop a probabilistic method to learn a profitable strategy to set the reserve price. We use historical auction data with features to fit a predictor of the best reserve price. This problem is delicate—the structure of the auction is such that a reserve price set too high is much worse than a reserve price set too low. To address this we develop *objective variables*, an approach for combining probabilistic modeling with optimal decision-making. Objective variables are “hallucinated observations” that transform the revenue maximization task into a regularized maximum likelihood estimation problem, which we solve with the EM algorithm. This framework enables a variety of prediction mechanisms to set the reserve price. As examples, we study objective variable methods with regression, kernelized regression, and neural networks on simulated and real data. Our methods outperform previous approaches both in terms of scalability and profit.

General Terms

Machine Learning; Economics

Keywords

Online Auctions; Second-Price Auctions; Graphical Models

1. INTRODUCTION

Many online companies earn money from auctions, selling advertisement space or other items. One of the most common auction mechanisms is second-price auctions with reserve [10]. This type of auction works as follows. First, the company sets a *reserve price*, the minimal price at which they are willing to sell the item. Then the auction is opened and potential buyers cast their bids; they cannot see each other’s bids. Finally, the auction is completed. If the highest bid is smaller than the reserve price then there is no transaction and the company does not earn any money. If any bid is larger than the reserve price then the highest bidding buyer wins the auction, paying the larger of the second-highest bid and

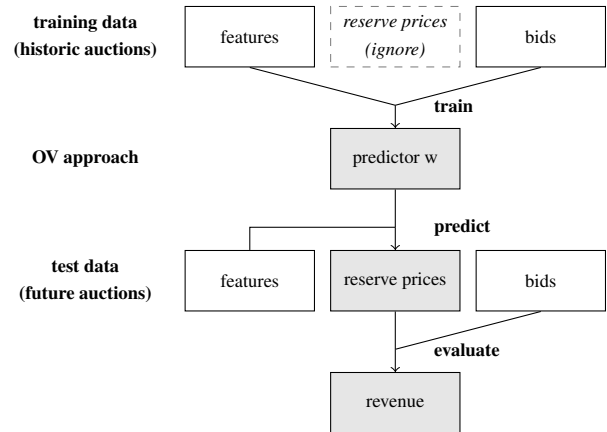


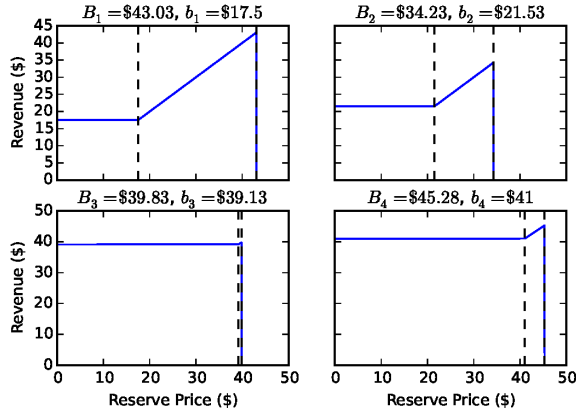
Figure 1: The objective variable approach uses features and bids from past auctions to learn a predictor of profitable reserve prices.

the reserve price. This mechanism is incentive compatible [3]—it rewards buyers who bid their true idea of what the item is worth.


A company that runs such auctions faces the important challenge of how to set the reserve price of each item in order to maximize profit. Specifically, a company maximizes its profit by setting the reserve as close as possible to the (future and unknown) highest bid, but no higher. This gives rise to an important machine learning challenge [13]—how can we use historical data to map the properties of a new item to a good reserve price? In this paper, we develop a new method to solve this problem.

Figure 1 gives a schematic of our method. We are given historical data about past auctions, namely features of the items and the two highest bids. These features might be properties of the product, such as the placement of the advertisement, properties of the potential buyers, such as each one’s average past bids, or other external features, such as time of day of the auction. Given the data, our method learns a predictor of reserve price that maximizes the profit of future auctions. Finally, we use the predictor to set the reserve price of new items. Note that our method only requires past items and bids. It does not require previously set reserve prices.

A traditional machine learning solution treats this problem as simple prediction: given a set of features x , use a parameterized predictor $f(x; w)$ to predict the high bids B (because setting the reserve to the high bid promises the maximum possible profit). There are many ways of solving such prediction problems, including linear regression, regularized linear regression, or nonlinear methods such as a neural network. At a high level, these methods all fit the parameters w to minimize a symmetric error function, such as mean



(a) Revenue function of 4 auctions from the eBay data set as a function of reserve price. In second-price auctions with reserve the revenue depends on the highest and the second highest bid (dashed lines).



(b) A scatter plot of the first and second highest bid in a test set of 2000 auctions taken from the eBay data set. Three example auctions are highlighted in red to emphasize the nuances of the reserve price prediction task.

Figure 2: Example auctions from the eBay data set we study in Section 3.

squared error, between a historical data set of true bids and the corresponding predictions. Note that the symmetric error function equally weights overestimating and underestimating the prediction.

But predicting a good reserve price is more delicate than this. The reason is that the revenue function for each auction—the amount of money that we make as a function of the reserve price y —is *asymmetric*. It remains constant to the second bid b , increases to the highest bid B , and is zero beyond the highest bid. Formally,

$$R(y, B, b) = \begin{cases} b & \text{if } y < b \\ y & \text{if } b \leq y \leq B \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

The revenue function reveals two nuances to forming a good prediction of the best reserve price. First, from the perspective of profit, overestimating the reserve price can be substantially worse than underestimating it. Thus a predictor should err on the side of underestimating the highest bid. Second, it is the size of the gap between the second bid and the highest bid that determines the potential to profit. If the gap is small then the company can comfortably underestimate the highest bid in order to receive the (almost optimal) second highest bid. Thus we would like a predictor to focus on cases where the gap tends to be large.

Figure 2a illustrates the revenue function for four sports collectibles from eBay, each one a function of the first and second bids. This figure puts the delicacy into relief. For example, consider the top left panel in Figure 2a, which might be a Sandy Koufax baseball card. (The items in our data are anonymized, but we use this example for concreteness.) The best reserve price in retrospect is \$43.03. A method that minimizes squared error (such as linear regression) is just as likely to overestimate as to underestimate. It fails to reflect that setting the price in advance to \$44.00 would yield zero earnings while setting it to \$40.00 would yield the full reserve.

Now consider the first panel in the second row. Here the gap between the first and second bid is small. For this auction, we profit nearly as much setting the risk-free reserve price of zero as trying to accurately predict the highest bid (which runs the risk of overestimation). While a technique like linear regression treats this auction equally important to the Sandy Koufax auction, a better predictor will focus on predicting well for items with a large gap.

In this paper, we develop the *objective variables approach* (OV approach) to setting the reserve price. Our method fits the parameters of a common predictor, such as linear regression or a neural network, but uses an alternative fitness function that implicitly accounts for the special structure of the revenue. (We describe the technical approach in more detail below.) We found that OV approach outperforms the existing state of the art.

Example results on eBay. In Section 3 we will study a large collection of eBay auctions of sports collectibles. (Figure 2a was taken from these data.) Figure 2b illustrates the highest and second highest bid in the test set, items that we did not observe when fitting the model. There is large variability in the bids as well as in the potential for profit in the different auctions.

Consider the three auctions in Figure 2b which are marked in red. Auction A is the most important type of auction. The margin between the highest and second highest bid is large, and so a good predictor will set the reserve price close to the highest bid. There is a smaller margin between the bids in auctions B and C. In these auctions, setting the reserve price close to zero is a better option—it still yields the second highest bid as revenue, while reducing the risk of losing the auction by overshooting the highest bid.

Figure 3 compares how different approaches for predicting the reserve price cope with the challenge. Each point in the scatter plot corresponds to an auction and is drawn at the predicted reserve price for this auction and with height equal to the highest bid. The grey scale indicates the profitability of each prediction, drawn in terms of the percentage of highest possible revenue (highest bid).

For all points below the diagonal the prediction is larger than the highest bid, resulting in 0 revenue. Points well above the diagonal are either bright when they resulted in revenue well less than the highest bid or dark. When they are dark, this is because the margin is small. The auction winner pays the second bid even though the predicted reserve price underestimates it.

The three panels compare three methods for setting the reserve price from features: linear regression with the highest bids as the response variables (left panel), DC [13], the current state of the art for setting reserve prices from features (middle panel), and OV neural network, a method we develop in this paper (right panel). The predictors were trained on a separate sample of 2,000 auctions.




Figure 3: Comparison of three methods for predicting the reserve price from features: linear regression against the highest bid (left panel), DC [13] (current state of the art, middle panel), and OV neural network (this paper, right panel.) Each point in the scatter plot corresponds to an auction in the test set of 2000 data points from the eBay data set. For each prediction we compute the revenue using Eq. 1. The grey scale corresponds to the fraction of the highest possible revenue each prediction generated (revenue divided by highest bid.) For each method we also report the total revenue generated as percentage of the highest possible revenue.

The left panel is linear regression. It treats all auctions equally and its loss function, the mean squared error, does not distinguish between overestimating and underestimating the reserve price. New predictions are as likely to lie below the highest bid as above the highest bid, and this results in 0 revenue in approximately half of the auctions. For this reason, the total revenue as a percentage of the total highest possible revenue is 49%.¹

The middle panel is the DC method [17, 13]. It is much more conservative than linear regression, hardly ever overestimating the reserve, and thus produces significantly improved results over linear regression (58% of the highest possible revenue.) The right panel is OV neural networks, one of the objective variable models we present here. It is less conservative than DC, but only with auctions that have small potential profit (i.e., a small highest bid). In exchange for its bolder predictions, it outperforms both methods on auctions with larger potential profit, and outperforms them overall.

We also note that the objective variables method scales better than DC, its closest competitor among the previous methods. In Section 3 we analyze a much larger set of the eBay data, which the DC algorithm is unable to process.

Technical summary. To solve this problem we develop a new idea, the *objective variable*. Objective variables use probabilistic models to reason about difficult prediction problems, such as one that optimizes Eq. 1. Specifically, objective variables enable us to formulate probabilistic models for which MAP estimation directly uncovers profitable decision-making strategies.

Our aim is to find a parameterized mechanism $f(x_i; w)$ to set the reserve price from the auction features x_i . In our study, we will consider a linear predictor, kernelized regression, and a neural network. We observe a historical data set of N auctions, where each contains features x_i and the auction’s two highest bids B_i and b_i . We would like to learn a good mechanism by optimizing the parameter w to maximize the total (retrospective) revenue $\sum_{i=1}^N R(f(x_i; w), B_i, b_i)$.

We solve this optimization problem by turning it into a maximum a posteriori (MAP) problem. For each auction we define new

binary variables—these are the objective variables—that are conditional on a reserve price. The probability of the objective variable being on (i.e., equal to one) is related to the revenue obtained from the reserve price; it is more likely on if the auction produces more revenue. We then set up a model that first assumes each reserve price is drawn from the parameterized mechanism $f(x_i; w)$ and then draws the corresponding objective variable. Note that this model is defined conditioned on our data, the features and the bids. It is a model of the objective variables.

With the model defined, we now imagine a “data set” where all of the objective variables are on, and then fit the parameters w subject to these data. Because of how we defined the objective variables, the model will prefer more profitable settings of the parameters. With this set up, fitting the parameters by MAP estimation is equivalent to finding the parameters that maximize revenue.

The spirit of our approach is that the objective variables are likely to be on when we make good decisions, that is, when we profit from our setting of the reserve price. When we imagine that they are all on, we are imagining that we made good decisions (in retrospect). When we fit the parameters to these data, we are using MAP estimation to find a mechanism that helps us make such decisions.

Related work. Second-price auctions with reserve are reviewed in [10]. Ref. [14] empirically demonstrates the importance of optimizing reserve prices; Their study quantifies the positive impact it had on Yahoo!’s revenue. However, most previous work on optimizing the reserve price are limited in that they do not consider features of the auction [14, 8].

Our work builds on the ideas in Ref. [13]. This research shows how to learn a linear mapping from auction features to reserve prices, and demonstrates that we can increase profit when we incorporate features into the reserve-price setting mechanism. We take a probabilistic perspective on this problem, and show how to incorporate nonlinear predictors. We show in Section 3 that our algorithms scale better and perform better than these approaches.

In our method we construct a graphical model to solve the revenue maximization problem. One part of the model corresponds to the reserve price setting mechanism, while the objective variables who ensure that the predicted reserve prices maximize revenue can be seen as a separate part of the model. The inherent modularity of

¹We compute this percentage by dividing the total revenue from all predictions by the sum of the highest bids and multiplying by one hundred. An oracle would generate the highest possible revenue by anticipating the highest bids.

graphical models allows us to easily substitute the reserve price setting mechanism with different alternatives. In this work, we study linear regression, kernelized regression, and neural networks [6].

The objective variable framework relates closely to recent ideas from reinforcement learning [19, 18]. Reinforcement learning seeks an action policy that maximizes expected future reward. Refs. [19, 18] introduce a binary reward variable (similar to an objective variable) and use maximum likelihood estimation to find such a policy. Our work solves a different problem with similar ideas. Further, reinforcement learning tends to focus on simple discrete policies; we use objective variables for continuously parameterized predictors.

Organization of this paper. We first develop the OV approach for linear predictors and show how to use the expectation-maximization algorithm [9] to solve the corresponding MAP problem. We then generalize our method to nonlinear predictors, such as kernel regression and neural networks. Finally, on both simulated data and real-world data from eBay, we show that the OV approach outperforms existing methods for setting the reserve price. It is both more profitable and more easily scales to large data sets.

2. OBJECTIVE VARIABLES FOR SECOND-PRICE AUCTIONS WITH RESERVE

We describe the problem setting and the objective function. Our data come from previous auctions. For each auction, we observe features x_i , the highest bid B_i , and the second bid b_i . The features represent various characteristics of the auction, such as the date, time of day, or properties of the item. For example, one of the auctions in the eBay sport collectibles data set might be for a Sandy Koufax baseball card; its features include the date of the auction and various aspects of the item, such as its condition and the average price of such cards on the open market.

When we execute an auction we set a reserve price before seeing the bids; this determines the revenue we receive after the bids are in. The revenue function (Eq. 1), which is indexed by the bids, determines how much money we make as a function of the chosen reserve price. We illustrate this function for 4 auctions from eBay in Figure 2a. Our goal is to use the historical data to learn how to profitably set the reserve price from auction features, that is, before we see the bids.

For now we will use a linear function to map auction features to a good reserve price. Given the feature vector x_i we set the reserve price with $f(x_i; w) = w^\top x_i$. (In Section 2.4 we consider nonlinear alternatives.) We fit the coefficients w from data, seeking w that maximizes the regularized revenue

$$w^* = \arg \max \sum_{i=1}^N R(f(x_i; w), B_i, b_i) - (\lambda/2) w^\top w. \quad (2)$$

We chose L_2 regularization with parameter λ ; other regularizers are possible.

Before we discuss how we solve this optimization, we make two points. First, previous reserve prices are *not* included in the data. Rather, our data tell us about the relationship between features and bids. All the information about how much we might profit from the auction is in the revenue function; the way previous sellers set the reserve prices is not relevant.

Second, our goal is not the same as learning a mapping from features to the highest bid. As we described in the introduction, not all auctions are made equal. Consider the difference between the top left auction in Figure 2a and the bottom left auction. On top left the highest and second bid are far apart, $B = \$43.03$ and $b = \$17.5$; on bottom left they are almost identical, $B = \$39.83$ and $b = \$39.17$. The margin between bids in the first auction is

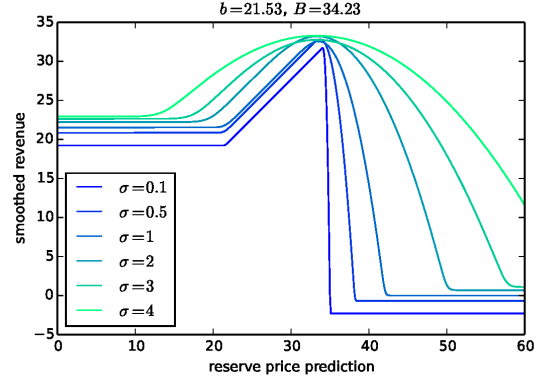


Figure 4: The effect of smoothing on the revenue function of an auction from the eBay data set. The smaller σ , the closer the smoothed revenue approximates the actual revenue function.

much larger than for the second, so there is more room to profit by correctly setting the reserve price. We account for this by directly maximizing revenue, rather than by only predicting the highest bid.

2.1 The Smoothed Revenue

The optimization problem in Eq. 2 is difficult to solve because $R(\cdot)$ is discontinuous (and thus non-convex). Previous work [13] addresses this problem by iteratively fitting differences of convex (DC) surrogate functions and solving the resulting DC-program [17]. Here we define an objective function related to the revenue, but that smooths out the troublesome discontinuity. We will optimize this objective with expectation-maximization.

We first place a Gaussian distribution on the reserve price centered around the linear mapping, $y_i \sim \mathcal{N}(f(x_i; w), \sigma^2)$. We define the *smoothed regularized revenue* as

$$\mathcal{L}(w) = \sum_{i=1}^N \log \mathbb{E}_{y_i} [\exp \{R(y_i, B_i, b_i)\}] - (\lambda/2) w^\top w. \quad (3)$$

Figure 4 shows the first term of Eq. 3 for one of the auctions in our data set. This figure illustrates how the smoothed revenue becomes closer to the original revenue function as σ^2 decreases. This approach was inspired by probit regression, where a Gaussian expectation is introduced to smooth the discontinuous 0-1 loss [2, 12].

The smoothed revenue is a well-defined and continuous objective function. In principle, we could use its gradient to fit the parameters of the predictor. However, we will take a different approach. We first recast the problem as a regularized likelihood under a latent variable model, and then using the expectation-maximization (EM) algorithm [9] to maximize that likelihood. This strategy enjoys closed-form computation in both the E and M steps, and facilitates replacing linear regression with more powerful nonlinear predictors.

2.2 Objective variables

To reformulate the optimization problem, we introduce the idea of the *objective variable*. An objective variable is part of a probabilistic model for which maximum-a-posteriori (MAP) estimation of w recovers the parameter that maximize the smoothed revenue in Eq. 3.

Objective variables are binary variables z_i for each auction, each conditioned on the reserve price y_i , the highest bid B_i , and second bid b_i . We can interpret these variables to indicate “Is the auction host satisfied with the outcome?” Concretely, the likelihood of sat-

isfaction is related to how profitable the auction was relative to the maximum profit, $p(z_i = 1 | y_i, B_i, b_i) = \pi(y_i, B_i, b_i)$ where

$$\pi(y_i, B_i, b_i) = \exp\{-(B_i - R(y_i, B_i, b_i))\}. \quad (4)$$

The revenue function $R(\cdot)$ is in Eq. 1. The revenue is bounded by B_i ; thus the probability is in $(0, 1]$.

What we will do is set up a probability model around the objective variables, assume that they are all “observed” to be equal to one (i.e., we are satisfied with all of our auction outcomes), and then fit the parameter w to maximize the posterior conditioned on this “hallucinated data.”

Concretely, the model is

$$w \sim \mathcal{N}(0, \lambda^{-1} I) \quad (5)$$

$$y_i | w, x_i \sim \mathcal{N}(f(x_i; w), \sigma^2) \quad i \in \{1, \dots, N\} \quad (6)$$

$$z_i | y_i, B_i, b_i \sim \text{Bernoulli}(\pi(y_i, B_i, b_i)) \quad (7)$$

where $f(x_i; w) = x_i^\top w$ is a linear map. Figure 5a shows the graphical model.

Now consider a data set \mathbf{z} where all of the objective variables z_i are equal to one. Conditional on this data, the log posterior of w marginalizes out the latent reserve prices y_i ,

$$\begin{aligned} \log p(w | \mathbf{z}, \mathbf{x}, \mathbf{B}, \mathbf{b}) &= \sum_{i=1}^N (\log \mathbb{E}[\exp\{R(y_i, B_i, b_i)\}] - B_i) \\ &\quad + \log p(w | \lambda) - C, \end{aligned} \quad (8)$$

where C is the normalizer.

This is the smoothed revenue of Eq. 3 plus a constant involving the top bids B_i in Eq. 4, constant components of the prior on w , and the normalizer. Thus, we can optimize the smoothed revenue by taking MAP estimates of w .

Intuitively, we have defined variables corresponding to the auction host’s satisfaction. With historical data of auction attributes and bids, we imagine that the host was satisfied with every auction. When we fit w , we ask for the reserve-price-setting mechanism that leads to such an outcome.

2.3 MAP estimation with EM

The EM algorithm is a technique for maximum likelihood estimation in the face of hidden variables [9]. (When there are regularizers, it is a technique for MAP estimation.) In the E-step, we compute the posterior distribution of the hidden variables given the current model settings; in the M-step, we maximize the expected complete regularized log likelihood, where the expectation is taken with respect to the previously computed posterior.

In the OV model, the latent variables are the reserve prices \mathbf{y} ; the observations are the objective variables \mathbf{z} ; and the model parameters are the coefficients w . We compute the posterior expectation of the latent reserve prices in the E-step and fit the model parameters in the M-step. This is a coordinate ascent algorithm on the expected complete regularized log likelihood of the model and the data. Each E-step tightens the bound on the likelihood and the new bound is then optimized in the M-step.

E-step. At iteration t , the E-step computes the conditional distribution of the latent reserve prices y_i given the objective variables $z_i = 1$ and the parameters $w^{(t-1)}$ of the previous iteration. It is

$$\begin{aligned} p(y_i | z_i = 1, w^{(t-1)}) \\ \propto p(z_i = 1 | y_i) p(y_i | w^{(t-1)}) \end{aligned} \quad (9)$$

$$\propto \exp\{-(B_i - R(y_i, B_i, b_i))\} \phi\left(\frac{y_i - f(x_i; w^{(t-1)})}{\sigma}\right). \quad (10)$$

where $\phi(\cdot)$ is the pdf of the standard normal distribution.

The normalizing constant for this expression is in the appendix in Eq. 15; we compute it by integrating Eq. 9 over the real line. Once the distribution is normalized we can compute the posterior expectation $\mathbb{E}[y_i | z_i, w^{(t-1)}]$ by using the moment generating function. For the derivations and the resulting updates see Eq. 18 in Appendix A.

M-step. The M-step maximizes the complete joint log-likelihood with respect to the model parameters w . When we use a linear predictor to set the reserve prices, i.e. $f(x_i; w) = x_i^\top w$, the M-step has a closed form update, which amounts to ridge regression against response variables $\mathbb{E}[y_i | z_i, w^{(t-1)}]$ (Eq. 18) computed in the E-step. The update is

$$w^{(t)} = \left(\lambda I + \frac{1}{\sigma^2} \mathbf{x}^\top \mathbf{x}\right)^{-1} \frac{1}{\sigma^2} \mathbf{x}^\top \mathbb{E}[\mathbf{y} | \mathbf{z}, w^{(t-1)}] \quad (11)$$

where $\mathbb{E}[\mathbf{y} | \mathbf{z}, w^{(t-1)}]$ denotes the vector with i^{th} entry

$\mathbb{E}[y_i | \mathbf{z}, w^{(t-1)}]$ and \mathbf{x} is a matrix of all feature vectors x_i .

Objective variables, EM, and mitigating risk. We visually examine the interplay between the E-step and the M-step in Figure 5b. For fixed $w^* = w^{(t-1)}$, the three terms of the posterior distribution of the latent variables in Eq. 9 are drawn in the picture for four example auctions: The posterior distribution of the latent variables (red) is proportional to the likelihood term (green) times the prior (blue.) In the E-step, we need to compute the expectation of the distribution proportional to the red curve.

In the bottom right panel, the current reserve price prediction (the peak of the blue Gaussian bump) is slightly smaller than the highest bid (the peak of the green likelihood). We can see that the mean of the red curve lies somewhere in between. Hence, in the next M-step the prediction task for this particular auction consists of predicting a slightly larger reserve price than is currently predicted.


In contrast, consider the top right panel. Here, the current reserve price prediction is almost aligned with the highest bid. To the right of the highest bid, the likelihood is very small ($p(z_2 = 1 | Y_2 \geq B_2) = e^{-B_2}$.) As a consequence, almost the entire mass of the red curve proportional to the posterior lies to the left of the highest bid. Its mean (the expected reserve price) will be smaller than the current reserve price prediction. Even though the current prediction achieves almost perfect revenue, in the E-step the expected reserve price is pulled away from the highest bid. This reduces the risk of overestimating.

After the expectations are updated for all auctions – this can happen in parallel – we run another M-step which is a linear regression against the expectations we just computed.

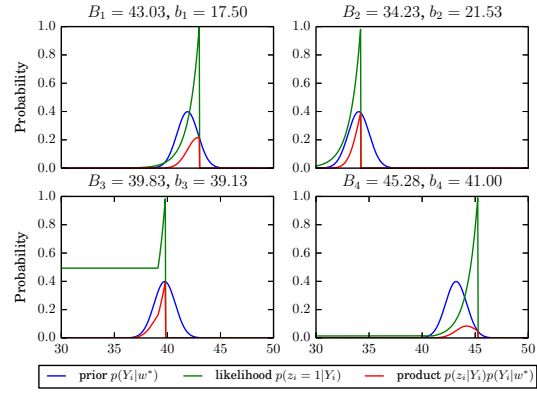
In summary, given a predictor and its predictions from the previous M-step, the E-step pulls the response variables for the next M-step either closer to the highest bid or away from it if the risk of overbidding becomes too large. The shapes of the individual priors and likelihood terms control the auction specific trade-off between pulling into one direction or into the other.

Algorithm details and computational complexity. To initialize, we set the expected reserve prices to be the highest bids $\mathbb{E}[y_i | z_i] = B_i$ and run an M-step. This deterministic initialization puts the initial predictions in a range we expect them to be, i.e., around the highest bids.

The algorithm then alternates between updating the weights using Eq. 11 (M-step) and computing posterior expectations of the latent reserve prices (E-step). It terminates when the change in



(a) The objective variable model (OV model). The objective variable is shaded with diagonal lines to distinguish that its value is not observed but rather set to our desired value.



(b) For fixed w the posterior of the latent reserve price (red) is proportional to the prior (blue) times the likelihood of the objective (green). MAP estimation uncovers profitable modes of the posterior.

Figure 5: The OV approach transforms the revenue maximization task into a MAP estimation task. The model and the hallucinated data are designed such that the modes of the model’s posterior are the local maxima of the smoothed revenue in Eq. 3

Algorithm 1 OV Regression

```

initialize  $\mathbb{E}[Y_{1:N}] = B_{1:N}$ 
initialize  $t = 1$ 
repeat
  update  $w^{(t)}$  using Eq. 11
  for  $i \in \{1, \dots, N\}$  do
    update  $\mathbb{E}[Y_i]$  using Eq. 18
  end for
  increment  $t$ 
  compute  $R_v^{(t)}$  revenue of predictor  $w^{(t)}$  on the validation set
  using Eq. 1
until  $|R_v^{(t)} - R_v^{(t-1)}| < 10^{-5}$ 
return predictor  $w^{(t)}$  and its revenue  $R$  on the test set

```

revenue on a validation set is below a threshold. (We use 10^{-5} .) Algorithm 1 contains pseudocode.

The E-step is linear in the number of auctions N . It is easily parallelized because the expected reserve prices are conditionally independent given the predictor. The M-step update has asymptotic complexity $O(d^2N)$ where d is the number of features.

2.4 Nonlinear Objective Variable Models

One of the advantages of our EM algorithm is that we can change the parameterized prediction technique $f(x_i; w)$ from which we map auction features to the mean of the reserve price. So far we have only considered linear predictors; here we show how we can adapt the algorithm to nonlinear predictors. As we will see in Section 3, these nonlinear predictors outperform the linear predictors.

In our framework, much of the model in Figure 5a and corresponding algorithm remains the same even when considering nonlinear predictors. The distribution of the objective variables is unchanged (Eq. 4) as well as the E-step update in the EM algorithm (Eq. 18). All of the changes are in the M-step.

Kernel regression. Kernel regression [1] maps the features x_i into a higher dimensional space through a feature map $\psi(\cdot)$; the mechanism for setting the reserve price becomes $f(x_i; w) = \psi(x_i)^T w$. In kernel regression we work with the $N \times N$ Gram ma-

trix K of inner products, where $K_{ij} = \psi(x_i)^T \psi(x_j)$. In this work we use a polynomial kernel of degree D , and thus compute the Gram matrix without evaluating the feature map $\psi(\cdot)$ explicitly, $K = (\mathbf{x}^T \mathbf{x} + 1)^D$.

Rather than learning the weights directly, kernel methods operate in the dual space $\alpha \in \mathbb{R}^N$. If K_i is the i^{th} column of the Gram matrix, then the mean of the reserve price is

$$f(x_i; w) = \psi(x_i)^T w = K_i^T \alpha. \quad (12)$$

The corresponding M-step in the algorithm becomes

$$\alpha^{(t)} = \left(\frac{1}{\sigma^2} K + \lambda I_N \right)^{-1} \frac{1}{\sigma^2} \mathbb{E}[\mathbf{y} | \mathbf{z}, \alpha^{(t-1)}]. \quad (13)$$

See [7] for the technical details around kernel regression.

We will demonstrate in Section 3 that replacing linear regression with kernel regression can lead to better reserve price predictions. Next, we consider neural networks as an alternative method for infusing nonlinearity into the model.

Neural networks. We also explore an objective variable model that uses a neural network [6] to set the mean reserve prices. We use a network with one hidden layer of H units and activation function $\tanh(\cdot)$. The parameters of the neural net are the weights of the first layer and the second layer: $w = \{w^{(1)} \in \mathbb{R}^{H \times d}, w^{(2)} \in \mathbb{R}^{1 \times H}\}$. The mean of the reserve price is

$$f(x_i; w) = w^{(2)} (\tanh(w^{(1)} x_i)). \quad (14)$$

The M-step is no longer analytic; Instead, the network is trained using stochastic gradient methods. See Appendix B for implementation details. We found the OV method with neural networks performed better than both linear regression and kernel regression.

3. EMPIRICAL STUDY

We study our algorithms with two simulated data sets and a large collection of real-world auction data from eBay. In each study, we fit a model on a subset of the data (using a validation set to set hyperparameters) and then test how profitable we would be if we used the fitted model to set reserve prices in a held out set. Our objective variable methods outperform the existing state of the art.

Table 1: The performance of the EM algorithms from Section 2 (OV regression, OV kernel regression with degree 2 and 4, OV neural networks) against the current state of the art (DC [13] and NoF [8]). We report results in terms of percentage of maximum possible revenue (computed by an oracle that knows the highest bid in advance). For each data set, we report mean and standard error aggregated from ten independent splits into training, validation, and test data. Our methods outperform the existing methods on all data.

	OV Reg	OV Kern (2)	OV Kern (4)	OV NN	DC [13]	NoF [8]
Linear Sim.	81.4 ± 0.2	81.2 ± 0.2	78.2 ± 0.6	72.2 ± 1.5	80.3 ± 0.3	49.9 ± 0.1
Nonlinear Sim.	50.3 ± 0.3	66.2 ± 0.4	70.1 ± 0.6	63.7 ± 2.9	59.4 ± 2.0	49.9 ± 0.2
eBay (s)	61.0 ± 0.7	63.7 ± 3.0	63.4 ± 2.8	74.4 ± 1.1	59.5 ± 1.1	55.8 ± 0.3
eBay (L)	62.4 ± 0.2	-	-	84.0 ± 0.2	-	56.0 ± 0.1

Data sets and replications. We evaluated our method on both simulated data and real-world data.

- *Linear simulated data.* Our simplest simulated data contains $d = 5$ auction features. We draw features $x_i \sim N(0, I) \in \mathbb{R}^d$ for 2,000 auctions; we draw a ground truth weight vector $\hat{w} \sim \mathcal{N}(0, I) \in \mathbb{R}^d$ and an intercept $\alpha \sim \mathcal{N}(0, 1)$; we draw the highest bids for each auction from the regression $B_i \sim \mathcal{N}(w^\top x_i + \alpha, 0.1)$ and set the second bids $b_i = B_i/2$. (Data for which B_i is negative are discarded and re-drawn.) We split into $N_{\text{train}} = 1000$, $N_{\text{valid}} = 500$, and $N_{\text{test}} = 500$.
- *Nonlinear simulated data.* These data contain features x_i , true coefficients \hat{w} , and intercept α generated as for the linear data. We generate highest bids by taking the absolute value of those generated by the regression and second highest bids by halving them, as above. Taking the absolute value introduces a nonlinear relationship between features and bids.
- *Data from eBay.*

Our real-world data is auctions of sports collectibles from eBay. It has been collected by the owners of sportscollectors.net and contains historical auction data from eBay of sport collectibles and memorabilia²

Second price auctions with reserve, as we study them in this paper are widely used for online auctions of advertisement space, but unfortunately there is no publicly available data set of advertisement auctions with features. In contrast to the auction mechanism we study here, the actual auction mechanism used on eBay is not incentive compatible and the bids are not sealed. Consequently, to keep sale prices low, experienced bidders refrain from bidding the true amount they are willing to pay until seconds before the bidding closes [16].

Additionally, eBay does not publish the highest bid which is necessary to construct the revenue function Eq. 1 and for training the objective variable models and their competitors. We used the derived data set provided here³ which has been synthetically augmented with highest and second highest bids by the authors of [13] so it can be used as a proxy for a real-world data set of sealed bid second price auctions with reserve.

The data set is focused on products which have been autographed by athletes playing baseball, basketball, football and hockey. It contains products from a wide variety of categories such as athletic gear, trading cards and ticket stubs and the athletes span from very famous players such as Michael Jordan (the basketball player), to virtually unknown players playing in local leagues. As a result there is high variability in the popularity of an auction and the prices the products actually sold for. The mean

and standard deviation of the highest and second highest bid are $\text{mean}(B) = 52.23$, $\text{std}(B) = 71.01$, $\text{mean}(b) = 29.31$, and $\text{std}(b) = 42.08$ respectively. Figure 2b contains a scatter plot of the highest bid vs. the second highest bid for this data set. There are $d = 78$ features ranging over information from eBay such as seller rating and indicators whether the listing is associated with a picture or not, to context from other sources about the product. There is for example an indicator whether the player who signed the item is in his sport’s hall of fame. All covariates are centered and rescaled to have mean zero and standard deviation one.

We analyze two data sets from eBay, one small and one large. On the small data set, the total number of auctions is 6,000, split into $N_{\text{train}} = N_{\text{valid}} = N_{\text{test}} = 2,000$. On the large data set, the total number is 70,000, split into $N_{\text{train}} = 50,000$, and $N_{\text{valid}} = N_{\text{test}} = 10,000$.

In our study, we fit each method on the training data, use the validation set to decide on hyperparameters, and then evaluate the fitted predictor on the test set, i.e., compute how much revenue we make when we use it to set reserve prices. For each data set, we replicate each study ten times, each time randomly creating the training set, test set, and validation set. The total test revenue is reported as percentage of the sum of the highest bids in the test set. These results are then averaged over the 10 replications.

Algorithms. We describe the objective variable algorithms from Section 2, all of which we implemented in Theano [5, 4], as well as the two previous methods we compare against.

- *OV Regression.* OV regression learns a linear predictor w for reserve prices using the algorithm in Section 2.3. We find a good setting for the smoothing parameter σ and regularization parameter λ using grid search.
- *OV Kernel Regression.* OV kernel regression uses a polynomial kernel to predict the mean of the reserve price; we study polynomial kernels of degree 2 and 4.
- *OV Neural Networks.* OV neural networks fits a neural net for predicting the reserve prices. As we discuss in Section 2.4, the M-step uses gradient optimization; we use stochastic gradient ascent with a constant learning rate and early stopping [15]. Further, we use a warm-start approach, where the next M-step is initialized with the results of the previous M-step. We set the number of hidden units to $H = 5$ for the simulated data and $H = 100$ for the eBay data. We use grid search to set the smoothing parameter σ , the regularization parameters, the learning rate, the batch size, and the number of passes over the data for each M-step.
- *Difference of Convex Functions (DC) [13].* The DC algorithm finds a linear predictor of reserve price with an iterative procedure based on DC-programming [17]. Grid search is used on the regularization parameter as well as the margin to select the surrogates for the auction loss.

²<http://jaygrossman.com/post/2013/06/10/Predicting-eBay-Auction-Sales-with-Machine-Learning.aspx>

³<http://cims.nyu.edu/~munoz/data/index.html>




Figure 6: Additional comparisons to Figure 3 of methods for predicting the reserve price from features: OV regression (left panel), OV kernel regression with $d = 2$ (middle panel), and OV kernel regression with $d = 4$ (right panel). Each point in the scatter plot corresponds to an auction in the test set of 2000 data points from the eBay data set. The grey scale corresponds to the fraction of the highest possible revenue each prediction generated (revenue divided by highest bid.) For each method, the total revenue generated as percentage of the highest possible revenue for this particular data split is reported in the title of the panel.




Figure 7: Auction outcomes on the small eBay test set from Figures 2b, 3 and 6 summarized in histograms of the highest bids. The histogram bins are color coded by the outcome of the corresponding auction. Bins in white color account for the highest bids of auctions where the prediction was larger than the highest bid. For these auctions, the resulting revenue is 0. The black bins represent the highest bids of auctions where the predicted reserve price was below the second bid and resulted in the second bid as revenue. The grey bins correspond to auctions where the prediction fell into the margin between the highest and second highest bid. Here, the predicted reserve price is earned as revenue.

- *No Features (NoF)* [8]. This is the state of the art approach to set the reserve prices when we do not consider the auction’s features. The algorithm iterates over the highest bids in the training set and evaluates the profitability of setting all reserve prices to this value on the training set. Ref. [13] gives a more efficient algorithm based on sorting.

Results. Table 1 gives the results of our study. The metric is the percentage of the highest possible revenue, where an oracle anticipates the bids and sets the reserve price to the highest bid. It is computed by summing the revenues generated from each prediction, dividing by the sum of the highest bids.

A trivial strategy (not reported) sets all reserve prices to zero, and earns the second highest bid on each auction. The algorithm using no features [8] does slightly better than this but not as well as the algorithms using features. OV regression [this paper] and DC [13] both fit linear mappings and exhibit similar performance. However, the DC algorithm does not scale to the large eBay data set.

The nonlinear OV algorithms (OV kernel regression and OV neural networks) outperform the linear models on the nonlinear simulated data and the real-world data.

The predictions on test auctions from the small eBay data of all feature based reserve price prediction methods we study in this paper are visualized in Figures 3 and 6.

The histograms in Figure 7 complement these scatter plots. All five panels are histograms of the highest bids. The average of the highest bids is around \$50 but most of the mass of the histogram is on auctions with smaller highest bids. The auctions corresponding to these bids are split into three cases which are color coded to distinguish the outcomes of the auctions.

The predictions of DC are very conservative and result in few overpredictions. The reserve price predictions of OV regression result in similar overall performance but their distribution is different. While DC makes barely any predictions over \$50, the predictions of OV regression are larger. OV regression generates less profit than DC in most individual auctions but compensates for this effect by doing much better in auctions with a large profit margin.

The mediocre performance of DC and OV regression on this data set indicates that there might not be a strong linear relationship between the auction features and a profitable setting of the reserve prices. The nonlinear methods outperform DC and OV regression on this data.

The two OV kernel methods produce similar predictions. Compared to the other methods, they are very bold and many of their




Figure 8: Reserve price predictions of OV neural networks on unseen test auctions from the large eBay data set. As in Figures 3 and 6 each prediction is color coded by the fraction of highest possible revenue it produces. The left panel shows a scatter plot of the second bid against the highest bid. The panel in the middle shows a scatter plot of the predicted reserve price against the highest bid and the panel on the right shows a scatter plot of the final revenue against the highest bid. In the central panel, the points below the diagonal correspond to auctions where the predicted reserve price is larger than the highest bid. They result in 0 revenue and are mapped to the y-axis in the revenue plot in the right panel. The predictions above the diagonal in the central panel are either smaller than the second bid and get pulled closer to the diagonal in their visualization in the right panel or they lie in the margin between the highest and second highest bid. In this case, their location is preserved in the center and the right panel meaning that the auction host got the exactly their reserve price prediction as revenue.

predictions are large. This leads to high revenue in some auctions and 0 profit in many others. Yet, the total revenue generated is greater than for the linear predictors.

OV neural networks produces the best results. It places many predictions into the profitable margin between second highest bid and highest bid while keeping the number of overpredictions lower than the kernel methods. Hence, it is a viable solution to maximizing profit in real-world second price auctions with reserve.

It has been observed in many application domains that neural networks thrive under the presence of more training data. We see this in our study as well; the results of OV neural networks on the eBay data sets improve when the model is trained with more data.

Figure 8 explores the reserve price predictions of OV neural networks on unseen test data from the large eBay data set. When we compare the central panel (a scatter plot of the predicted reserve price against the highest bid) to the predictions in the right panel of Figure 3 we can see that the predictor trained on the larger data set is making larger predictions for auctions with large potential revenue (highest bid above 300.) By using more training data the model is able to uncover additional structure in the relationship between the eBay auction features and the potential for profit as encoded by the objective variables.

4. SUMMARY AND DISCUSSION

We develop the objective variable approach for combining probabilistic modeling with optimal decision making. We use this method to solve the problem of how to set the reserve price in second-price auctions with reserve. This auction paradigm is widely used to sell online advertisement space. Our algorithms scales better and outperforms the current state of the art on both simulated and real-world data.

Acknowledgments

This work is supported by IIS-1247664, ONR N00014-11-1-0651, DARPA FA8750-14-2-0009, Facebook, Adobe, Amazon, and the John Templeton Foundation.

5. REFERENCES

- [1] A. Aizerman, E. Braverman, and L. Rozoner. Theoretical foundations of the potential function method in pattern recognition learning. *Automation and remote control*, 25, 1964.
- [2] J. Albert and S. Chib. Bayesian analysis of binary and polychotomous response data. *Journal of the American statistical Association*, 88(422):669–679, 1993.
- [3] Z. Bar-Yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *ACM-SIAM symposium on Discrete algorithms*, pages 964–970, 2002.
- [4] F. Bastien et al. Theano: new features and speed improvements. Deep Learning and Unsupervised Feature Learning NIPS 2012 Workshop, 2012.
- [5] J. Bergstra et al. Theano: a CPU and GPU math expression compiler. In *Proceedings of the Python for Scientific Computing Conference (SciPy)*, June 2010.
- [6] C. Bishop et al. Neural networks for pattern recognition. 1995.
- [7] C. Bishop et al. *Pattern recognition and machine learning*, volume 4. springer New York, 2006.
- [8] N. Cesa-Bianchi, C. Gentile, and Y. Mansour. Regret minimization for reserve prices in second-price auctions. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 1190–1204, 2013.
- [9] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39:1–38, 1977.
- [10] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, New York, NY, USA, 2010.
- [11] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *International conference on artificial intelligence and statistics*, pages 249–256, 2010.

- [12] C. Holmes, L. Held, et al. Bayesian auxiliary variable models for binary and multinomial regression. *Bayesian Analysis*, 1:145–168, 2006.
- [13] A. Medina and M. Mohri. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *International Conference on Machine Learning*, 2014.
- [14] M. Ostrovsky and M. Schwarz. Reserve prices in internet advertising auctions: A field experiment. In *ACM conference on Electronic commerce*, pages 59–60, 2011.
- [15] L. Prechelt. Early stopping - but when? In *Neural Networks: Tricks of the Trade*, pages 53–67. Springer, 2012.
- [16] A. E. Roth and A. Ockenfels. Last minute bidding and the rules for ending second price auctions: evidence from ebay and amazon auctions on the internet. *American economic review*, 92(4):1093–1103, 2002.
- [17] P. Tao and L. An. A DC optimization algorithm for solving the trust-region subproblem. *SIAM Journal on Optimization*, 8(2):476–505, 1998.
- [18] M. Toussaint, L. Charlin, and P. Poupart. Hierarchical pomdp controller optimization by likelihood maximization. In *UAI*, volume 24, pages 562–570, 2008.
- [19] M. Toussaint, S. Harmeling, and A. Storkey. Probabilistic inference for solving (po) mdps. 2006.

APPENDIX

A. UPDATE EQUATIONS FOR EM

The normalizing constant C_i of Eq. 9 can be computed by integrating Eq. 9 over the real line. Let $\mu_i = f(x_i; w^{(t-1)})$. Up to a constant factor of e^{B_i} the normalizing constant C_i then equals

$$\begin{aligned}
 C_i e^{B_i} &= \int_{-\infty}^{b_i} e^{b_i} \phi\left(\frac{y_i - \mu_i}{\sigma}\right) dy_i + \int_{b_i}^{B_i} e^{y_i} \phi\left(\frac{y_i - \mu_i}{\sigma}\right) \\
 &\quad + \int_{B_i}^{\infty} \phi\left(\frac{y_i - \mu_i}{\sigma}\right) \\
 &= \sigma e^{b_i} \Phi\left(\frac{b_i - \mu_i}{\sigma}\right) + \sigma \left[1 - \Phi\left(\frac{B_i - \mu_i}{\sigma}\right)\right] \\
 &\quad + \sigma e^{\mu_i + \frac{\sigma^2}{2}} \left[\Phi\left(\frac{B_i - (\mu_i + \sigma^2)}{\sigma}\right) - \Phi\left(\frac{b_i - (\mu_i + \sigma^2)}{\sigma}\right)\right].
 \end{aligned} \tag{15}$$

Computing the expectation of the latent reserve price $\mathbb{E}[y_i]$ entails evaluating the moment generating function $M_i(s) = \mathbb{E}[e^{sy_i}]$, where expectation is taken w.r.t. the posterior $p(y_i|z_i = 1, w^{(t-1)})$. Taking the derivative with respect to s and setting $s = 0$ then yields the desired expectation.

$$\begin{aligned}
 \mathbb{E}[y_i] &= \left. \frac{dM_i(s)}{ds} \right|_{s=0} \\
 &= \frac{\sigma e^{b_i}}{C_i} \mu_i \Phi\left(\frac{b_i - \mu_i}{\sigma}\right) - \frac{\sigma^2 e^{b_i}}{C_i} \phi\left(\frac{b_i - \mu_i}{\sigma}\right) \\
 &\quad + \frac{\sigma}{C_i} (\mu_i + \sigma^2) e^{\mu_i + \frac{\sigma^2}{2}} \Phi\left(\frac{B_i - (\mu_i + \sigma^2)}{\sigma}\right) \\
 &\quad - \frac{\sigma}{C_i} (\mu_i + \sigma^2) e^{\mu_i + \frac{\sigma^2}{2}} \Phi\left(\frac{b_i - (\mu_i + \sigma^2)}{\sigma}\right) \\
 &\quad - \frac{\sigma^2}{C_i} e^{\mu_i + \frac{\sigma^2}{2}} \left[\phi\left(\frac{B_i - (\mu_i + \sigma^2)}{\sigma}\right) - \phi\left(\frac{b_i - (\mu_i + \sigma^2)}{\sigma}\right)\right] \\
 &\quad + \frac{\sigma}{C_i} \mu_i \left[1 - \Phi\left(\frac{B_i - \mu_i}{\sigma}\right)\right] + \frac{\sigma^2}{C_i} \phi\left(\frac{-\mu_i}{\sigma}\right)
 \end{aligned} \tag{17}$$

B. IMPLEMENTATION OF OV NEURAL NETWORKS

Here we give the implementation details for training OV neural networks. The E-step is identical to OV regression and is given in Eq. 18.

In the M-step, we train a neural network with H hidden layers and activation function $\tanh(\cdot)$ to predict the expectations computed in the previous E-step from auction features. This is done using stochastic gradient descent.

Initialization. The weights of the input layer are initialized by drawing each entry from

$$\text{Uniform}\left[-\frac{\sqrt{6}}{\sqrt{H+d}}, \frac{\sqrt{6}}{\sqrt{H+d}}\right] \tag{19}$$

where d is the number of features, while the weights of the hidden layer are drawn element wise from

$$\text{Uniform}\left[-\frac{\sqrt{6}}{\sqrt{H+1}}, \frac{\sqrt{6}}{\sqrt{H+1}}\right]. \tag{20}$$

See [11] for mathematical considerations of this initialization scheme.

Learning rate. We use a constant learning rate. It is determined using grid search over the values $[10^{-2}, 10^{-3}, 10^{-4}]$.

Regularization. The weights of the network are both L1 and L2 regularized. Both regularization parameters are given the same value (to reduce the complexity of the grid search). We determine a value in $[10^{-3}, 10^{-4}, 10^{-5}]$ using cross validation on a validation set.

Early Stopping. To avoid overfitting to the training set we use early stopping. Every 50 epochs the prediction accuracy is evaluated on the validation set. When the validation accuracy stops improving, the algorithm stops.