Modeling and Decision Optimization in Real-time Bidding Display Advertising
实时竞价广告中的建模与决策优化

Kan Ren

Apex Data and Knowledge Mangement Lab
Shanghai Jiao Tong University

Aug. 2018

Outline

1. Background
   - Online Advertising
   - Real-time Bidding
   - Research Topics
2. Research Problems
   - User Response Prediction
   - Bidding Strategy Optimization
   - Reinforcement Learning for Advertising
   - Conversion Attribution
   - Bid Landscape Forecasting
3. Related Literatures

Outline

Online Service and Marketing

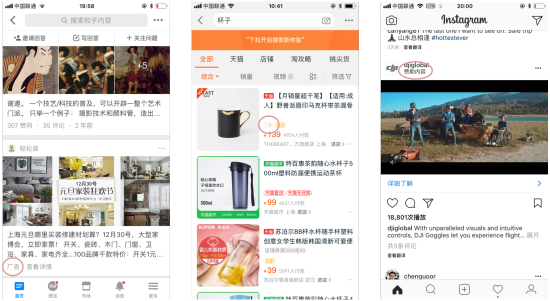

Online service has applied everywhere in our life

- Recommendation: Douban Music, Taobao product, etc.
- Aggregation: News feed, Search Engine, etc.
- Community: QA Websites, Social Media, etc.

Advertising has become the major income source for online services.

Online Advertising

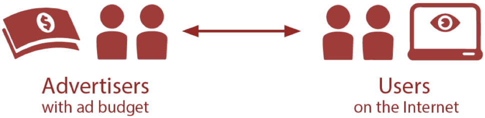

Online Advertising as a Service

Bridge the gap between the user and the product seller in a more flexible, effective and accurate paradigm.

Online Advertising

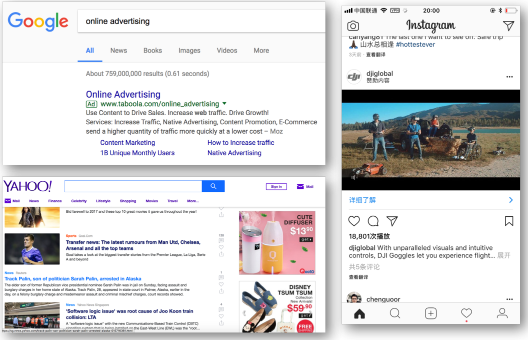

Types of Online Advertising:

- Search engine advertising;
- Display advertising;
- Mobile advertising, etc.

Nowadays, performance-based ads has drawn huge attentions.

Goal of Computer

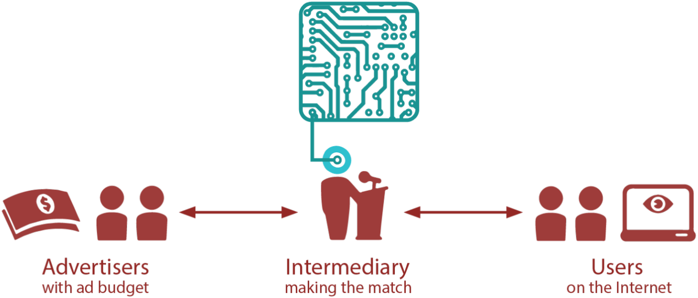

Goal of Computer

Address the **right user** with the **right message** in the **right context** and at the **right prices**.

Display Advertising

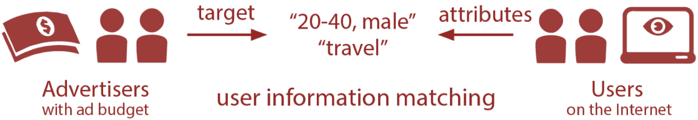

Example

User Profiling: model the attributes of different users.

User Targeting: buy a bundle of user volume with targeted attributes.

Outline

Real-time Bidding (RTB)

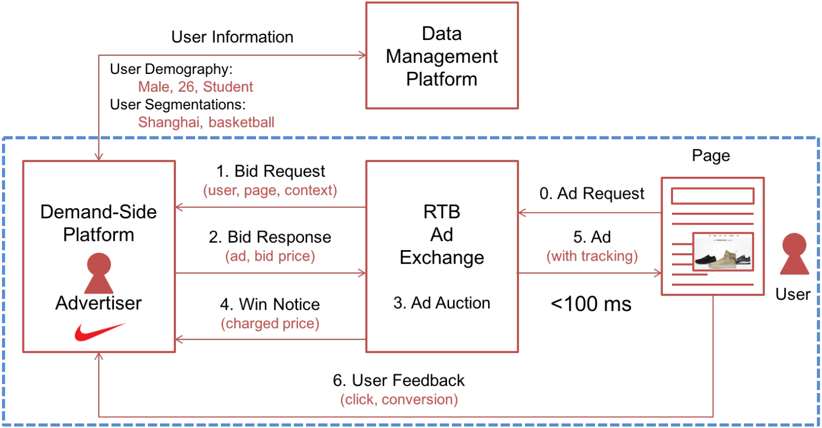

We are mainly focusing on the demand side (advertiser side).

Second Price Auction in RTB

Win and pay the **second** highest price, lose otherwise do nothing.

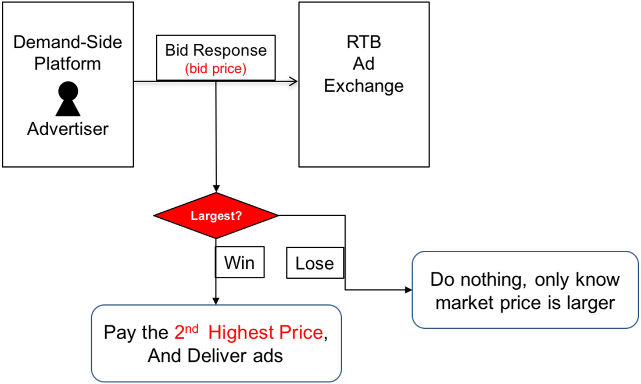

Outline

Components of Bidding Strategy

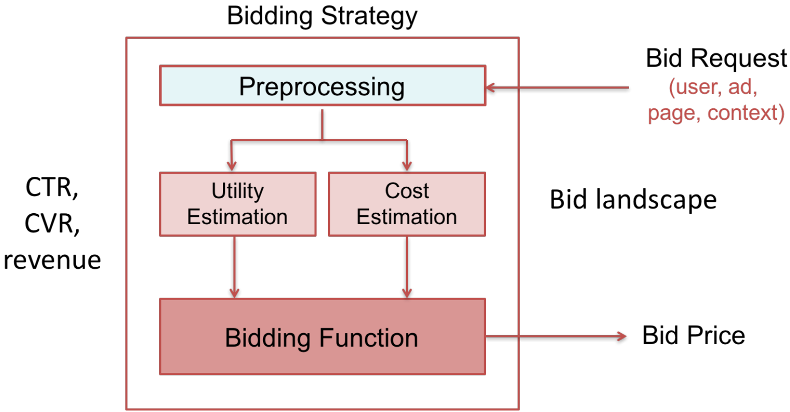

CTR: Click Through Rate.
CVR: Conversion Rate.
Bid Landscape: The probability density function of the market price.

Research Problems

Utility: User Response Prediction

Model the behavior patterns of the user and predict the user response on the given ad impression.

Decision: the Bidding Function

Sequential decision making (bidding in the RTB auction) with the consideration of total budgets.

Cost: Bid Landscape Forecasting

Estimate the cost (market price) for the given ad request, and predict the winning probability of the given bid price.

Outline

Utility Estimation: User Response Prediction

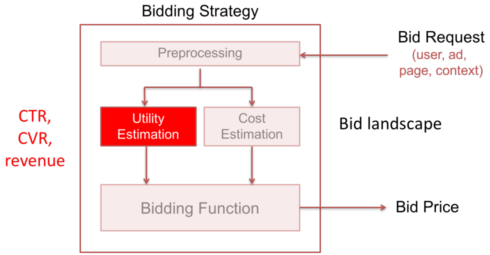

Problem Definition

Given feature $\boldsymbol{x}$ of the user and the ad, predict the probability of user taking action (click or conversion) on the proposed ads $\Pr(y = 1|\boldsymbol{x})$.

Data Challenges

Categorical Data: { Location=Shanghai, Gender=Male ... }

Sparse Input: $\boldsymbol{x} = [0, 0, 1, 0, \ldots, 0, 1, 0, \ldots]$

Related Work of User Response Prediction

Regression Model

Logistic Regression (LR)

K.-c Lee et al. Estimating Conversion Rate in Display Advertising from Past Performance Data. KDD 2012

Tree-based Model

X. He et al. Practical Lessons from Predicting Clicks on Ads at Facebook. ADKDD 2014

Factorization Machines

A.K. Menon et al. Response prediction using collaborative filtering with hierarchies and side-information. KDD 2011

Related Work

Other Variants

Bayesian Probit Regression

Web-scale Bayesian Click-through Rate Prediction for Sponsored Search Advertising in Microsoft's Bing Search Engine, T. Graepel et al. ICML 2010

Factorization Machine with FTRL

A.-P. Ta. Factorization Machines with Follow-The-Regularized-Leader for CTR prediction in Display Advertising. Big Data 2015

Deep Neural Networks

Q. Liu et al. A Convolutional Click Prediction Model. CIKM 2015.
W. Zhang et al. Deep Learning over Multi-field Categorical Data – A Case Study on User Response Prediction. ECIR 2016.
Y. Qu, H. Cai, **K. Ren**. Product-based Neural Networks for User Response Prediction. ICDM 2016.
G. Zhou et al. Deep Interest Network for Click-Through Rate Prediction. KDD 2018

Related Work (cont.)

Objective Function

Squared Error

$$\mathcal{L}_{SE} = \frac{1}{2}(y - \hat{y})^2, y \in \{0, 1\}, \hat{y} \in [0, 1]$$

Cross Entropy

$$\mathcal{L}_{CE} = -y \log \hat{y} - (1 - y) \log(1 - \hat{y}), y \in \{0, 1\}, \hat{y} \in [0, 1]$$

Related Work (cont.)

Evaluation Measurements

- Area under ROC Curve (AUC)
- Relative Information Gain (Cross Entropy)

Related Work (cont.)

Traditional Bidding Function

Truthful Bidding Function

$$b(\boldsymbol{x}) = V_{action} \cdot f(\boldsymbol{x}),$$

where f is the utility estimation function, such as pCTR.

Linear Bidding Function

$$b(\boldsymbol{x}) = \phi \cdot V_{action} \cdot f(\boldsymbol{x}) = b_0 \cdot f(\boldsymbol{x}).$$

C. Perlich et al. Bid optimizing and inventory scoring in targeted online advertising. KDD 2012.

Problem Setup

Motivation

To make the user response prediction more precisely *accurate* considering the context.

Problem Setup

Motivation

To make the user response prediction more precisely *accurate* considering the context.

Rethinking about CTR estimation

- Why do we regard the CTR estimation as a classification task?
- What is the optimization objective for the advertiser? The accuracy of pCTR?
- Why is the bid price linear or positive correlated to the pCTR?

Problems

Isolation Optimization and Joint Optimization

The related works only take classification error as the loss.
They do not consider the subsequent usage of the prediction model.

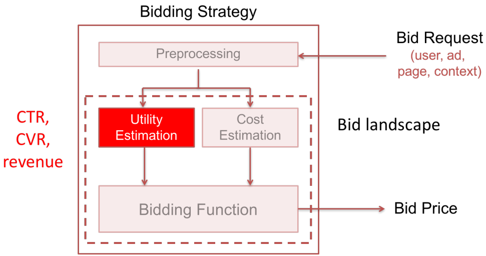

Our Solution

Embed the user response prediction model into the whole procedure of the bidding. And take the overall profit as our learning objective function, to maximize the gains of the advertiser.

K. Ren et al. User Response Learning for Directly Optimizing Campaign Performance in Display Advertising. CIKM, 2016.

Notations and descriptions

Notation	Description
y	The true label of user response.
$\boldsymbol{x}$	The bid request represented by its features.
$\boldsymbol{\theta}$	The parameter of CTR estimation function.
$f_{\boldsymbol{\theta}}(\boldsymbol{x})$	the CTR estimation function to learn.
$b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))$	The bid price determined by the estimated CTR, b for short.
$R_{\boldsymbol{\theta}}(\cdot)$	The utility function.

Market Modeling

Market Price

The second highest price proposed during an RTB auction (2^{nd} price).

Market Price Distribution (p.d.f.)

$$p_z(z), z \in \mathcal{N}.$$

Market Modeling

Market Price

The second highest price proposed during an RTB auction (2^{nd} price).

Market Price Distribution (p.d.f.)

$$p_z(z), z \in \mathcal{N}.$$

Winning Probability when Bidding at price b (c.d.f.)

$$w(b) = \int_0^b p_z(z) dz. \tag{1}$$

Market Modeling

Market Price

The second highest price proposed during an RTB auction (2^{nd} price).

Market Price Distribution (p.d.f.)

$$p_z(z), z \in \mathcal{N}.$$

Winning Probability when Bidding at price b (c.d.f.)

$$w(b) = \int_0^b p_z(z)dz. \quad (1)$$

Expected Cost under 2^{nd} Price Auction (if winning)

$$c(b) = \frac{\int_0^b z p_z(z)dz}{w(b)} = \frac{\int_0^b z p_z(z)dz}{\int_0^b p_z(z)dz}. \quad (2)$$

Objective Function

Objective Function

$$\theta^* = \arg\max_{\theta} \int_{\boldsymbol{x}} R_{\theta}(\boldsymbol{x}, y; b, v, c, w) p_{\boldsymbol{x}}(\boldsymbol{x}) d\boldsymbol{x}. \tag{3}$$

- $R_{\theta}(\cdot)$ is the *Utility Function.*
- Constant click value v limits the max bid.
- We will propose two variants of $R_{\theta}(\cdot)$.

Expected Utility Model

Expected Utility (EU)

$$R_{\boldsymbol{\theta}}^{\mathrm{EU}}(\boldsymbol{x}, y) = [vy - c(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))] \cdot w(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))). \tag{4}$$

Profit = [gain - (expected cost)] × winning probability

gain = click value × click indicator .

Objective of EU

The overall expected direct profit of all the auctions can be calculated by replacing the winning probability funcion $w(b(\cdot))$ and the expected cost function $c(b(\cdot))$ into EU objective function as

$$\sum_{(\boldsymbol{x},y)\in D} R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x},y) = \sum_{(\boldsymbol{x},y)\in D} [vy - c(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))] \cdot w(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))$$
$$= \sum_{(\boldsymbol{x},y)\in D} \left[vy - \frac{\int_0^{b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))} z \cdot p_z(z)dz}{\int_0^{b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))} p_z(z)dz} \right] \cdot \int_0^{b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))} p_z(z)dz$$
$$= \sum_{(\boldsymbol{x},y)\in D} \int_0^{b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))} (vy - z) \cdot p_z(z)dz. \tag{5}$$

Optimal Parameter

Taking Eq. (5) into Eq. (3) with a regularization term turns our learning problem into convex optimization:

$$\begin{aligned}\boldsymbol{\theta}^{\text{EU}} &= \arg\min_{\boldsymbol{\theta}} \; -\sum_{(\boldsymbol{x},y)\in D} R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x},y) + \frac{\lambda}{2}\|\boldsymbol{\theta}\|_2^2 \\ &= \arg\min_{\boldsymbol{\theta}} \; \sum_{\boldsymbol{x}} \int_0^{b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))} (z - vy)\cdot p_z(z)dz + \frac{\lambda}{2}\boldsymbol{\theta}^T\boldsymbol{\theta}.\end{aligned} \tag{6}$$

where the optimal value of $\boldsymbol{\theta}$ is obtained by taking a gradient descent algorithm.

Gradient of EU

The gradient of $R_{\boldsymbol{\theta}}^{\text{EU}}(\boldsymbol{x}, y)$ with regard to $\boldsymbol{\theta}$ is calculated as

$$\frac{\partial R_{\boldsymbol{\theta}}^{\text{EU}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = \overbrace{\big(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})) - vy\big)}^{\text{bid error}} \cdot \overbrace{p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))}^{\text{market sensitivity}} \cdot \frac{\partial b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))}{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})} \frac{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})}{\partial \boldsymbol{\theta}} + \lambda \boldsymbol{\theta}. \tag{7}$$

and we update for each data instance as $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \eta \frac{\partial R_{\boldsymbol{\theta}}^{\text{EU}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}}$ by above chain rule. (SGD)

Illustration of EU Update

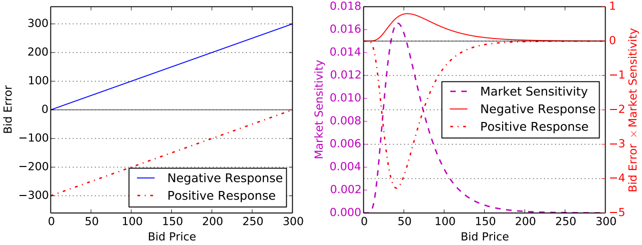

Figure: The illustration of the impact from the bid and market price of Expected Utility (EU); click value $v = 300$.

Gradient of EU

$$\frac{\partial R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = \overbrace{\left(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})) - vy\right)}^{\text{bid error}} \cdot \overbrace{p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))}^{\text{market sensitivity}} \cdot \frac{\partial b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))}{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})} \frac{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})}{\partial \boldsymbol{\theta}} + \lambda \boldsymbol{\theta}.$$

Question

Gradient of EU

$$\frac{\partial R^{\mathrm{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = \overbrace{(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})) - vy)}^{\text{bid error}} \cdot \overbrace{p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))}^{\text{market sensitivity}} \cdot \frac{\partial b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))}{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})} \frac{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})}{\partial \boldsymbol{\theta}} + \lambda \boldsymbol{\theta}.$$

Question

Gradient of EU

$$\frac{\partial R^{\mathrm{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = \overbrace{(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})) - vy)}^{\text{bid error}} \cdot \overbrace{p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))}^{\text{market sensitivity}} \cdot \frac{\partial b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))}{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})} \frac{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})}{\partial \boldsymbol{\theta}} + \lambda \boldsymbol{\theta}.$$

Gradient of RR

$$\frac{\partial R^{\mathrm{RR}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = \overbrace{\left(-\frac{vy}{b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))} + \frac{v(1-y)}{v - b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))} \right)}^{\text{bid error}} \cdot \overbrace{p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))}^{\text{market sensitivity}} \cdot \frac{\partial b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))}{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})} \frac{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})}{\partial \boldsymbol{\theta}} + \lambda \boldsymbol{\theta}.$$

Question

Gradient of EU

$$\frac{\partial R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = \overbrace{(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})) - vy)}^{\text{bid error}} \cdot \overbrace{p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))}^{\text{market sensitivity}} \cdot \frac{\partial b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))}{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})} \frac{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})}{\partial \boldsymbol{\theta}} + \lambda \boldsymbol{\theta}.$$

Gradient of RR

$$\frac{\partial R^{\text{RR}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = \Big(\overbrace{-\frac{vy}{b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))} + \frac{v(1-y)}{v - b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))}}^{\text{bid error}} \Big) \cdot \overbrace{p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))}^{\text{market sensitivity}} \cdot \frac{\partial b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))}{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})} \frac{\partial f_{\boldsymbol{\theta}}(\boldsymbol{x})}{\partial \boldsymbol{\theta}} + \lambda \boldsymbol{\theta}.$$

$f_{\boldsymbol{\theta}}(\boldsymbol{x})$? $b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))$? $p_z(z)$?

Realization of Model

Response Prediction Model $f_{\boldsymbol{\theta}}(\boldsymbol{x})$

$$f_{\boldsymbol{\theta}}(\boldsymbol{x}) \equiv \sigma(\boldsymbol{\theta}^T \boldsymbol{x}) = \frac{1}{1 + e^{-\boldsymbol{\theta}^T \boldsymbol{x}}} . \tag{8}$$

Linear Bidding Strategy

$$b(f_{\boldsymbol{\theta}}(\boldsymbol{x})) \equiv \phi \cdot v \cdot f_{\boldsymbol{\theta}}(\boldsymbol{x}), \tag{9}$$

where ϕ is the scaling parameter.

Linear Gradient

Linear Gradient of EU

$$\frac{\partial R^{\mathrm{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = \phi v^2 (\sigma(\boldsymbol{\theta}^T \boldsymbol{x}) - y) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \sigma(\boldsymbol{\theta}^T \boldsymbol{x})(1 - \sigma(\boldsymbol{\theta}^T \boldsymbol{x}))\boldsymbol{x} + \lambda \boldsymbol{\theta}. \tag{10}$$

Linear Gradient of RR

$$\frac{\partial R^{\mathrm{RR}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = \phi v \Big(- \frac{y}{\sigma(\boldsymbol{\theta}^T \boldsymbol{x})} + \frac{1 - y}{1 - \sigma(\boldsymbol{\theta}^T \boldsymbol{x})} \Big) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \sigma(\boldsymbol{\theta}^T \boldsymbol{x})(1 - \sigma(\boldsymbol{\theta}^T \boldsymbol{x}))\boldsymbol{x} + \lambda \boldsymbol{\theta}. \tag{11}$$

Recall of Traditional Logistic Regression

Squared Error LR (SE)

$$\mathcal{L}^{\text{SE}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y) = \frac{1}{2}(y - \sigma(\boldsymbol{\theta}^T \boldsymbol{x}))^2,$$
$$\frac{\partial \mathcal{L}^{\text{SE}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = (\sigma(\boldsymbol{\theta}^T \boldsymbol{x}) - y)\sigma(\boldsymbol{\theta}^T \boldsymbol{x})(1 - \sigma(\boldsymbol{\theta}^T \boldsymbol{x}))\boldsymbol{x}. \tag{12}$$

Cross Entropy LR (CE)

$$\mathcal{L}^{\text{CE}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y) = -y \log \sigma(\boldsymbol{\theta}^T \boldsymbol{x}) - (1 - y)\log(1 - \sigma(\boldsymbol{\theta}^T \boldsymbol{x})),$$
$$\frac{\partial \mathcal{L}^{\text{CE}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = (\sigma(\boldsymbol{\theta}^T \boldsymbol{x}) - y)\boldsymbol{x}. \tag{13}$$

Discussion 1: Truthful Bidding Simplification

Simplification: Truthful Bidding, $\phi = 1$

$$b(f_{\boldsymbol{\theta}}(\boldsymbol{x})) = v \cdot f_{\boldsymbol{\theta}}(\boldsymbol{x}). \tag{14}$$

Corresponding Gradient of EU & RR

$$\frac{\partial R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = v^2(\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \sigma(\boldsymbol{\theta}^T\boldsymbol{x})(1 - \sigma(\boldsymbol{\theta}^T\boldsymbol{x}))\boldsymbol{x} + \lambda\boldsymbol{\theta}, \tag{15}$$

$$\frac{\partial R^{\text{RR}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = v(\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y)p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))\boldsymbol{x} + \lambda\boldsymbol{\theta}, \tag{16}$$

Adopting Truthful Bidding function, EU & RR have one more component (market sensitivity) than SE & CE, respectively!

Discussion 2: Uniform Market Price Distribution Simplification

Simplification: Uniform Market Price Distribution

$$p_z(z) = l. \tag{17}$$

Corresponding Gradient of EU & RR

$$\frac{\partial R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = v^2 l(\sigma(\boldsymbol{\theta}^T \boldsymbol{x}) - y) \cdot \sigma(\boldsymbol{\theta}^T \boldsymbol{x})(1 - \sigma(\boldsymbol{\theta}^T \boldsymbol{x}))\boldsymbol{x} + \lambda \boldsymbol{\theta}, \tag{18}$$

$$\frac{\partial R^{\text{RR}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = vl(\sigma(\boldsymbol{\theta}^T \boldsymbol{x}) - y)\boldsymbol{x} + \lambda \boldsymbol{\theta}. \tag{19}$$

Adopting Truthful Bidding and uniform market price distribution, EU & RR have totally degenerated to SE & CE!

Summary of the Discussion

Table: The comparison of the model gradients (without regularization). LR: logistic regression, TB: truthful bidding, LB: linear bidding, UM: uniform market price distribution. LR and LR+TB+UM are equivalent (LR+TB reduces to the baseline LR when assuming the uniform market price distribution).

Model Setting	EU (SE) Gradient	RR (CE) Gradient
LR (baseline)	$\frac{\partial \mathcal{L}^{\text{SE}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = (\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y) \cdot \sigma(\boldsymbol{\theta}^T\boldsymbol{x})(1-\sigma(\boldsymbol{\theta}^T\boldsymbol{x}))\boldsymbol{x}$	$\frac{\partial \mathcal{L}^{\text{CE}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = (\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y)\boldsymbol{x}$
LR+TB	$-\frac{\partial R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = v^2(\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \sigma(\boldsymbol{\theta}^T\boldsymbol{x})(1-\sigma(\boldsymbol{\theta}^T\boldsymbol{x}))\boldsymbol{x}$	$-\frac{\partial R^{\text{RR}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = v(\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \boldsymbol{x}$
LR+TB+UM	$-\frac{\partial R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = v^2 l(\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y) \cdot \sigma(\boldsymbol{\theta}^T\boldsymbol{x})(1-\sigma(\boldsymbol{\theta}^T\boldsymbol{x}))\boldsymbol{x}$	$-\frac{\partial R^{\text{RR}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = vl(\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y)\boldsymbol{x}$
LR+LB	$-\frac{\partial R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = \phi v^2(\phi\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \sigma(\boldsymbol{\theta}^T\boldsymbol{x})(1-\sigma(\boldsymbol{\theta}^T\boldsymbol{x}))\boldsymbol{x}$	$-\frac{\partial R^{\text{RR}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = \phi v\Big(-\frac{y}{\phi\sigma(\boldsymbol{\theta}^T\boldsymbol{x})} + \frac{1-y}{1-\phi\sigma(\boldsymbol{\theta}^T\boldsymbol{x})}\Big) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \sigma(\boldsymbol{\theta}^T\boldsymbol{x})(1-\sigma(\boldsymbol{\theta}^T\boldsymbol{x}))\boldsymbol{x}$

Evaluation Flow

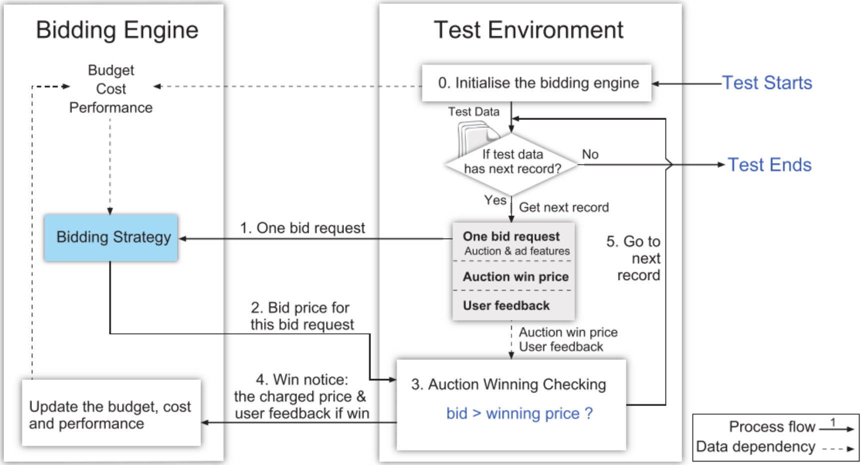

Evaluation Measures

- AUC
- RMSE

Evaluation Measures

AUC

RMSE

profit $=$ gain - cost $= V_{click} \cdot$ click $\#$ - $\sum cost$

ROI $=$ profit / cost

CTR $=$ click $\#$ / impression $\#$

eCPC $=$ cost / click $\#$

CPM $=$ cost / impression $\#$

Dataset

iPinYou

64.75M bids, 19.5 imps, 14.79 clicks and 16K expense on 9 camps over 10 days.

YOYI

443M imps, 362K clicks and 210K CNY expense over 8 days.

Compared Settings

User response prediction (truthful bidding function $b(\boldsymbol{x}) = v \cdot f(\boldsymbol{x})$)

- CE - Cross entropy loss logistic regression
- SE - Squared loss logistic regression
- EU - Expected utility model
- RR - Risk return model

Accuracy of CTR Estimation

Table: Regression performances over campaigns. AUC: the higher, the better. RMSE: the smaller, the better.

	AUC				RMSE ($\times 10^{-2}$)			
iPinYou	SE	CE	EU	RR	SE	CE	EU	RR
1458	.948	.987	.987	.977	3.01	1.94	2.42	2.32
2259	.542	.692	.674	.691	2.01	1.77	1.76	1.79
2261	.490	.569	.622	.619	1.84	1.68	1.71	1.68
2821	.511	.620	.608	.639	2.56	2.43	2.39	2.46
2997	.543	.610	.606	.608	5.98	5.82	5.84	5.82
3358	.863	.974	.970	.980	3.07	2.47	3.32	2.67
3386	.593	.768	.761	.778	2.95	2.84	3.32	2.85
3427	.634	.976	.976	.960	2.78	2.20	2.61	2.34
3476	.575	.957	.954	.950	2.50	2.32	2.39	2.33
Average	.633	.794	.795	**.800**	2.97	**2.61**	2.86	2.69
YOYI	.882	.891	**.912**	**.912**	11.9	11.7	11.8	**11.6**

Campaign Profit Evaluation (baselines)

Table: Direct campaign profit over baselines.

	profit($\times 10^7$)		ROI	
iPinYou	SE	CE	SE	CE
1458	3.2	3.6	4.2	6.6
2259	-0.32	0.40	-0.080	0.18
2261	0.29	0.63	0.26	0.40
2821	0.11	0.08	0.21	0.023
2997	0.11	0.14	0.42	0.71
3358	1.76	2.4	5.4	5.2
3386	0.51	1.6	0.16	1.2
3427	0.33	2.9	0.11	3.4
3476	0.65	3.1	0.36	3.5
Average	0.74	**1.7**	1.2	**2.3**
YOYI	665.6	**669.5**	1.8	**1.9**

Campaign Profit Evaluation

Table: Campaign profit improvement over baseline CE.

	Profit gain		ROI gain	
iPinYou	EU	RR	EU	RR
1458	7.10%	9.00%	233%	267%
2259	81.6%	99.3%	233%	472%
2261	26.3%	31.1%	44.4%	91.2%
2821	573%	615%	1334%	943%
2997	5.00%	0.700%	-3.60%	-11.4%
3358	1.70%	6.70%	77.1%	77.7%
3386	-1.20%	2.50%	20.6%	58.3%
3427	5.50%	8.70%	52.0%	175%
3476	4.20%	8.60%	16.0%	91.1%
YOYI	9.04%	0.600%	14.8%	2.11%
Average	+71.2%	+78.2%	+202%	+217%

Overall Statistics

	CTR ($\times 10^{-4}$)				eCPC			
iPinYou	SE	CE	EU	RR	SE	CE	EU	RR
1458	34	33	59	190	17	11	4.3	3.4
2259	3.3	3.6	3.7	5.8	303	235	172	136
2261	2.4	2.7	3.0	2.8	234	212	188	168
2821	5.5	5.9	4.8	7.0	116	137	105	112
2997	31	25	26	27	9.8	8.2	8.3	8.6
3358	51	41	69	61	18	19	12	12
3386	7.8	11	13	15	90	48	43	36
3427	7.2	25	29	72.8	98	25	17.3	10
3476	6.4	16	17	33.1	111	34	30	20
Average	16	18	25	**46**	110	81	64	**57**
YOYI	16	18	**26**	24	12.9	12.4	**11.3**	12

	CPM				Win Rate			
iPinYou	SE	CE	EU	RR	SE	CE	EU	RR
1458	57	37	25	65	0.22	0.24	0.13	.041
2259	100	84	64	78	0.89	0.63	0.44	0.24
2261	57	56	56	46	0.55	0.81	0.71	0.67
2821	63	80	50	78	0.12	0.63	0.48	0.45
2997	30	20	21	22	0.55	0.63	0.65	0.63
3358	92	77	80	70	0.11	0.20	0.11	0.13
3386	71	54	55	55	0.82	0.45	0.36	0.29
3427	70	60	49	75	0.75	0.26	0.22	.082
3476	71	55	50	65	0.49	0.31	0.31	0.15
Average	68	58	**50**	62	**0.50**	0.46	0.38	0.30
YOYI	**20**	23	29	30	**0.36**	0.30	0.22	0.22

Linear Gradient

Linear Gradient of EU

$$\frac{\partial R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = \phi v^2 (\sigma(\boldsymbol{\theta}^T \boldsymbol{x}) - y) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \sigma(\boldsymbol{\theta}^T \boldsymbol{x})(1 - \sigma(\boldsymbol{\theta}^T \boldsymbol{x}))\boldsymbol{x} + \lambda \boldsymbol{\theta}. \tag{20}$$

Linear Gradient of RR

$$\frac{\partial R^{\text{RR}}_{\boldsymbol{\theta}}(\boldsymbol{x}, y)}{\partial \boldsymbol{\theta}} = \phi v \Big(- \frac{y}{\sigma(\boldsymbol{\theta}^T \boldsymbol{x})} + \frac{1-y}{1 - \sigma(\boldsymbol{\theta}^T \boldsymbol{x})} \Big) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \sigma(\boldsymbol{\theta}^T \boldsymbol{x})(1 - \sigma(\boldsymbol{\theta}^T \boldsymbol{x}))\boldsymbol{x} + \lambda \boldsymbol{\theta}. \tag{21}$$

Bidding Analysis

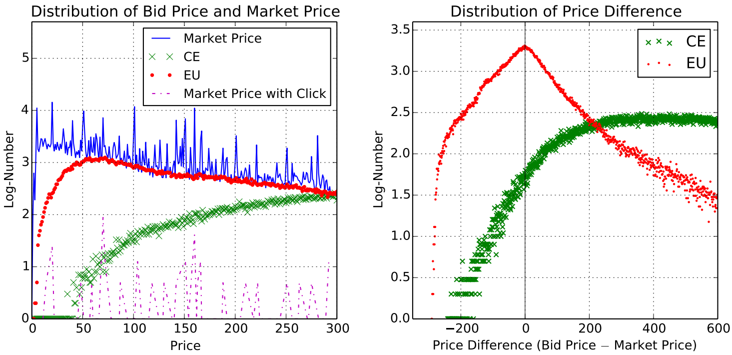

Figure: Analysis of bid price and market price distribution (iPinYou campaign 2259)

Online A/B Testing

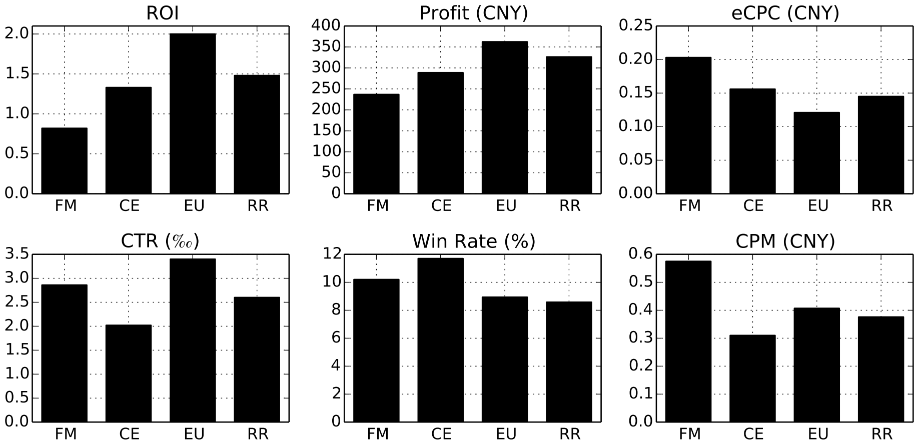

Figure: Online A/B testing results on YOYI PLUS, up to 25% improvements.

Outline

1. Background
   - Online Advertising
   - Real-time Bidding
   - Research Topics
2. Research Problems
   - User Response Prediction
   - Bidding Strategy Optimization
   - Reinforcement Learning for Advertising
   - Conversion Attribution
   - Bid Landscape Forecasting
3. Related Literatures

Decision Optimization: the Bidding Function

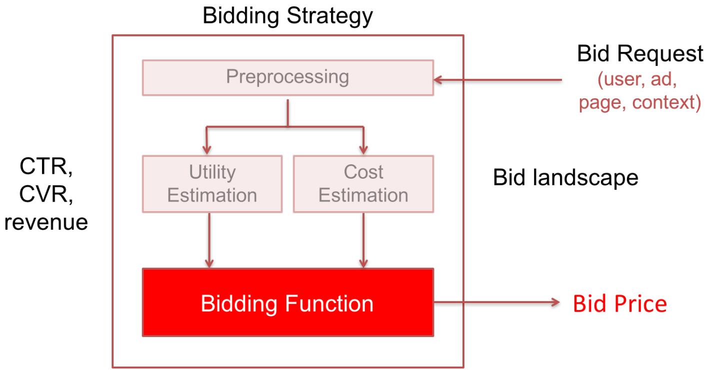

Problem Definition

Propose the optimal bidding function $b(\boldsymbol{x})$ to maximize the overall gains (clicks, conversions or profits), under the constraints of budget B.

Related Work (cont.)

Traditional Bidding Function

Truthful Bidding Function

$$b(\boldsymbol{x}) = V_{action} \cdot f(\boldsymbol{x}),$$

where f is the utility estimation function, such as pCTR.

Linear Bidding Function

$$b(\boldsymbol{x}) = \phi \cdot V_{action} \cdot f(\boldsymbol{x}) = b_0 \cdot f(\boldsymbol{x}).$$

C. Perlich et al. Bid optimizing and inventory scoring in targeted online advertising. KDD 2012.

Related Work (cont.)

Non-linear Bidding Function

Optimal Real-time Bidding Strategy

$$b()_{ORTB} = \arg\max_{b()} \int_{\boldsymbol{x}} \text{clicks } d\boldsymbol{x}$$

$$\text{subject to } \int_{\boldsymbol{x}} \text{expected\_costs } d\boldsymbol{x} \le B \ .$$

$\Rightarrow$

$$b()_{ORTB} = \arg\max_{b()} \int_{\boldsymbol{x}} f(\boldsymbol{x}) w(b(f(\boldsymbol{x}))) p_{\boldsymbol{x}}(\boldsymbol{x}) d\boldsymbol{x},$$

$$s.t. \int_{\boldsymbol{x}} b(f(\boldsymbol{x})) w(b(f(\boldsymbol{x}))) p_{\boldsymbol{x}}(\boldsymbol{x}) d\boldsymbol{x} \le B.$$

W. Zhang et al. Optimal Real-Time Bidding for Display Advertising. KDD 2014.

Related Work (cont.)

Non-linear Bidding Function

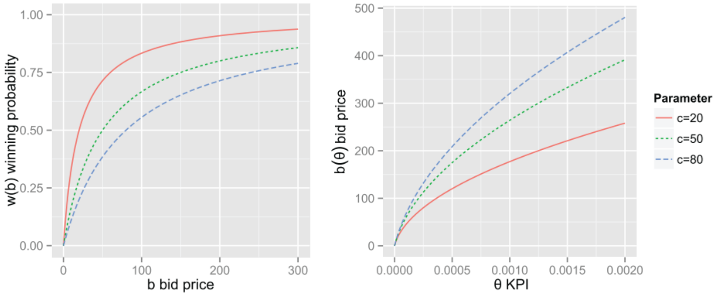

(a) Winning function 1. (b) Bidding function 1.

θ is the pCTR function, $w(\cdot)$ is the winning probability estimation function, $b(\cdot)$ is the bidding funciton.

W. Zhang et al. Optimal Real-Time Bidding for Display Advertising. KDD 2014.

Problems of the Related Work

Naive assumption for the bidding function.

Zhang's paper only considers first-price auction, which is not appropriate in practice.

Our Solution

Unified learning objective of the overall profits for utility estimation, cost estimation and bidding strategy optimization.

K. Ren et al. Bidding Machine: Learning to Bid for Directly Optimizing Profits in Display Advertising. TKDE 2018

Bidding Machine - Joint Optimization Framework

- $(\boldsymbol{x}, y)$: feature and label
- v: value of click (constant)
- $f_{\boldsymbol{\theta}}$: utility (CTR) estimation function
- b: bidding function (strategy)
- w_{ϕ}: winning probability
- c: expected cost

The expected profit formulation is

$$R(b, \boldsymbol{\theta}, \phi) = \int_{\boldsymbol{x}} [vy - c] w_{\phi} \cdot p_{\boldsymbol{x}}(\boldsymbol{x}) d\boldsymbol{x}$$

Bidding Machine - Joint Optimization Framework

$$w_{\phi}(b|\boldsymbol{x}) = \int_0^b p_z(z|\boldsymbol{x};\phi)dz,$$

$$c(b) = \frac{\int_0^b z p_z(z)dz}{\int_0^b p_z(z)dz},$$

$$\begin{aligned} R(b, \boldsymbol{\theta}, \boldsymbol{\phi}) &= \int_{\boldsymbol{x}} [vy - c(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))] w_{\boldsymbol{\phi}}(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) p_{\boldsymbol{x}}(\boldsymbol{x}) d\boldsymbol{x} \\ &= \sum_{(\boldsymbol{x},y)\in D} [vy - c(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))] w_{\boldsymbol{\phi}}(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))). \end{aligned}$$

Derivation Result 1

Some interesting findings

Table: The comparison of the model gradients (without regularization). LR: logistic regression, TB: truthful bidding, LB: linear bidding, UM: uniform market price distribution. LR and LR+TB+UM are equivalent (LR+TB reduces to the baseline LR when assuming the uniform market price distribution).

Model Setting	EU (SE) Gradient	RR (CE) Gradient
LR (baseline)	$\frac{\partial \mathcal{L}^{\text{SE}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = (\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y) \cdot \sigma(\boldsymbol{\theta}^T\boldsymbol{x})(1-\sigma(\boldsymbol{\theta}^T\boldsymbol{x}))\boldsymbol{x}$	$\frac{\partial \mathcal{L}^{\text{CE}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = (\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y)\boldsymbol{x}$
LR+TB	$-\frac{\partial R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = v^2(\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \sigma(\boldsymbol{\theta}^T\boldsymbol{x})(1-\sigma(\boldsymbol{\theta}^T\boldsymbol{x}))\boldsymbol{x}$	$-\frac{\partial R^{\text{RR}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = v(\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \boldsymbol{x}$
LR+TB+UM	$-\frac{\partial R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = v^2 l(\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y) \cdot \sigma(\boldsymbol{\theta}^T\boldsymbol{x})(1-\sigma(\boldsymbol{\theta}^T\boldsymbol{x}))\boldsymbol{x}$	$-\frac{\partial R^{\text{RR}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = vl(\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y)\boldsymbol{x}$
LR+LB	$-\frac{\partial R^{\text{EU}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = \phi v^2(\phi\sigma(\boldsymbol{\theta}^T\boldsymbol{x}) - y) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \sigma(\boldsymbol{\theta}^T\boldsymbol{x})(1-\sigma(\boldsymbol{\theta}^T\boldsymbol{x}))\boldsymbol{x}$	$-\frac{\partial R^{\text{RR}}_{\boldsymbol{\theta}}(\boldsymbol{x},y)}{\partial \boldsymbol{\theta}} = \phi v\Big(-\frac{y}{\phi\sigma(\boldsymbol{\theta}^T\boldsymbol{x})} + \frac{1-y}{1-\phi\sigma(\boldsymbol{\theta}^T\boldsymbol{x})}\Big) \cdot p_z(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) \cdot \sigma(\boldsymbol{\theta}^T\boldsymbol{x})(1-\sigma(\boldsymbol{\theta}^T\boldsymbol{x}))\boldsymbol{x}$

Derivation Result 2

Optimal Bidding Strategy for Profit Optimization under 2nd Price Auction **without Budget Constraint**

$$
\begin{aligned}
R(b, \boldsymbol{\theta}, \boldsymbol{\phi}) &= \int_{\boldsymbol{x}} [vy - c(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))] w_{\boldsymbol{\phi}}(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))) p_{\boldsymbol{x}}(\boldsymbol{x}) d\boldsymbol{x} \\
&= \sum_{(\boldsymbol{x}, y) \in D} [vy - c(b(f_{\boldsymbol{\theta}}(\boldsymbol{x})))] w_{\boldsymbol{\phi}}(b(f_{\boldsymbol{\theta}}(\boldsymbol{x}))).
\end{aligned}
$$

Theorem

We can theoretically prove that the optimal bidding function is the truthful bidding for profit maximization under second-price auction.

$$
b(\boldsymbol{x}) = v \cdot f(\boldsymbol{x}) ,
$$

Derivation Result 3

Optimal Bidding Strategy for Profit Optimization under 2nd Price Auction

Theorem

The optimal bidding function under a symmetric game of repeated auctions with budget constraints is linear to the estimated utility.

$$\begin{aligned} \max_{b()} \quad & T \int_r [u(r) - c(b(\tau))] w_b(b(\tau)) p_r(r) dr \ , \\ s.t. \quad & T \int_r c(b(\tau)) w_b(b(\tau)) p_r(r) dr = B \ , \end{aligned} \tag{22}$$

here we assume that the bidding is based on a *signal* τ related with the CTR $r = f(\boldsymbol{x})$.

Here v is the click value of the advertiser.We derive in the paper that

$$b(r) = \frac{vr}{\lambda + 1} \ . \tag{23}$$

Derivation Result 4

What if all the advertisers adopt the same bidding strategy?

Derivation Result 4

What if all the advertisers adopt the same bidding strategy?

Theorem

The bid price is monotonously **increasing** *w.r.t. the number of the participating advertiser bidders, and the* **tragedy of the commons** *will occur in the market.*

$$b(r) = \frac{vr}{\lambda + 1} \Rightarrow \frac{Br}{T \int_r \int_0^r t(n-1)F_r(t)^{n-2} p_r(t)\, dt \; p_r(r) dr}. \tag{24}$$

Derivation Result 4

What if all the advertisers adopt the same bidding strategy?

Theorem

The bid price is monotonously **increasing** *w.r.t. the number of the participating advertiser bidders, and the* **tragedy of the commons** *will occur in the market.*

$$b(r) = \frac{vr}{\lambda + 1} \Rightarrow \frac{Br}{T \int_r \int_0^r t(n-1)F_r(t)^{n-2} p_r(t)\, dt \; p_r(r) dr}. \tag{24}$$

The profit of the platform will increase :)

Bidding Machine - Joint Optimization Framework

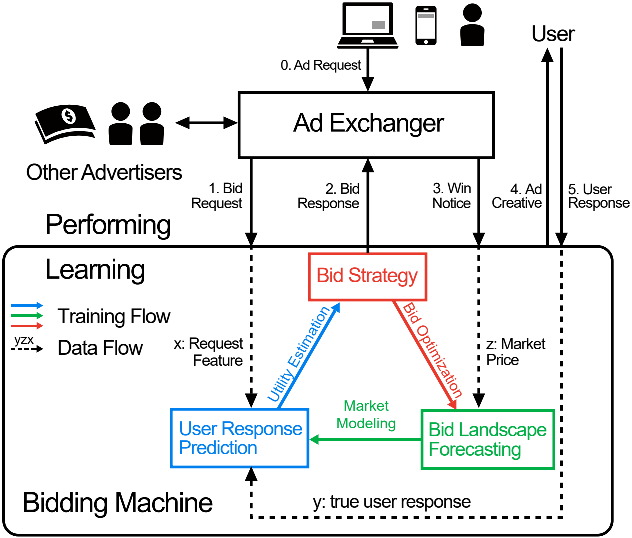

Bidding Machine

Exp. Results

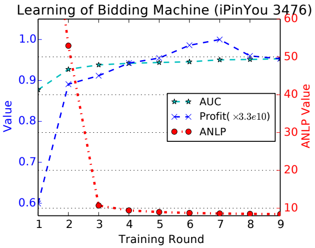

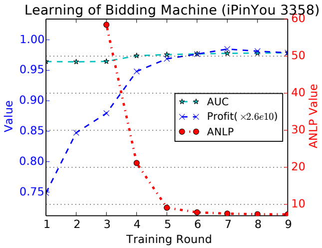

Bidding Machine

Offline Results

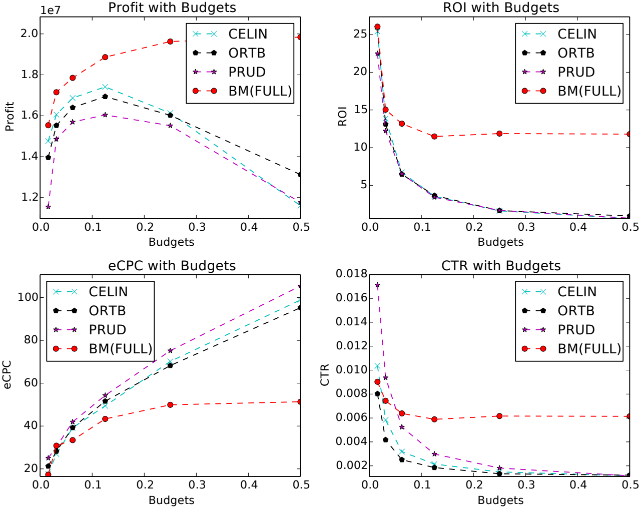

Bidding Machine

Offline Results

Table: Campaign profit for Single CTR estimation and Binary Optimization with market modeling.

		1458	2259	2261	2821	2997	3358	3386	3427	3476	Average
AUC	EU	.987	.674	.622	.608	.606	.970	.761	.976	.954	.795
	RR	.977	.691	.619	.639	.608	.980	.778	.960	.950	.800
	BM(MKT)	.981	.678	.647	.620	.603	.980	.788	.973	.955	**.803**
Profits ($\times 10^7$)	EU	3.91	.732	.797	.539	.147	2.42	1.58	3.05	3.25	1.82
	RR	3.98	.803	.827	.572	.141	2.54	1.64	3.14	3.39	1.89
	BM(MKT)	4.02	.766	.863	.669	.148	2.57	1.73	3.18	3.31	**1.91**
ROI	EU	19.2	.607	.582	.333	.679	9.26	1.46	5.30	4.02	4.60
	RR	24.3	1.03	.771	.247	.624	9.29	1.90	9.57	6.63	6.04
	BM(MKT)	31.7	.829	.692	.476	.733	8.83	1.08	9.70	5.40	**6.61**
eCPC	EU	4.27	172	187	104	8.33	11.4	42.5	17.3	30.0	64.3
	RR	3.39	136	167	112	8.61	11.4	36.1	10.3	19.7	**56.1**
	BM(MKT)	2.62	151	175	94.7	8.07	11.9	50.2	10.1	23.5	58.7

Bidding Machine

Online Results 1

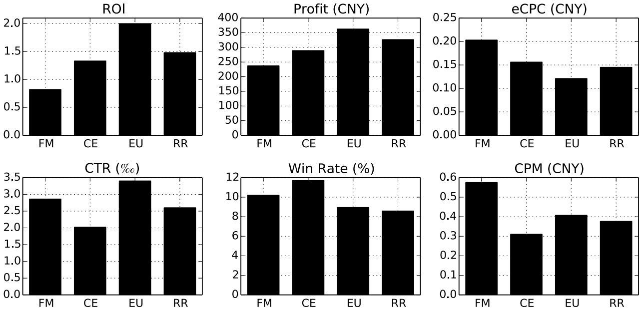

Figure: Online results on YOYI MOBILE (Phase I in 2016). Up to 25% improvement over traditional CTR model on profits.

Bidding Machine

Online Results 2

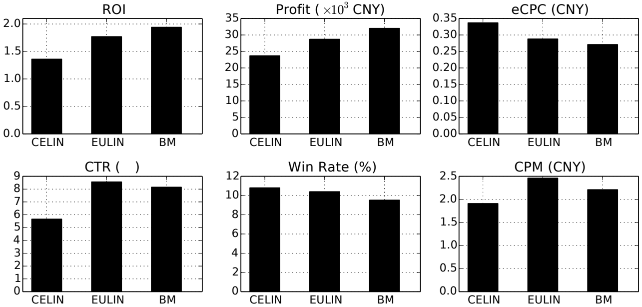

Figure: Online results on YOYI MOBILE (Phase II in 2017). Up to 8% improvements over traditional linear bidding methods.

Outline

Bidding as Sequential Decision Making

Relationship between RTB & RL

- Real-time Bidding is a sequence of decision making.
- The goal is to maximize the cumulative rewards (clicks, etc.) of the advertiser (bidder).
- The constraint is the total budget of the advertiser.

RL for Bidding

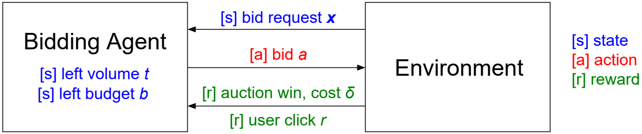

Figure: Real-time Bidding as reinforcement learning.

MDP Formulation

state: $(t, b, \boldsymbol{x})$ with remaining auction num. t, remaining budget b and the received auction feature $\boldsymbol{x}$.

action: bid price.

reward: predicted CTR (as model-baed RL).

$P = Pr((t-1, b-\delta, \boldsymbol{x}_{t-1}) | (t, b, \boldsymbol{x}_t), \delta)$ which is the winning probability, where *delta* is the bid price.

1) H. Cai, **K. Ren**, et al. Real-Time Bidding by Reinforcement Learning in Display Advertising. WSDM 2017
2) Y. Song, **K. Ren**, et al. Volume Ranking and Sequential Selection Programmatic Display Advertising. CIKM 2017

Learned Value Function

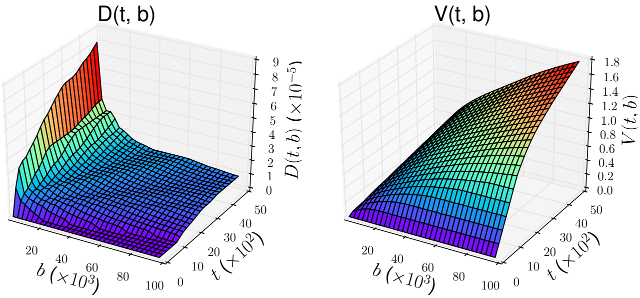

Figure: The learned value function over states.

Online Results

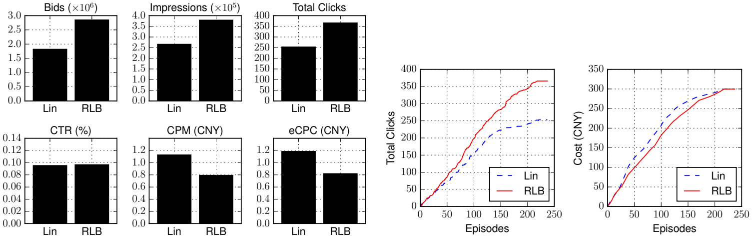

(a) The overall results. (b) Performance over time.

Figure: The online results on VLion ad platform.

Outline

Conversion Attribution

Problem Definition

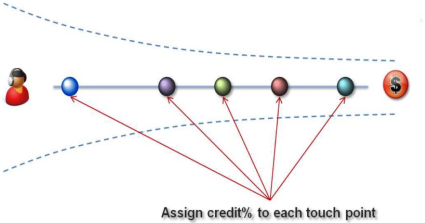

Two view of the problem

Horizontal View: Given a sequence of user activities leading to a conversion, assign the attribution credits to each touch point for the (negative) contribution on the final conversion.
Vertical View: Calculate the conversion attribution over difference channels or subcampaigns.

Problem Challenge: Multi-touch Conversion Attribution

Cons of the traditional methods

Prediction upon single point: ignore the sequential data patterns in model training.

Rule-based method: heuristically assign the conversion credits on to the multiple touches.

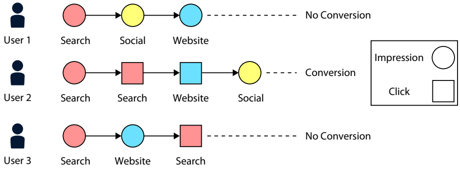

Our Solution

Use recurrent neural network to model the sequential user activities.

Assign "attention" to the touch points to model the conversion attributions.

Simultaneously model impression-level and click-level patterns for conversion estimation.

Dual-attention Mechanism for Conversion Attribution

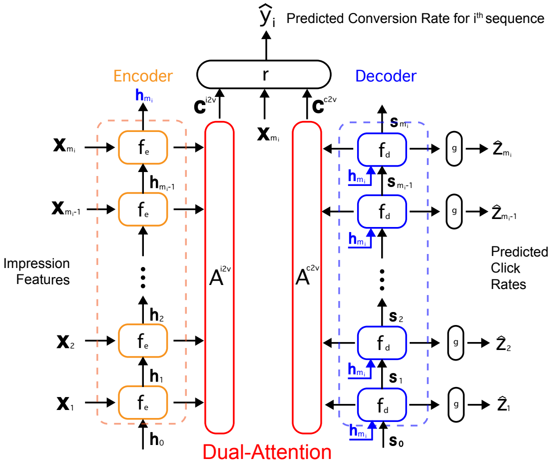

K. Ren et al. Learning Conversion Attribution with Dual-attention Mechanism for Online Advertising. CIKM, 2018.

Attention Implementation

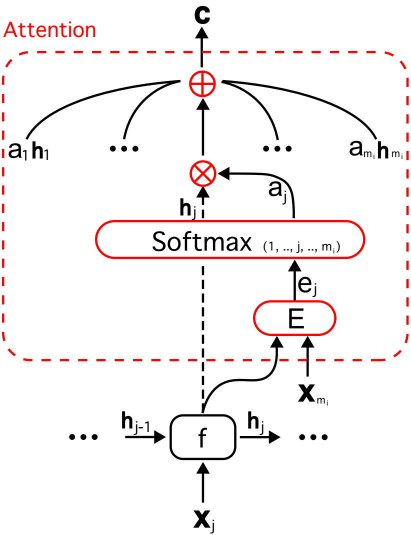

Visualization of the Attribution

Horizontal Sequence Level

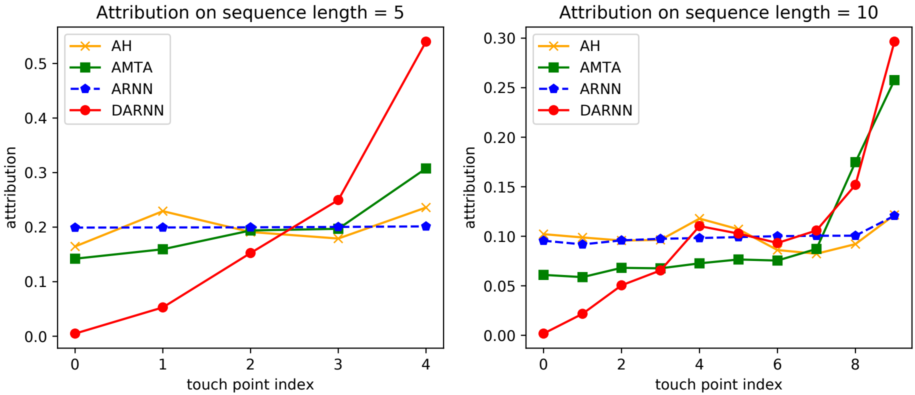

Figure: Touch point level attribution statistics (Miaozhen).

Visualization of the Attribution

Vertical Channel Level

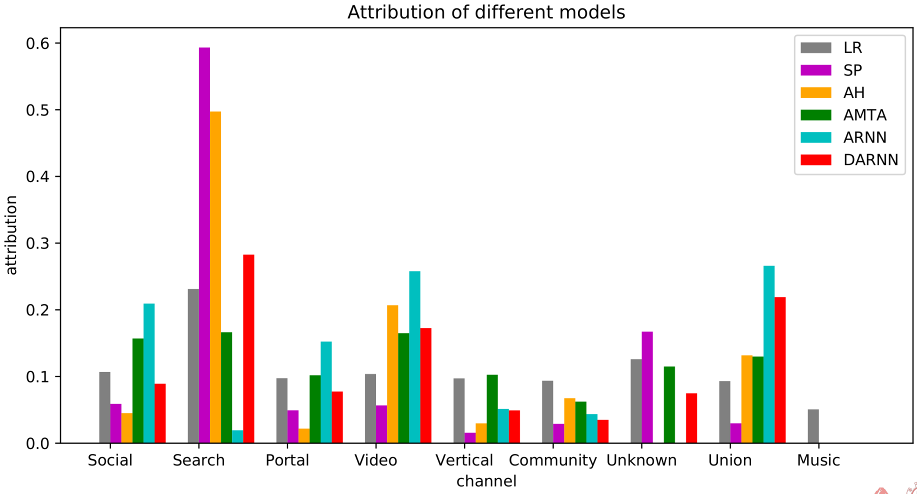

Figure: Attribution of different channels on Miaozhen.

Visualization of the Attribution Preferences

Click-level v.s. Impression-level

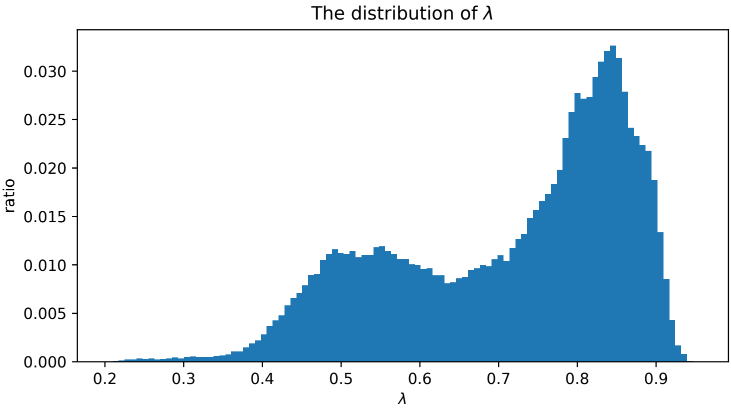

Figure: The distribution of λ over Criteo dataset.

Outline

Cost Estimation: Bid Landscape Forecasting

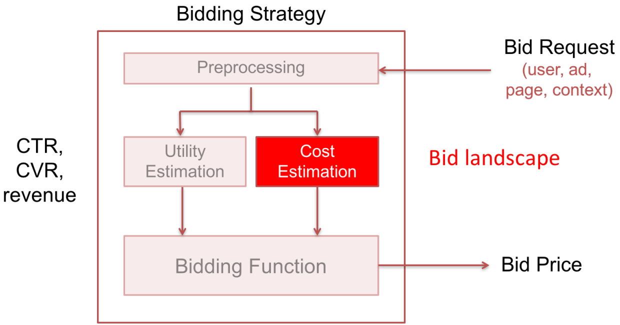

Challenge: Modeling Market Price Distribution

Win and pay the **second** highest price, lose otherwise do nothing.

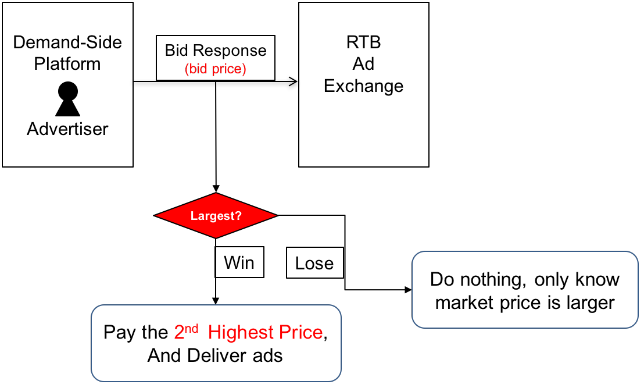

Cost Estimation: Bid Landscape Forecasting

Problem Definition

Model the probability density function $p_z(z; \boldsymbol{x})$ of the market price z w.r.t. the given feature $\boldsymbol{x}$.

- Date: 20160320
- Hour: 14
- Weekday: 7
- IP: 119.163.222.*
- Region: England
- City: London
- Country: UK
- Ad Exchange: Google
- Domain: yahoo.co.uk
- URL: http://www.yahoo.co.uk/abc/xyz.html
- OS: Windows
- Browser: Chrome
- Ad size: 300*250
- Ad ID: a1890
- User tags: Sports, Electronics

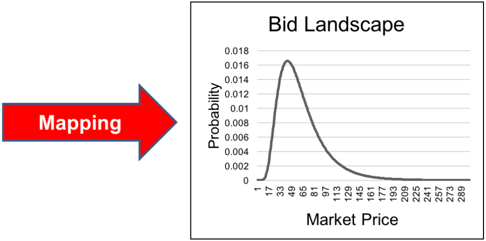

Problem Analysis

- No ground truth for either P.D.F. or C.D.F. of the market price.
- There are censored data to handle (without knowledge of the true market price).
- Fine-grained forecasting for each individual sample.

Related Work: Heuristic Form

Log-normal Form

$$p_z(z) = \frac{1}{z\sigma\sqrt{2\pi}} e^{\frac{-(\ln z - \mu)^2}{2\sigma^2}}, z > 0 \ .$$

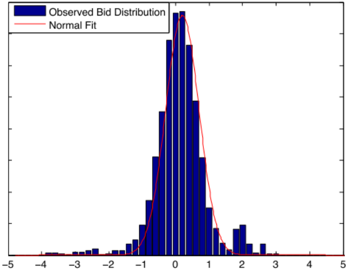

Y. Cui et al. Bid landscape forecasting in online ad exchange marketplace. KDD 2011

Related Work: Regression Model

z_i as the predicted winning price,

$$z_i \approx \beta^T x_i + \epsilon_i \ ,$$

$$\epsilon_i \sim N(0, \sigma^2) \ ,$$

$$\text{maximize} \sum_{i \in W} \log(\phi(\frac{w_i - \beta_W^T x_i}{\sigma})) \ .$$

and mixinize the log likelihood.

W. Wu et al. Predicting Winning Price in Real Time Bidding with Censored Data. KDD 2015

Challenge: Modeling Right Censored Data

Right Censored

Right Censorship

As in 2^{nd} price auction, if you *lose*, you only know that the *market price* is higher than your bidding price, which result in right censorship.

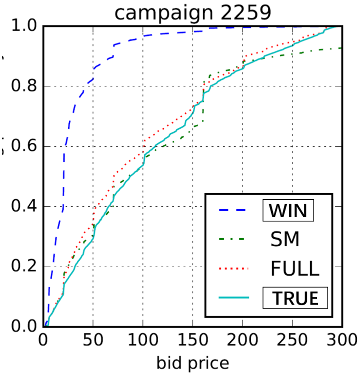

Handling Censorship with Kaplan-Merier Estimator

For winning auctions: We have the true market price value.
For lost auctions: We only know our proposed bid price and know that the true market price is higher than that.

Intuition

More than considering Winning Logs, but also utilize Losing Logs.

Idea: Modeling Winning (Dying) Likelihood

$$w(b_x) = 1 - \prod_{b_j < b_x} \frac{n_j - d_j}{n_j}, \quad p(z) = w(z+1) - w(z). \tag{25}$$

$b_j < b_{j+1}$, d_j is number of winning auctions by $b_j - 1$, n_j is number of lost auctions by $b_j - 1$. So

$$l(b_x) = \prod_{b_j < b_x} \frac{n_j - d_j}{n_j}.$$

Tree-based Mapping

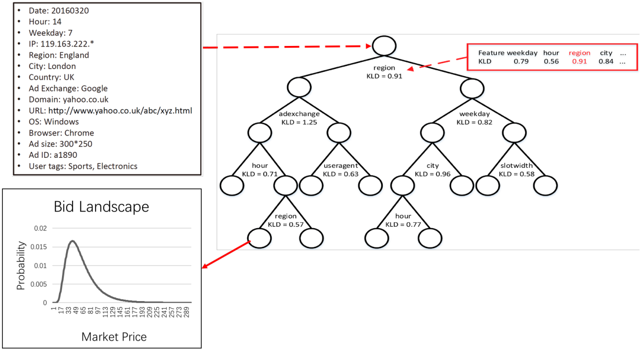

Censorship Handling

Using Kaplan Meier estimator to capture the right censored patterns.

Y. Wang, **K. Ren**, W. Zhang, Y. Yu. Functional Bid Landscape Forecasting for Display Advertising. ECML-PKDD, 2016.

Results

Table: Performance illustration. Average negative log probability (likelihood) of five compared settings. ANLP: the smaller, the better.

	ANLP				
Campaign	MM	NM	SM	NTM	STM
1458	5.7887	5.3662	4.7885	4.7160	4.3308
2259	7.3285	6.7686	5.8204	5.4943	5.4021
2261	7.0205	5.5310	5.1053	4.4444	4.3137
2821	7.2628	6.5508	5.6710	5.4196	5.3721
2997	6.7024	5.3642	5.1411	5.1626	5.0944
3358	7.1779	5.8345	5.2771	4.8377	4.6168
3386	6.1418	5.2791	4.8721	4.6698	4.2577
3427	6.1852	4.8838	4.6453	4.1047	4.0580
3476	6.0220	5.2884	4.7535	4.3516	4.2951
overall	6.5520	5.6635	5.0997	4.7792	**4.6065**

Related Work: Censorship Handling with Mixture Model

$$
\begin{aligned}
z_i &= [Pr(z_i < b_i)\beta_W + (1 - Pr(z_i < b_i))\beta_L]^T x_i \\
&= \beta_{mix}^T x_i \ , \\
Pr(z_i < b_i) &= p(\boldsymbol{x}) \ .
\end{aligned}
\tag{26}
$$

W. Wu et al. Predicting Winning Price in Real Time Bidding with Censored Data. KDD 2015

Related Work

DeepHit Model for Survival Analysis

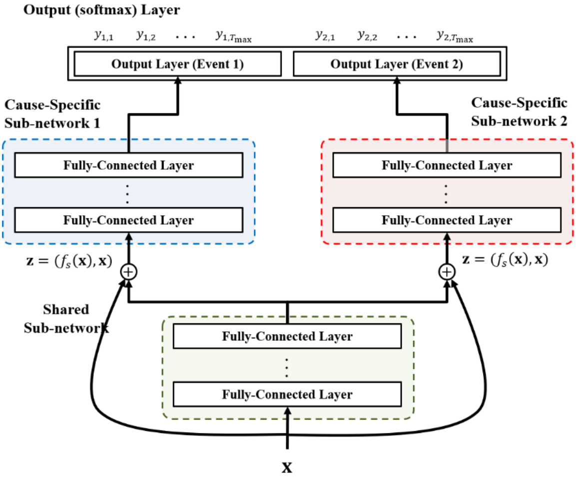

C. Lee et al. DeepHit: A Deep Learning Approach to Survival Analysis with Competing Risks. AAAI 2018

Problems in the Related Work

- Heuristic assumption on the distribution.
- Naive ensorship handling
  - Mixture model: combines probability and cumulative probability in a simple way.
  - Tree-based model: uses counting-based statistics for censorship handling.
- DeepHit model: Sparse gradient signals, without consideration of sequential patterns along time.

Our Solution (Under Review)

- Using deep recurrent neural network to model the event rate at each timestep (price).
- Using maximum partial likelihood for censorship handling.

Deep Survival Analysis

- We utilize recurrent neural network to model sequential patterns in the time series space.
- We also adopt partial likelihood for censorship handling.
- The model achieves state-of-the-art performance.
- It can also inspire the survival analysis in other fields such as clinical research.

Related Literatures

Ren K, Fang Y, Zhang W, et al. Learning Multi-touch Conversion Attribution with Dual-attention Mechanisms for Online Advertising[C]//Proceedings of the 27th ACM International on Conference on Information and Knowledge Management.

Ren K, Zhang W, Chang K, et al. Bidding Machine: Learning to Bid for Directly Optimizing Profits in Display Advertising[J]. IEEE Transactions on Knowledge and Data Engineering, 2018.

Ren K, Zhang W, Rong Y, et al. User response learning for directly optimizing campaign performance in display advertising[C]//Proceedings of the 25th ACM International on Conference on Information and Knowledge Management. ACM, 2016: 679-688.

Song Y, **Ren K**, Cai H, et al. Volume Ranking and Sequential Selection in Programmatic Display Advertising[J]. 2017.

Qu Y, Cai H, **Ren K**, et al. Product-based neural networks for user response prediction[J]. 2016 IEEE 16th International Conference on Data Mining (ICDM), Barcelona, 2016, pp. 1149-1154.

Wang Y, **Ren K**, Zhang W, et al. Functional bid landscape forecasting for display advertising[C]//Joint European Conference on Machine Learning and Knowledge Discovery in Databases. Springer International Publishing, 2016: 115-131.

Cai H, **Ren K**, Zhang W, et al. Real-Time Bidding by Reinforcement Learning in Display Advertising[C]//10th ACM International Conference on Web Search and Data Mining (WSDM), 2017.

Other and Working Paper

Ren K, et al. Deep Survival Analysis. Working paper.

Lantao Y*, Xuejian W* (Equal Contribution), **Ren K**, et al. A Dynamic Attention Deep Model for Article Recommendation by Learning Human Editors' Demonstration[C]// Proceedings of the 23th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM, 2017.

Zhou Z, Cai H, Rong S, Song Y, **Ren K**, Zhang W, Wang J, Yu Y. Activation Maximization Generative Adversarial Nets[J]. 2018.

Zhu C, **Ren K**, Liu X, et al. A Graph Traversal Based Approach to Answer Non-Aggregation Questions Over DBpedia[C]//Joint International Semantic Technology Conference. Springer, Cham, 2015: 219-234.

Thank you for your attention!

http://saying.ren
rk.ren@outlook.com