Learning to Bid with AuctionGym

Olivier Jeunen, Sean Murphy and Ben Allison
Amazon, Edinburgh, United Kingdom

August 15th 2022

jeunen@amazon.com

ACM SIGKDD Workshop on Computational Advertising (AdKDD ’22)

The Problem Setting

- **Ad impression opportunities** are sold off in real-time auctions.

The Problem Setting

- **Ad impression opportunities** are sold off in real-time auctions.
- **Goal:** participate in auctions, maximise some notion of *utility*.

The Problem Setting

- **Ad impression opportunities** are sold off in real-time auctions.
- **Goal:** participate in auctions, maximise some notion of *utility*.
- **Complication:** advertising auctions are rarely *incentive compatible*.

The Problem Setting

- **Ad impression opportunities** are sold off in real-time auctions.
- **Goal:** participate in auctions, maximise some notion of *utility*.
- **Complication:** advertising auctions are rarely *incentive compatible*.
- **Task:** learn a bidding strategy that maximises *utility*.

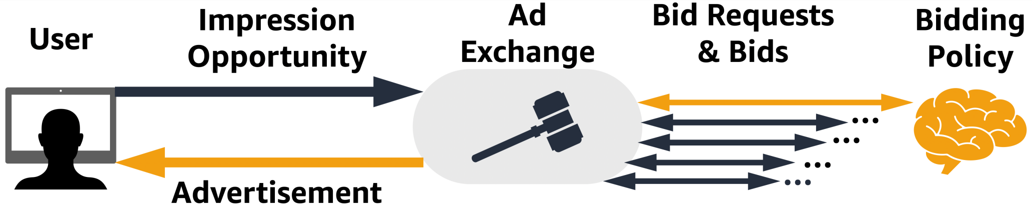

Figure 1: High-level overview of a real-time-bidding flow in computational advertising.

How do you *“learn to bid”*?

The Bidding Objective

We want to maximise *utility*:

"*Value minus price for impression opportunities we win*"

$$U = W(V - P) \tag{1}$$

V depends on the ad, W and P depend on the bid

The Bidding Objective

Goal

Maximise *utility*, given that we sample bids according to some contextual policy $\pi(B|A;X)$:

$$\mathop{\mathbb{E}}_{b \sim \pi(B|A;X)}[U] \tag{2}$$

How do you estimate this, using samples from π_0?

Optimising The Bidding Objective

Choosing a Counterfactual Estimator

1. **Value-based Estimation** (The *"Direct Method"*) High Bias
   (Existing work fits this paradigm — model P$(\mathrm{win}|\mathrm{bid})$)

Optimising The Bidding Objective

Choosing a Counterfactual Estimator

1. **Value-based Estimation** (The *"Direct Method"*) High Bias
   (Existing work fits this paradigm — model $\mathrm{P}(\text{win}|\text{bid})$)

2. **Policy-based Estimation** (*IPS*) High Variance
   (Novel)

Optimising The Bidding Objective

Choosing a Counterfactual Estimator

1. **Value-based Estimation** (The *"Direct Method"*) High Bias
   (Existing work fits this paradigm — model P($\mathrm{win}|\mathrm{bid}$))
2. **Policy-based Estimation** (*IPS*) High Variance
   (Novel)
3. **Doubly Robust Estimation** Unbiased, lower variance
   (Novel)

How do you evaluate this?

How do you evaluate this?

- **Offline**: use counterfactual estimators . . .

How do you evaluate this?

- **Offline**: use counterfactual estimators . . .
  - *"When a measure becomes a target, it ceases to be a good measure"* (Goodhart's Law)

How do you evaluate this?

- **Offline**: use counterfactual estimators . . .
  - *"When a measure becomes a target, it ceases to be a good measure"* (**Goodhart's Law**)
- **Online**: A/B-tests span weeks, require production-level prototypes, ...

How do you evaluate this?

- **Offline**: use counterfactual estimators ...
  - *"When a measure becomes a target, it ceases to be a good measure"* (**Goodhart's Law**)
- **Online**: A/B-tests span weeks, require production-level prototypes, ...
- The **RL** community is well aware of these issues: **embrace simulation**!

Introducing

AuctionGym

Simulating Advertising Auctions End-to-End

1. An **impression opportunity arises**, with features $x \sim \mathrm{P}(X)$,

Simulating Advertising Auctions End-to-End

1. An **impression opportunity arises**, with features $x \sim \mathrm{P}(X)$,
2. the **auctioneer presents** this opportunity to some **bidders**,

Simulating Advertising Auctions End-to-End

1. An **impression opportunity arises**, with features $x \sim \mathrm{P}(X)$,
2. the **auctioneer presents** this opportunity to some **bidders**,
3. **bidders** internally **decide** on an **ad** to show and a **bid** to place,

Simulating Advertising Auctions End-to-End

1. An **impression opportunity arises**, with features $x \sim P(X)$,
2. the **auctioneer presents** this opportunity to some **bidders**,
3. **bidders** internally **decide** on an **ad** to show and a **bid** to place,
4. the **auctioneer decides** on the auction **winner** and **price**,

Simulating Advertising Auctions End-to-End

1. An **impression opportunity arises**, with features $x \sim \mathrm{P}(X)$,
2. the **auctioneer presents** this opportunity to some **bidders**,
3. **bidders** internally **decide** on an **ad** to show and a **bid** to place,
4. the **auctioneer decides** on the auction **winner** and **price**,
5. the winning ad is shown and **possibly** leads to
   a **conversion event** that is observable by the **winning bidder**.

Simulating Advertising Auctions End-to-End

1. An **impression opportunity arises**, with features $x \sim P(X)$,
2. the **auctioneer presents** this opportunity to some **bidders**,
3. **bidders** internally **decide** on an **ad** to show and a **bid** to place,
4. the **auctioneer decides** on the auction **winner** and **price**,
5. the winning ad is shown and **possibly** leads to
   a **conversion event** that is observable by the **winning bidder**.

Bidders **update** their **models** every Δ_{r} auction **rounds**

Goals and Ambitions for AuctionGym

To be used as a research & validation tool

that does not rely on

sensitive, proprietary data.

. . . so which estimator do I use?

Core Research Question:

How does my learning method affect my profit?

(Assuming first-price auctions, details in the paper)

Simulated Auctions over Time

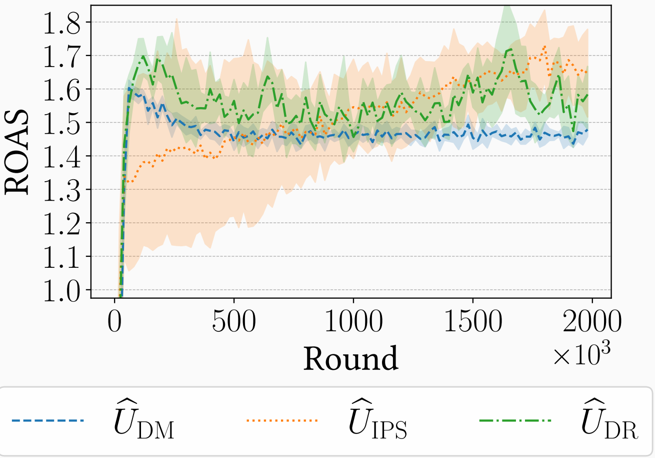

Simulated Auctions over Time

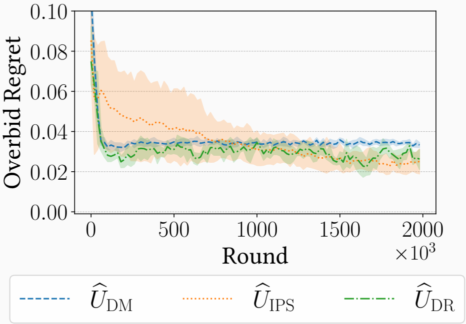

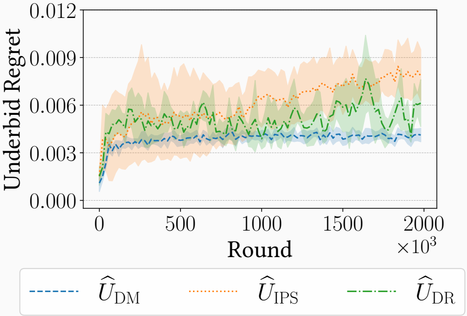

Contributions

Contributions

1. A **general framework** for bandit-based *“learning to bid”*

Contributions

1. A **general framework** for bandit-based *“learning to bid”*
2. Proposed a **novel approach**, leveraging **doubly robust** estimators

Contributions

1. A **general framework** for bandit-based *"learning to bid"*
2. Proposed a **novel approach**, leveraging **doubly robust** estimators
3. **AuctionGym**, a tool that can benefit **research** as well as **practitioners**

Contributions

1. A **general framework** for bandit-based *"learning to bid"*
2. Proposed a **novel approach**, leveraging **doubly robust** estimators
3. **AuctionGym**, a tool that can benefit **research** as well as **practitioners**
4. Insights from **AuctionGym** that we cannot extract from logged data.

Contributions

1. A **general framework** for bandit-based *"learning to bid"*
2. Proposed a **novel approach**, leveraging **doubly robust** estimators
3. **AuctionGym**, a tool that can benefit **research** as well as **practitioners**
4. Insights from **AuctionGym** that we cannot extract from logged data.

   (And much more in the paper!)

Thank you for listening!

Questions?

jeunen@amazon.com

amazon

AuctionGym