AdKDD
August 23, 2020

Contextual Bandits for Advertising Budget Allocation

Benjamin Han
Applied Machine Learning

Jared Gabor
Data Science

Driver Acquisition at Lyft

Lyft Advertising

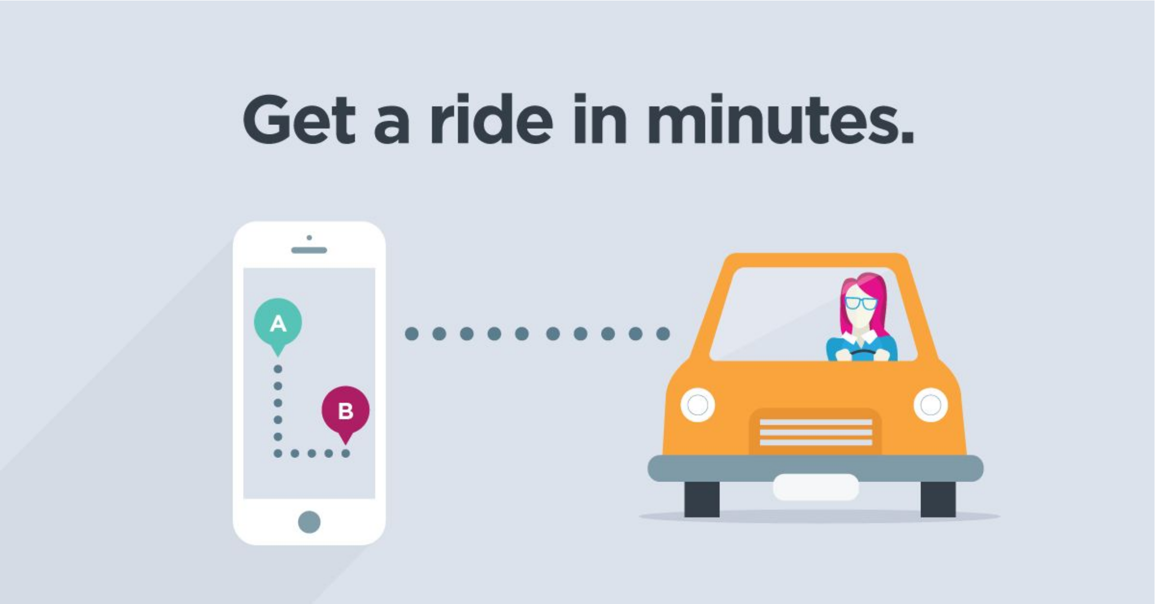

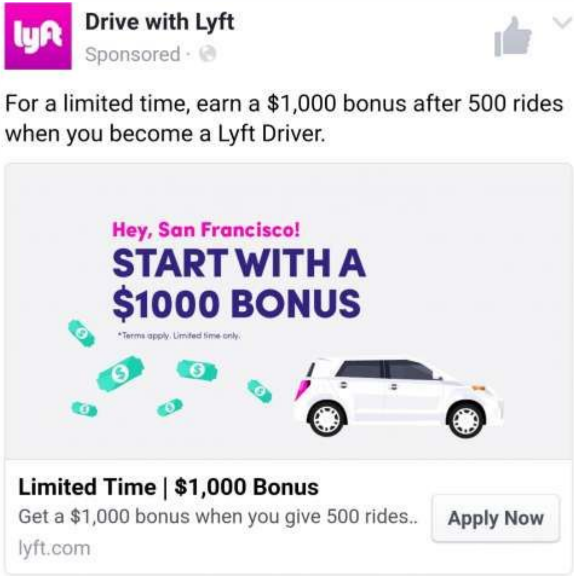

Han and Gabor. 2020. Contextual Bandits for Advertising Budget Allocation.

Objective

- Maximize total payouts f(x) over N campaigns given total budget B

- *Equivalently*, maximize total payouts given minimum **marginal ROI**, $\frac{df(x)}{dx}$

Maximum Budget Constraint

$$\begin{aligned} \underset{x \in \mathbb{R}}{\text{maximize}} \quad & \sum_{i=1}^{N} f_i(x_i) \\ \text{subject to} \quad & \sum_{i=1}^{N} x_i \leq B \\ & x_i \geq 0, \quad i \in \{1, \ldots, N\} \end{aligned}$$

Budget management

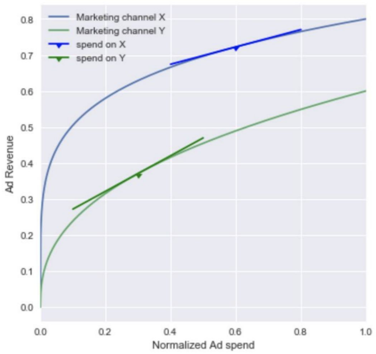

[D'Elia, AdKDD '19]

Budget allocation

Align marginal ROIs

How many Actions?

What is the revenue of the Ad campaign?

What should we predict?

Han and Gabor. 2020. Contextual Bandits for Advertising Budget Allocation.

Objective

- Maximize total payouts f(x) over N campaigns given total budget B
- *Equivalently*, maximize total payouts given minimum **marginal ROI**, $\frac{df(x)}{dx}$
- *Equivalently*, maximize total payouts given maximum **Cost Per Incremental Acquisition**, $\left(\frac{df_i(x_i)}{dx}\right)^{-1}$

Maximum Budget Constraint

$$\begin{aligned} \underset{x \in \mathbb{R}}{\text{maximize}} \quad & \sum_{i=1}^{N} f_i(x_i) \\ \text{subject to} \quad & \sum_{i=1}^{N} x_i \leq B \\ & x_i \geq 0, \quad i \in \{1, \ldots, N\} \end{aligned}$$

Maximum Cost Per Incremental Acquisition

$$\begin{aligned} \underset{x \in \mathbb{R}}{\text{maximize}} \quad & \sum_{i=1}^{N} f_i(x_i) \\ \text{subject to} \quad & \left(\frac{df_i(x_i)}{dx}\right)^{-1} \leq C, \quad && i \in \{1, \ldots, N\} \\ & x_i \geq 0, && i \in \{1, \ldots, N\} \end{aligned}$$

Han and Gabor. 2020. Contextual Bandits for Advertising Budget Allocation.

CPIA

$$\begin{aligned} \underset{x \in \mathbb{R}}{\text{maximize}} \quad & \sum_{i=1}^{N} f_i(x_i) \\ \text{subject to} \quad & \left(\frac{df_i(x_i)}{dx}\right)^{-1} \leq C, \qquad i \in \{1, \ldots, N\} \\ & x_i \geq 0, \qquad i \in \{1, \ldots, N\} \end{aligned} \longrightarrow \frac{df_i(x_i)}{dx} = C^{-1}, \quad i \in \{1, \ldots, N\}$$

- Value of Driver = (Net Profit per Ride)(Incremental Rides generated by Driver)
  - Profitable if value exceeds **Cost Per Incremental Acquisition (CPIA)**
- Regional CPIA Targets are set externally based on forecast supply/demand

Problem: Given historical campaign performance and regional CPIA targets, Spend Allocator must yield a daily budget for each campaign. Our goal is to **maximize total acquisition**.

Intuition: Strongly performing campaigns should get large budgets, weak campaigns should get reduced budgets and eventually removed.

What about new or inconsistent campaigns with high risk or uncertainty?

Han and Gabor. 2020. Contextual Bandits for Advertising Budget Allocation.

Spend Allocator

1. Contextual Payout Modeling (interpolation)
2. Bayesian Linear Regression (extrapolation)
3. Thompson Sampling (exploration/exploitation)

Contextual Payout Modeling (interpolation)

Use ad context and proposed budget to predict future acquisition

- Dominating feature is proposed budget (more spend, more acquisition). Also yesterday's acquisition, age of creative, target demographic, impressions/clicks/leads, incentive, day of week
- Global model trained on entire performance history, across regions and campaigns
- Excellent performance *near-sample*

Approach	Bias	MAE	MSE
MCMC	0.1575	0.2144	0.5881
RF	-0.0340	0.0788	0.0656
LGBM	-0.0764	0.1894	0.3463

Limitations

- Predictions are not differentiable
- Don't always extrapolate well (especially tree models)

$$min(y_i) \leq \hat{y}(\theta, x) \leq max(y_i)$$

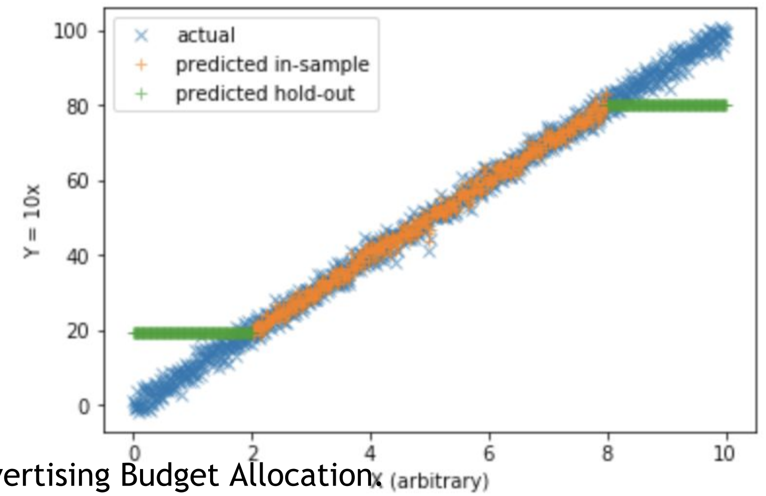

lyft Science

Han and Gabor. 2020. Contextual Bandits for Advertising Budget Allocation.

Desirable Payout Function Properties

Differentiable

Can be solved using the proposed portfolio maximization strategy

Monotonically Increasing

More allocated budget forecast more driver activations

Diminishing Returns

Cannot expect linear (or superlinear) driver activations as a function of budget. response rates after repeat exposures.

$$y = w_1 x^{w_2}$$
$$\log(y) = w_1 + w_2 \log(x)$$

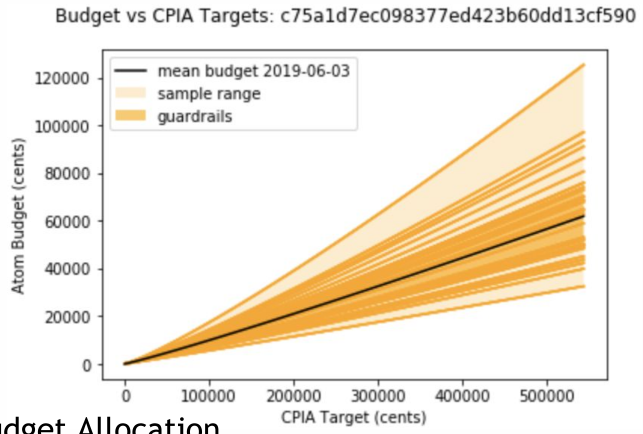

Han and Gabor. 2020. Contextual Bandits for Advertising Budget Allocation.

Bayesian Linear Regression (extrapolation)

Linear Regression using historical observations augmented by model point-predictions

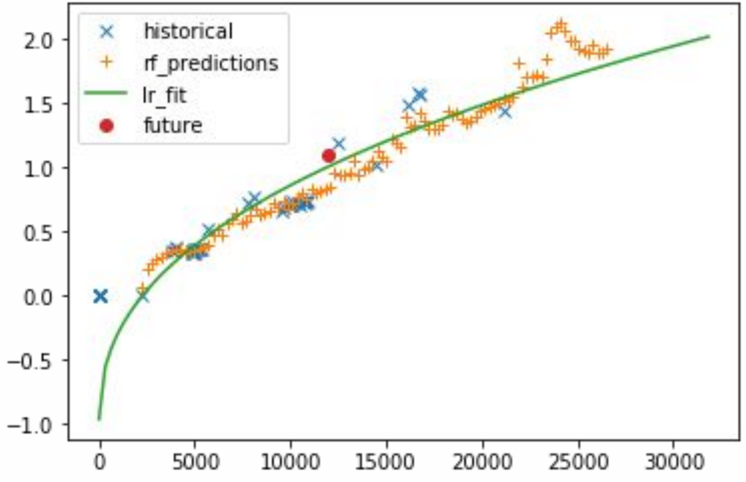

Predicted Driver Acquisition vs Daily Spend (cents)

Bayesian Linear Regression yields Curve Distribution with measure of uncertainty

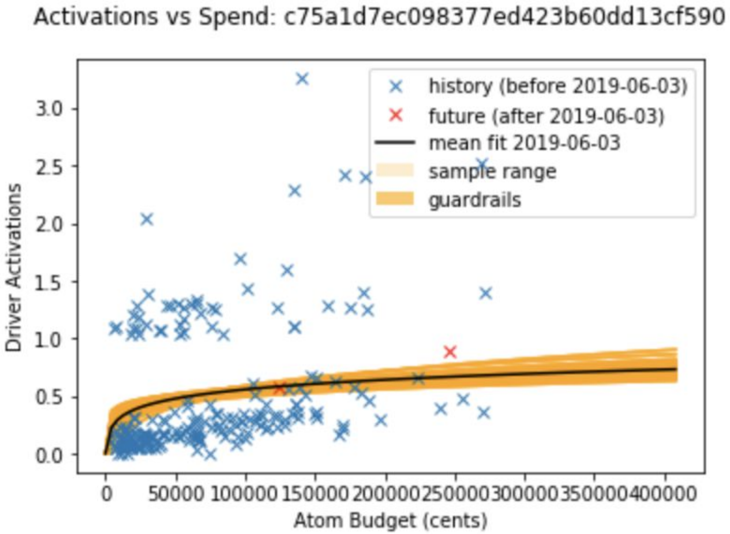

Han and Gabor. 2020. Contextual Bandits for Advertising Budget Allocation.

Thompson Sampling (exploration/exploitation)

What about new or inconsistent campaigns with high variance or risk?

Exploitation

Campaigns with established history are more predictable and can be priced accurately.

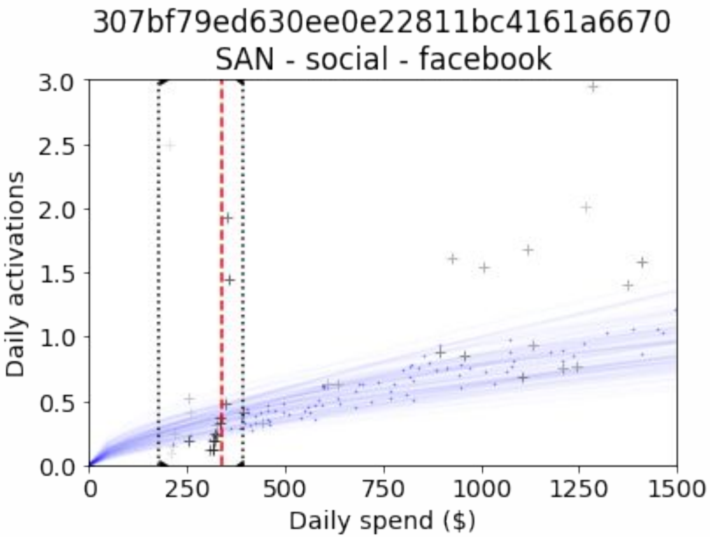

Exploration

Campaigns with inconsistent performance or with limited history have greater uncertainty.

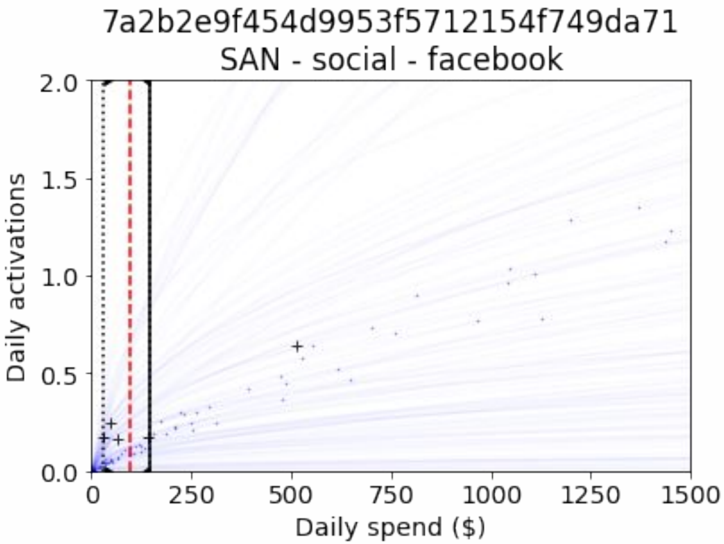

Han and Gabor. 2020. Contextual Bandits for Advertising Budget Allocation.

Experimental Results

22% CPA Savings, or 1.5¢ Profit per ride

- (21.8 +/- 10.2)% reduction in mean Cost Per (driver) Activation controlling for location, (21.5 +/- 13.1)% controlling for time

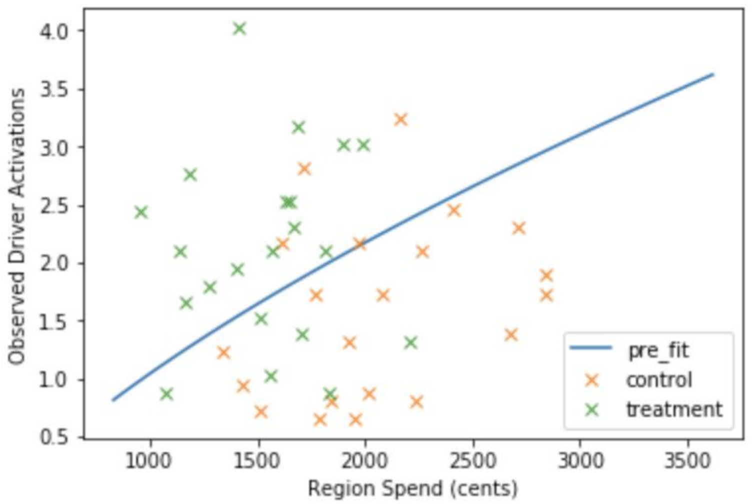


```
(Post Period / Pre Period) Spend Efficiency (higher is better):
        Mean Treatment Q: 0.8316 +/- 0.0654
        Mean Control Q: 0.6842 +/- 0.0503
Relative Improvement Q: 1.215408 +/- 0.130842
```

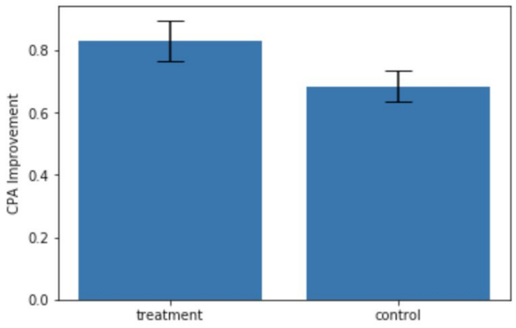

region		
BOS	1.137084	0.240171
CHI	1.325651	0.572299
DCA	1.099174	0.234972
DEN	0.988644	0.119454
PHI	1.155682	0.208775
PIT	1.790309	0.107267
SFO	1.030943	0.099320

```
Mean Regional improvement: 1.218212 +/- 0.102981
```


Han and Gabor. 2020. Contextual Bandits for Advertising Budget Allocation.

Summary

- Solve profitable budget allocation for known returns using Cost per Incremental Acquisition (CPIA)
- For unknown payouts, **RF_BLR** to explore uncertainty
  - Random Forest model provides accurate point estimates *near-sample*
  - Bayesian Linear Regression provides differentiable curve *far-from-sample*
  - Thompson Sampling explores budget-action space based on uncertainty
- Measured a 22% mean Cost Per Acquisition reduction, or ~$30 million annually, over **MCMC**
- Deployed globally, managing hundreds of millions of dollars annually

Han and Gabor. 2020. Contextual Bandits for Advertising Budget Allocation.

Thank you!

Lyft is hiring!

Benjamin Han
Applied Machine Learning
bhan@lyft.com
benjamin.han90@gmail.com

Please reach out with questions/comments/feedback