WSDM 2021
THE 14TH ACM INTERNATIONAL CONFERENCE ON WEB SEARCH AND DATA MINING
MARCH 8-12, 2021 • JERUSALEM, ISRAEL

Alibaba Group 阿里巴巴集团

Optimizing Multiple Performance Metrics with Deep GSP Auctions for E-commerce Advertising

Zhilin Zhang*, Xiangyu Liu*, Zhenzhe Zheng, Chenrui Zhang, Miao Xu, Junwei Pan,
Chuan Yu, Fan Wu, Jian Xu and Kun Gai

Alibaba Group

in collaboration with

SHANGHAI JIAO TONG UNIVERSITY 1896 · PEKING UNIVERSITY 1898 · YAHOO! RESEARCH

Background

What do we focus on ?

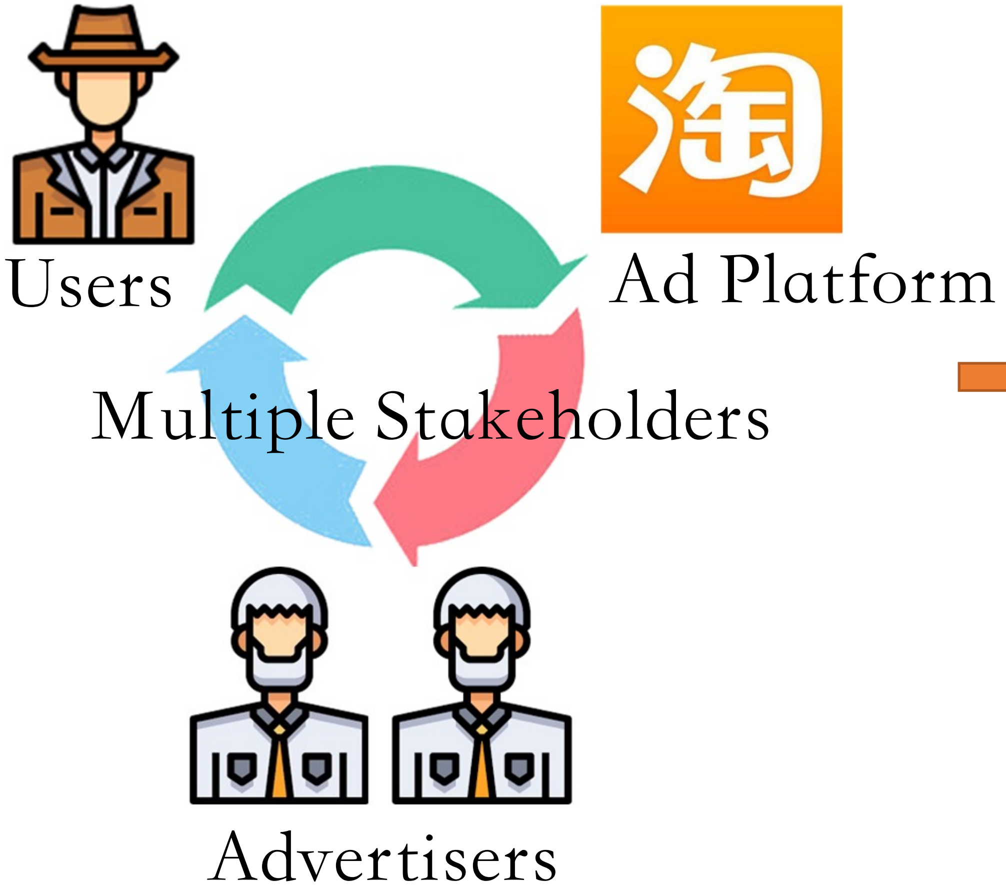

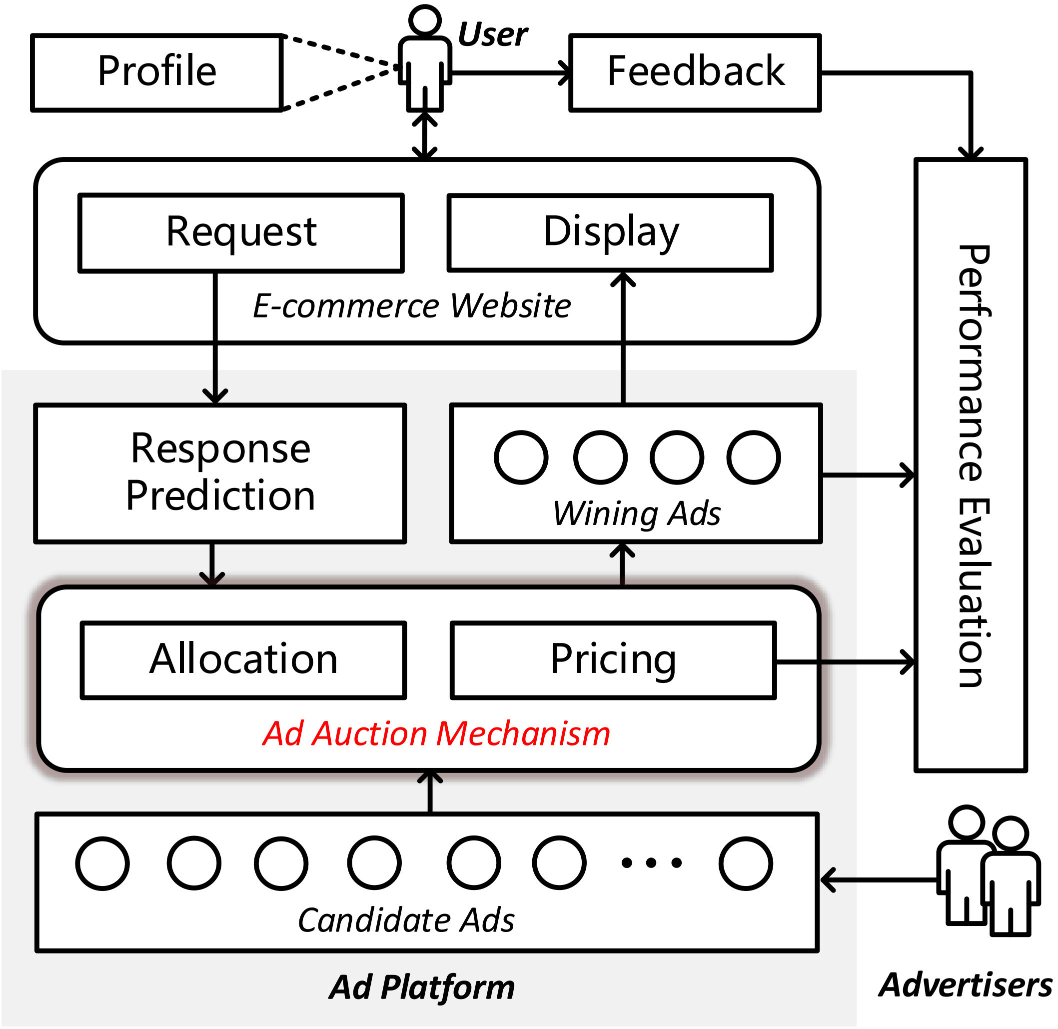

What do we focus on ?

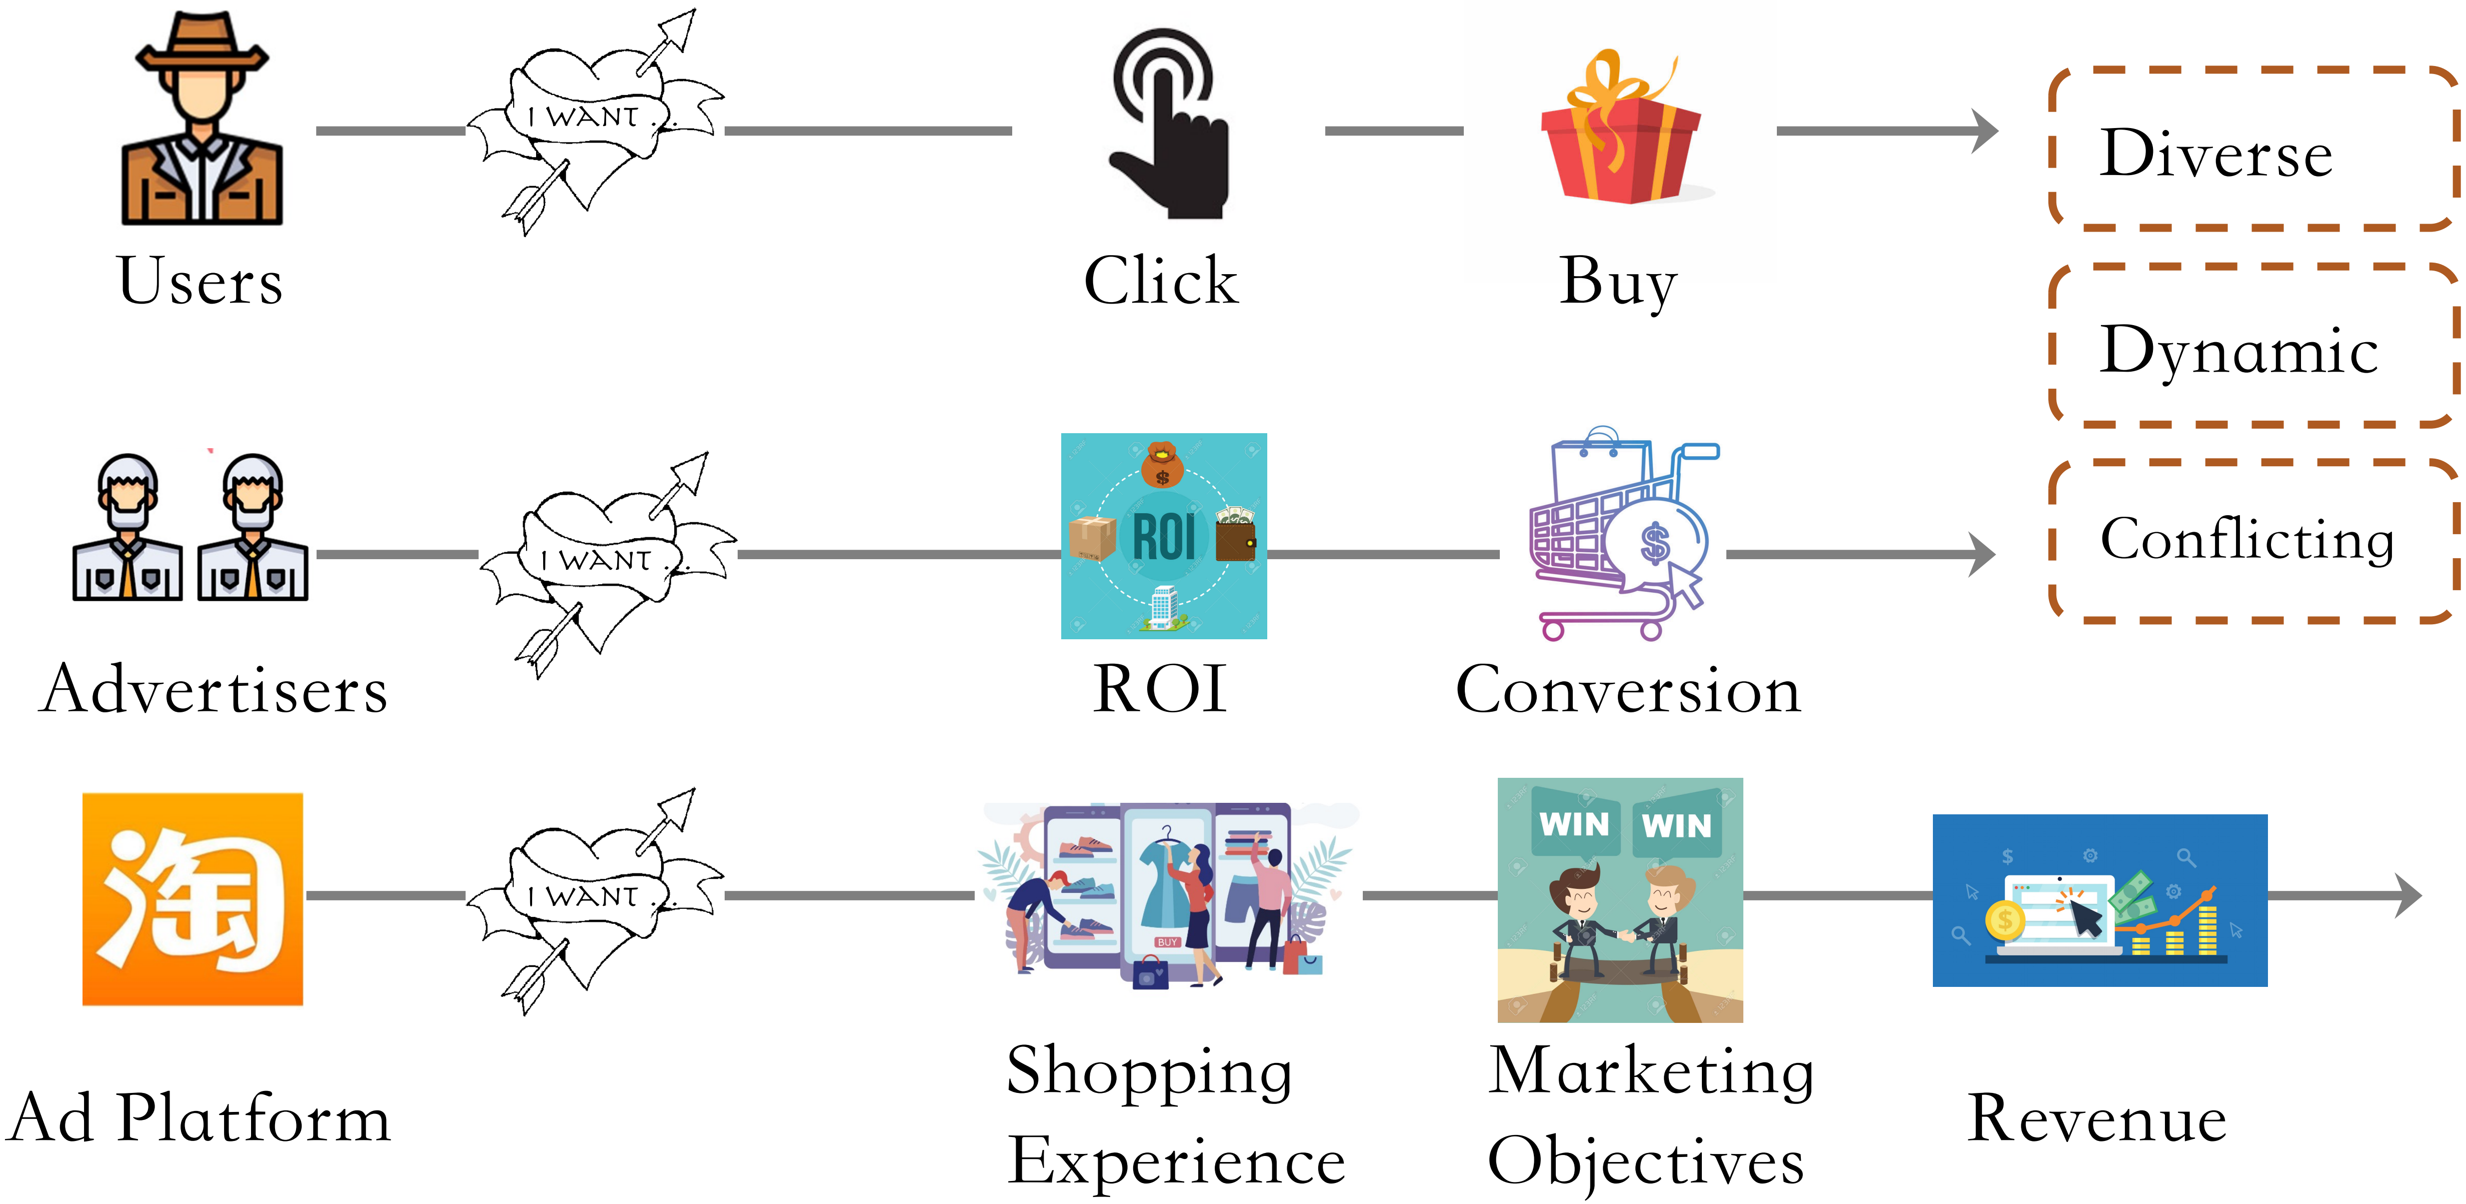

Challenges

✓The performance metrics from different stakeholders often **conflict** with each other.

✓Some metrics are **difficult to estimate** with prediction models.

✓It is substantially different from the traditional **multi-objective optimization**.

✓While the auction theory provides a rich set of tools for optimizing social welfare or revenue, few of them can be used to optimize the above mentioned **diverse, dynamic, conflicting and feedback-based performance metrics**.

Our Solution

GSP

+

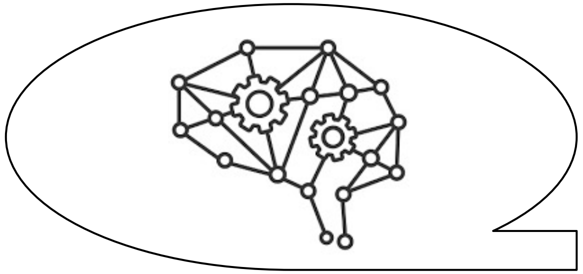

Deep Learning

=

Deep GSP

Problem Formulation

Formulation:

Multiple performance metrics optimization in the competitive advertising environments

Mechanism: $\mathcal{M}\langle\mathcal{R},\mathcal{P}\rangle$ with allocation $\mathcal{R}$ and pricing $\mathcal{P}$

- $\mathcal{R}$: Select K ads from N candidates
- $\mathcal{P}$: Calculating payment on K wining ads

Goals:

1. Optimizing L performance metrics (RPM, CTR, CVR, GMV, etc.)
2. Desirable mechanism properties: Game Equilibrium & Smooth Transition

$$
\begin{aligned}
\underset{\mathcal{M}}{\text{maximize}} \quad & \mathbb{E}_{\mathbf{b}\sim\mathcal{D}}\left[\sum_{j=1}^{L} w_j \times f_j(\mathbf{b};\mathcal{M})\right] \\
\text{s.t.} \quad & \textit{Game Equilibrium constraints}, \\
& \textit{Smooth Transition constraints},
\end{aligned}
$$

Game Equilibrium (GE)

- Single-slot: **Incentive Compatible (IC)**
  - Myerson Theorem

$$\mathcal{R}_i(z, \boldsymbol{b}_{-i}) \geq \mathcal{R}_i(b_i, \boldsymbol{b}_{-i}) \text{ if } z > b_i$$

$$\mathcal{P}_i = \inf_{z \mid \mathcal{R}_i(z, \boldsymbol{b}_{-i}) = \mathcal{R}_i(\boldsymbol{b})}$$

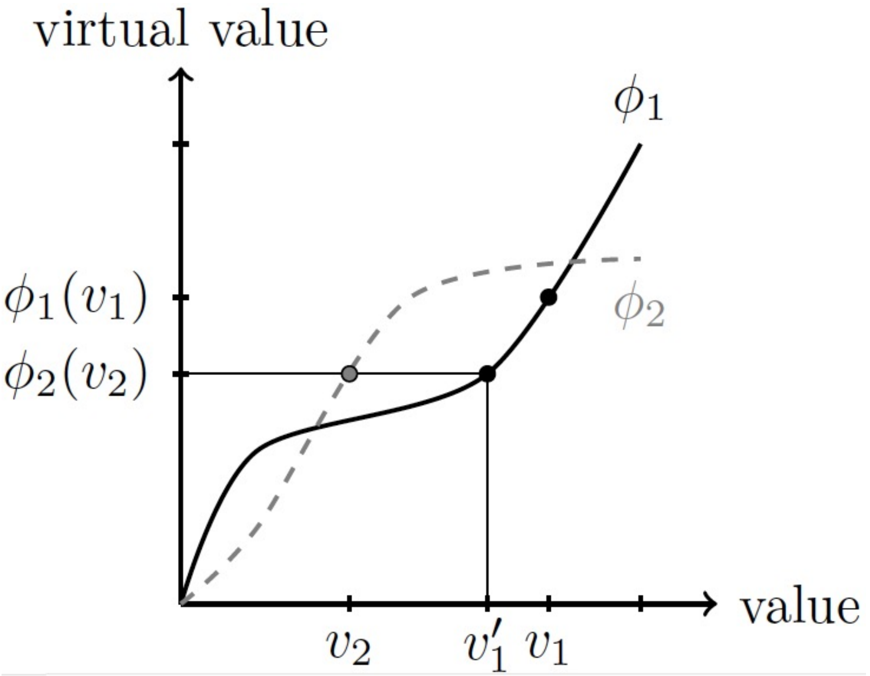

- Multi-slots: symmetric Nash equilibrium (SNE)

$$\beta_i(v_i - p_i) \geq \beta_j(v_i - p_j)$$

- **For Deep GSP**:
  - $\mathcal{R}$: $r_i = R_\theta(b_i, \mathbf{x}_i)$
  - $\mathcal{P}$: $p_i = R_\theta^{-1}(r_{i+1}, \mathbf{x}_i)$

Roger B Myerson. Optimal auction design. Mathematics of operations research, 1981.

Smooth Transition (ST)

The advertiser's utility would not fluctuate too much when the auction mechanism is switched towards optimizing another objective.

$$u_i(\mathcal{M}) \geq (1-\epsilon) \times \bar{u}_i(\mathcal{M}_0)$$

- u_i : advertiser i's utility
- $\mathcal{M}_0$: a benchmark mechanism
- ϵ : a tolerant utility loss ratio for advertisers
- $\bar{u}(\mathcal{M}_0)$: the lower bound of utility, set as the average utility over a certain period under the benchmark mechanism

Our Approach: Deep GSP

$\mathcal{R}$: Monotone Allocation

Point-wise monotonicity constraint

- Deep rank score function:

$$r_i = R_\theta(b_i, \mathbf{x}_i) = b_i \times \pi_\theta(b_i, \mathbf{x}_i),$$

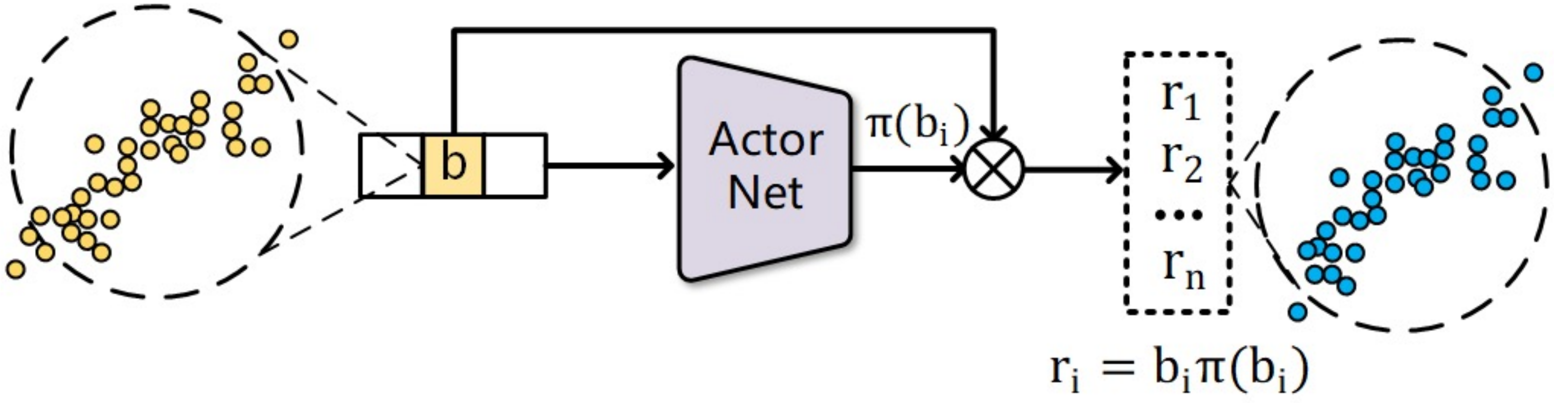

$$\mathcal{L}_{mono} = \sum_{i=1}^{N} \max(0, -\nabla_b R_\theta(b_i, \mathbf{x}_i)) = \sum_{i=1}^{N} \max(0, -(\pi_\theta(b_i, \mathbf{x}_i) + b_i \nabla_b \pi_\theta(b_i, \mathbf{x}_i)))$$

Monotonicity Evaluation

- Red markers: the real reported bid values

$$\mathcal{T}_m = \frac{1}{n} \sum \rho_{rank_{bids}, rank_{outputs}}$$

Exp	Metrics Configuration	$\mathcal{T}_m$
1	(1,0,0,0,0)	0.991
2	(0.5,0.5,0,0,0)	0.960
3	(0.5,0,0.5,0,0)	0.978
4	(0.5,0,0,0.5,0)	0.972
5	(0.5,0,0,0,0.5)	0.982
6	(0.6,0.1,0.1,0.1,0.1)	0.975

Offline Monotonicity performance

- $\mathcal{T}_m$: all above 0.96

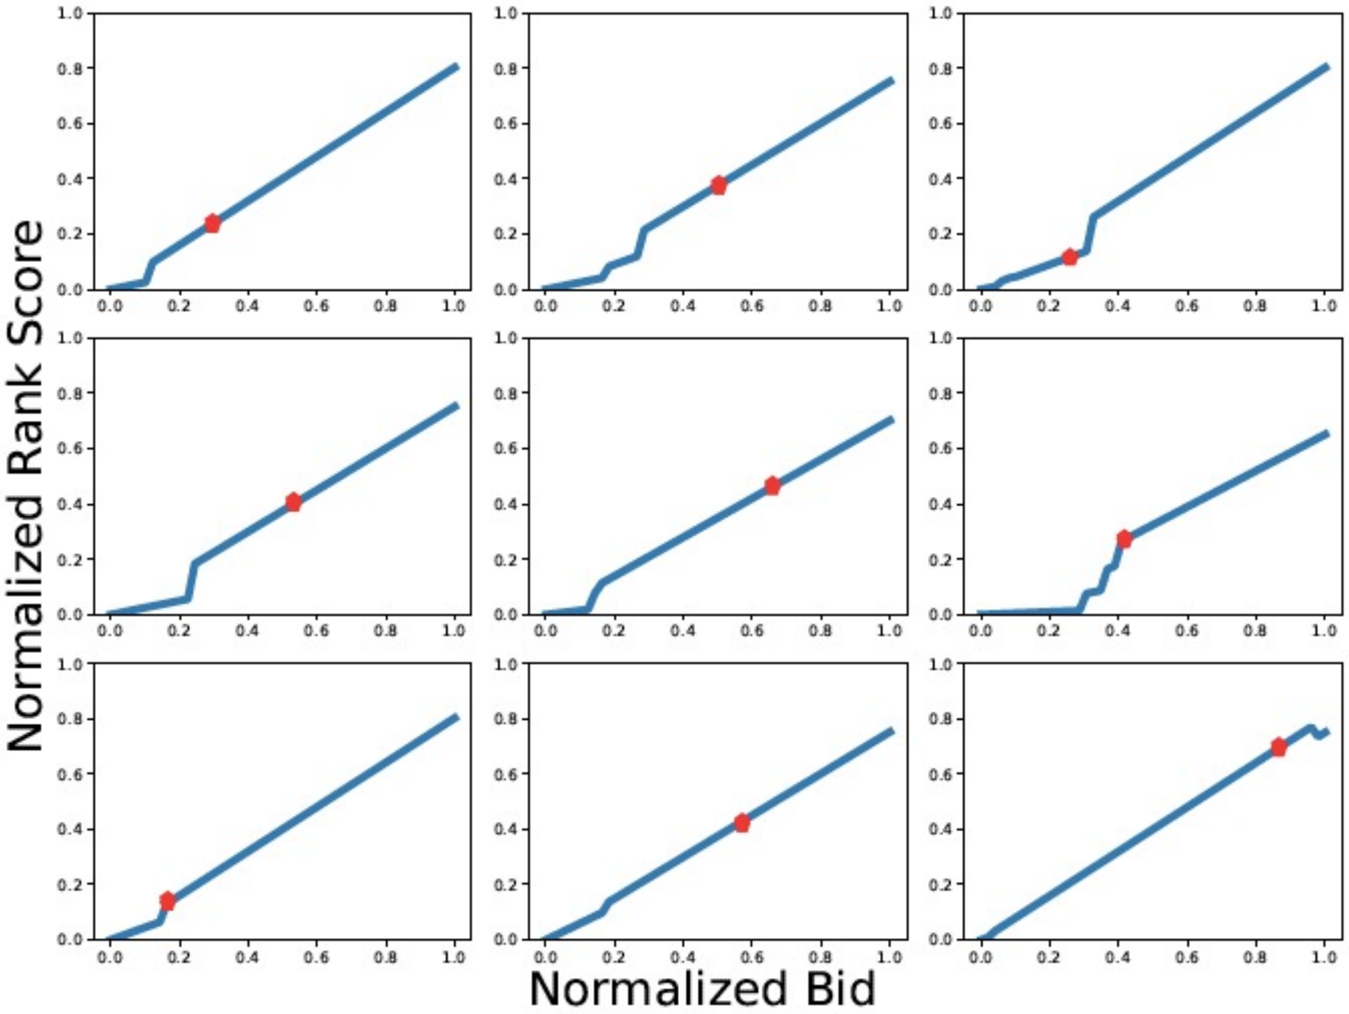

$\mathcal{P}$: Approximate Inverse Solution

- Challenge of DNN-based critical bid pricing:
  - pseudo-inverse matrices layer-by-layer
  - weight matrices are singular

- $\pi_\theta(b_i, \mathbf{x}_i)$ is not sensitive to the bid → regard it as a constant

$$r_i = R_\theta(b_i, \mathbf{x}_i) = b_i \times \pi_\theta(b_i, \mathbf{x}_i),$$

$$\boxed{p_i = R_\theta^{-1}(r_{i+1}, \mathbf{x}_i) = \frac{r_{i+1}}{\pi_\theta(b_i, \mathbf{x}_i)}}$$

- The approximate inverse solution pi does not introduce much bias.

Exp	Metrics Configuration	$\mathcal{T}_m$	PER
1	(1,0,0,0,0)	0.991	1.009
2	(0.5,0.5,0,0,0)	0.960	0.994
3	(0.5,0,0.5,0,0)	0.978	0.988
4	(0.5,0,0,0.5,0)	0.972	0.995
5	(0.5,0,0,0,0.5)	0.982	0.999
6	(0.6,0.1,0.1,0.1,0.1)	0.975	0.995

PER: error ratio

$\mathcal{R}$ & $\mathcal{P}$: Incentive Compatibility Evaluation

- Individual Stage-IC Metric: [0, 1]
- {0, 1} → a more nuanced comparison

applies small perturbations to bids

$$\text{i-SIC} = \lim_{\alpha \to 0} \frac{\mathbb{E}_{v\sim F}\left[\hat{u}((1+\alpha)v)\right] - \mathbb{E}_{v\sim F}\left[\hat{u}((1-\alpha)v)\right]}{2\alpha \cdot \mathbb{E}_{v\sim F}\left[v \cdot x(v)\right]}$$

$$\hat{u}(b) = b \times x(b) - p(b)$$

Exp	Metrics Configuration	$\mathcal{T}_m$	PER	IC
1	(1,0,0,0,0)	0.991	1.009	0.9878
2	(0.5,0.5,0,0,0)	0.960	0.994	0.9910
3	(0.5,0,0.5,0,0)	0.978	0.988	0.9903
4	(0.5,0,0,0.5,0)	0.972	0.995	0.9817
5	(0.5,0,0,0,0.5)	0.982	0.999	0.9856
6	(0.6,0.1,0.1,0.1,0.1)	0.975	0.995	0.9941

- Deep GSP can guarantee the IC property to some extent, which is meaningful to benefit the long-term healthy development of the whole advertising ecology.

Deng et al. A Data-Driven Metric of Incentive Compatibility, WWW 2020

Optimization

- Model-free RL-based Optimization Framework
  - We can only evaluate these metrics via actual feedback after deploying the auction mechanism.
  - This phenomenon is similar to the exploration process in reinforcement learning
- State:
  - Ad information: bid, $pCTR$, $pCVR$, and ad category, etc
  - Advertisers' information: current budget, the price of products, marketing intent, etc
  - User features: gender, age, income level, shopping preferences, etc
- Action:
  - the outcome of the deep rank score model
- Reward:

$$re_i = \sum_i w_j \times f_j - \eta \times \max(0, (1-\epsilon) \times \bar{u}(\mathcal{M}^0) - u(\mathcal{M}))$$

- No transition
- Goal:

$$R_\theta^* = \underset{R_\theta}{\operatorname{argmax}} \mathbb{E}_{\mathbf{b} \sim \mathcal{D}}[re_i | R_\theta]$$

Deep GSP Framework

- Actor-Critic: DDPG algorithm
- End-end training

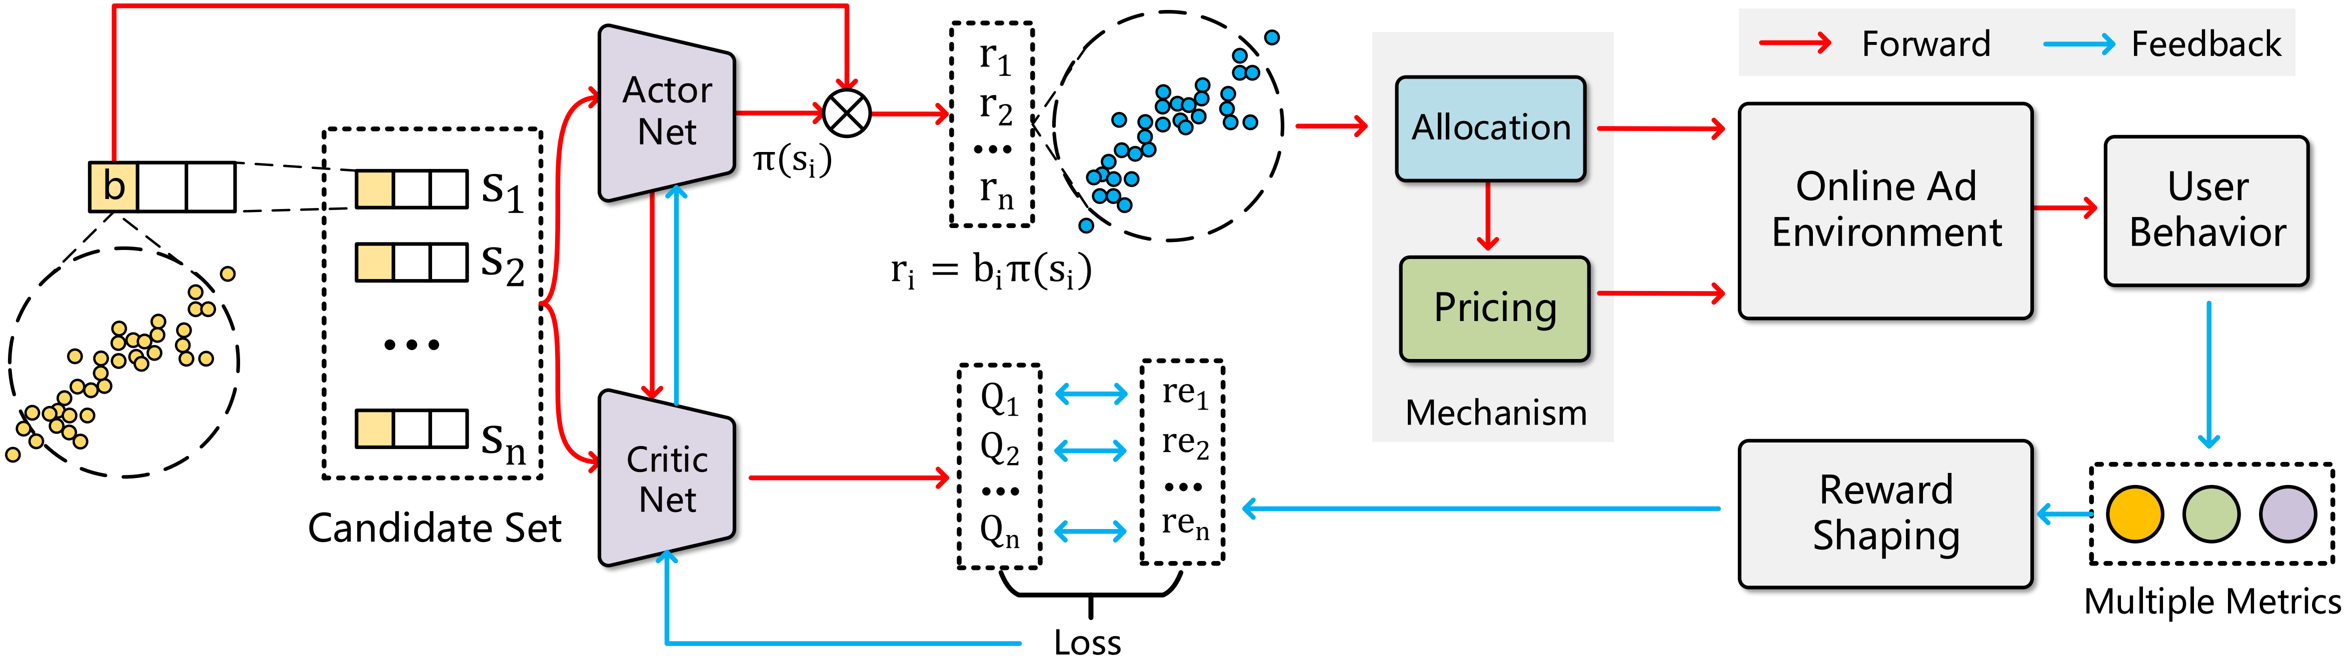

$$y_i = re_i \qquad \mathcal{L}(Q_\theta) = \frac{1}{N}\sum_i (y_i - Q_\theta(s_i, a_i))^2 \qquad \mathcal{L}(R_\theta) = \frac{1}{N}\sum_i (-Q_\theta(s_i, R_\theta(s_i)) + \gamma \times \mathcal{L}_{mono})$$

Experiments

Offline Experiments

- Data Set: 5870k records logged data from Taobao.
- Baselines: GSP, uGSP
- Performance metrics: RPM, CTR, ACR (Add-to-Cart Rate), CVR, GMV
- Intuitive comparisons: $\lambda \times RPM + (1 - \lambda) \times X$

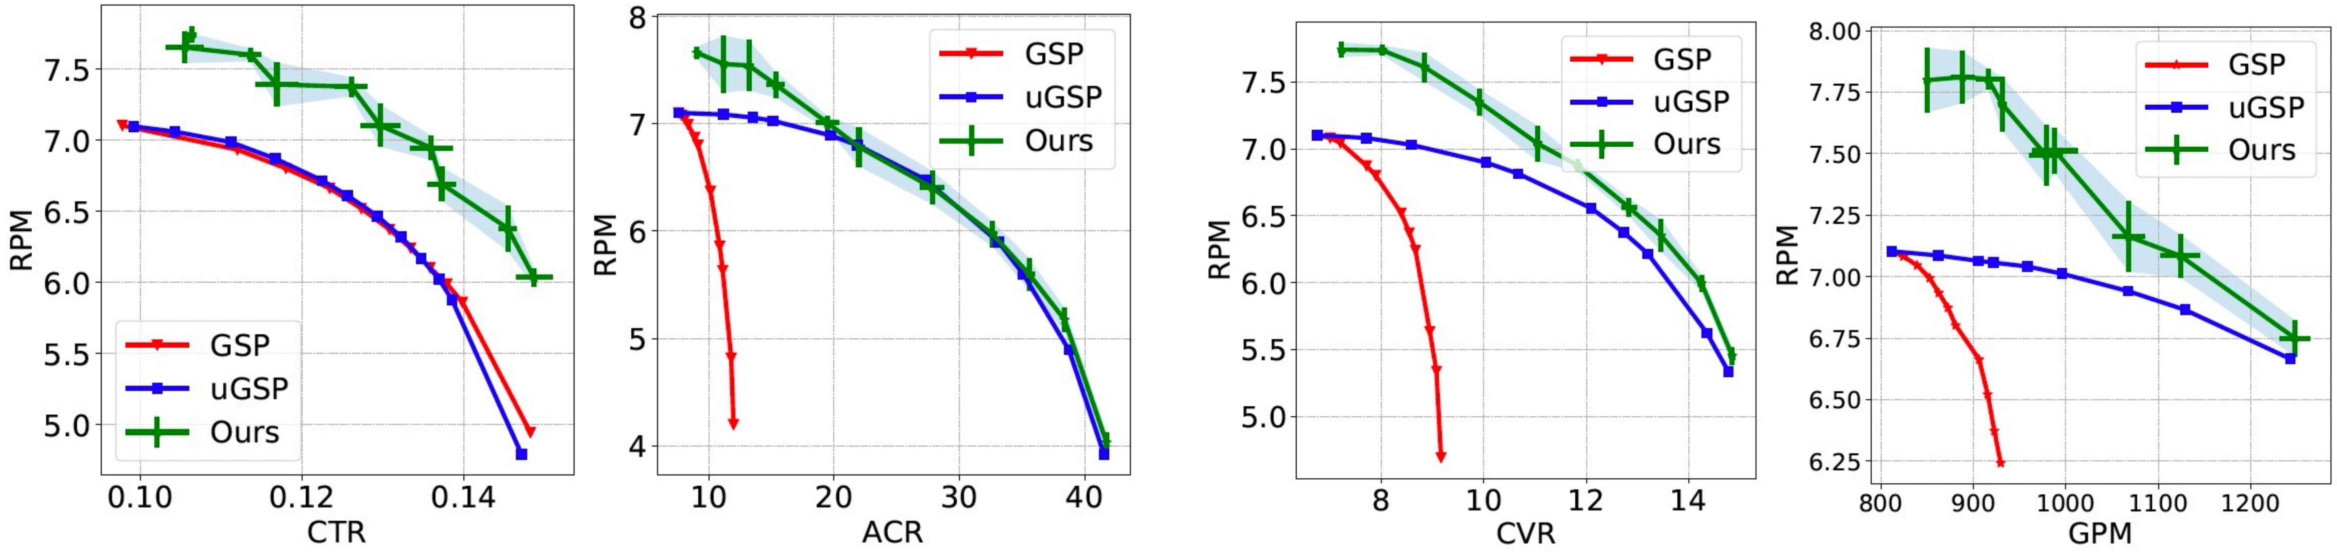

Online Experiments

Table 4: Online A/B test on different metrics configurations (August 1, 2020, 1% production flow).

Exp	Metrics	RPM	CTR	ACR	CVR	GPM
1	RPM	**+5.2%**	+3.1%	-1.5%	+0.8%	-2.0%
2	RPM&CTR	-0.3%	**+12.8%**	+5.6%	+20.0%	+7.5%
3	RPM&ACR	+0.7%	+1.5%	**+6.6%**	+6.8%	+8.1%
4	RPM&CVR	+0.0%	+1.4%	+3.6%	**+7.5%**	+31.0%
5	RPM&GPM	+0.2%	+3.3%	+2.4%	+3.6%	**+38.7%**
6	All	+1.8%	+6.2%	+1.4%	+5.9%	+3.7%

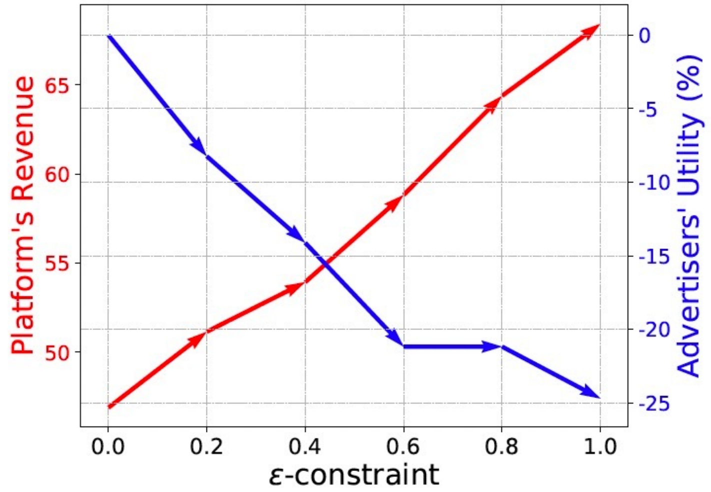

Figure 5: *Smooth transition* between mechanisms (from CTR to RPM) by increasing ϵ from 0.0 to 1.0.

Conclusion

Conclusions

✓We focus on the problem of optimizing multiple performance metrics in online e-commerce.

✓We leverage the deep learning technique to design a new rank score function and integrate it into the GSP auction framework, i.e., Deep GSP auction.

✓Both offline and online experimental results on a real-world e-commerce ad platform validate the effectiveness of the proposed auction mechanism.

Optimizing Multiple Performance Metrics with Deep GSP Auctions for E-commerce Advertising

in collaboration with