Learning-Augmented Algorithms for the Bahncard Problem

Hailiang Zhao [1] Xueyan Tang [2] Peng Chen [1] Shuiguang Deng [1]

[1]Zhejiang University

[2]Nanyang Technological University

Abstract

We study learning-augmented algorithms for the Bahncard problem. In the Bahncard problem, a traveler needs to irrevocably and repeatedly decide between a cheap short-term solution and an expensive long-term one with an unknown future. We develop a new learning-augmented algorithm, named PFSUM, that incorporates both history and short-term future to improve online decision making. We derive the competitive ratio of PFSUM as a function of the prediction error as follows:

$$\mathtt{CR}_{\mathsf{PFSUM}}(\eta) = \begin{cases} \frac{2\gamma+(2-\beta)\eta}{(1+\beta)\gamma+\beta\eta} & 0 \le \eta \le \gamma, \\ \frac{(3-\beta)\gamma+\eta}{(1+\beta)\gamma+\beta\eta} & \eta > \gamma. \end{cases} \quad (1)$$

We also conduct extensive experiments to show that PFSUM outperforms the SOTA algorithms.

Notations and Problem Formulation

BP(C, β, T) with $C > 0$, $T > 0$, and $0 \le \beta < 1$ is a request-answer game between an online algorithm ALG and an *adversary*. The adversary presents a finite sequence of travel requests $\sigma = \sigma_1\sigma_2\cdots$, where each σ_i is a tuple (t_i, p_i) that contains the travel time $t_i \ge 0$ and the regular ticket price $p_i \ge 0$.

ALG needs to react to each travel request σ_i. If ALG does not have a valid Bahncard, it can opt to buy the ticket with the regular price p_i, or first purchase a Bahncard which costs C, and then pay the ticket price with a β-discount, i.e., βp_i. A Bahncard purchased at time t is valid during the time interval $[t, t+T)$.

We say σ_i is a *reduced request* of ALG if ALG has a valid Bahncard at time t_i. Otherwise, σ_i is a *regular request* of ALG. We use ALG(σ_i) to denote ALG's cost on σ_i:

$$\mathsf{ALG}(\sigma_i) = \begin{cases} \beta p_i & \text{ALG has a valid Bahncard at } t_i, \\ p_i & \text{otherwise.} \end{cases}$$

We denote by ALG(σ) the total cost of ALG for reacting to all the travel requests in σ. The competitive ratio of ALG is defined by

$$\mathtt{CR}_{\mathsf{ALG}} := \max_{\sigma} \mathsf{ALG}(\sigma)/\mathsf{OPT}(\sigma),$$

where OPT is an offline optimum for BP(C, β, T). We use ALG$(\sigma; \mathcal{I})$ to denote the partial cost incurred during a time interval $\mathcal{I}$. Additionally, we use $c(\sigma; \mathcal{I})$ to denote the total regular cost in $\mathcal{I}$.

Given a time length l, we define the *l-recent-cost* of σ at time t as $c(\sigma; (t-l, t])$. Similarly, we define the *l-future-cost* of σ at time t as $c(\sigma; [t, t+l))$.

We further define the *regular l-recent-cost* of ALG on σ at time t as

$$\mathsf{ALG}^r\big(\sigma; (t-l, t]\big) := \sum_{i:\sigma_i \text{ is a regular request of ALG in } (t-l,t]} p_i.$$

Learning-Augmented Algorithms

To represent prediction errors, we use $\hat{c}(\sigma; [t, t+T))$ to denote the predicted total regular cost in $[t, t+T)$.

We take the competitive ratio of a learning-augmented online algorithm ALG as a function $\mathtt{CR}_{\mathsf{ALG}}(\eta)$ of the prediction error η. **ALG is δ-*consistent* if $\mathtt{CR}_{\mathsf{ALG}}(0) = \delta$, and ϑ-*robust* if $\mathtt{CR}_{\mathsf{ALG}}(\eta) \le \vartheta$ for all η.**

Lemma

For any time t, if $c\big(\sigma; [t, t+T)\big) \ge \gamma := C/(1-\beta)$, OPT has at least one reduced request in $[t, t+T)$. The same holds for the time interval $(t, t+T]$.

SUM

SUM purchases a Bahncard at a regular request (t, p) whenever its regular T-recent-cost at time t is at least γ, i.e., SUM$^r\big(\sigma; (t-T, t]\big) \ge \gamma$. SUM is the best deterministic online algorithm, whose competitive ratio is $2/(1+\beta)$.

FSUM

FSUM *purchases a Bahncard at a regular request (t, p) whenever the predicted T-future-cost at time t is at least γ, i.e., $\hat{c}\big(\sigma; [t, t+T)\big) \ge \gamma$.* FSUM is $2/(1+\beta)$-consistent, but its robustness is ∞.

PFSUM

FSUM fails to achieve any bounded robustness because it completely ignores the historical information in the Bahncard purchase condition. Thus, the worst case is that the actual ticket cost in the prediction window is close to 0, while the predictor forecasts that it exceeds γ, in which case hardly anything benefits from the Bahncard purchased.

We also note that SUM achieves a decent competitive ratio because a Bahncard is purchased only when the regular T-recent-cost is at least γ.

Motivated by this observation, we introduce a new algorithm PFSUM (Past and Future SUM), in which the Bahncard purchase condition incorporates the ticket costs in both a past time interval and a future prediction window, ***but uses them separately rather than taking their sum***. PFSUM purchases a Bahncard at a regular request (t, p) whenever

1. the T-recent-cost at t is at least γ, i.e., $c(\sigma; (t-T, t]) \ge \gamma$, and
2. the predicted T-future-cost at t is also at least γ, i.e., $\hat{c}(\sigma; [t, t+T)) \ge \gamma$.

Note that PFSUM considers the T-recent-cost, but SUM considers only the regular T-recent-cost.

Competitive Analysis for PFSUM

We adopt a *divide-and-conquer* approach to analyze PFSUM. We focus on the time intervals in which at least one of PFSUM and OPT has a valid Bahncard, and analyze the cost ratio between PFSUM and OPT in these intervals.

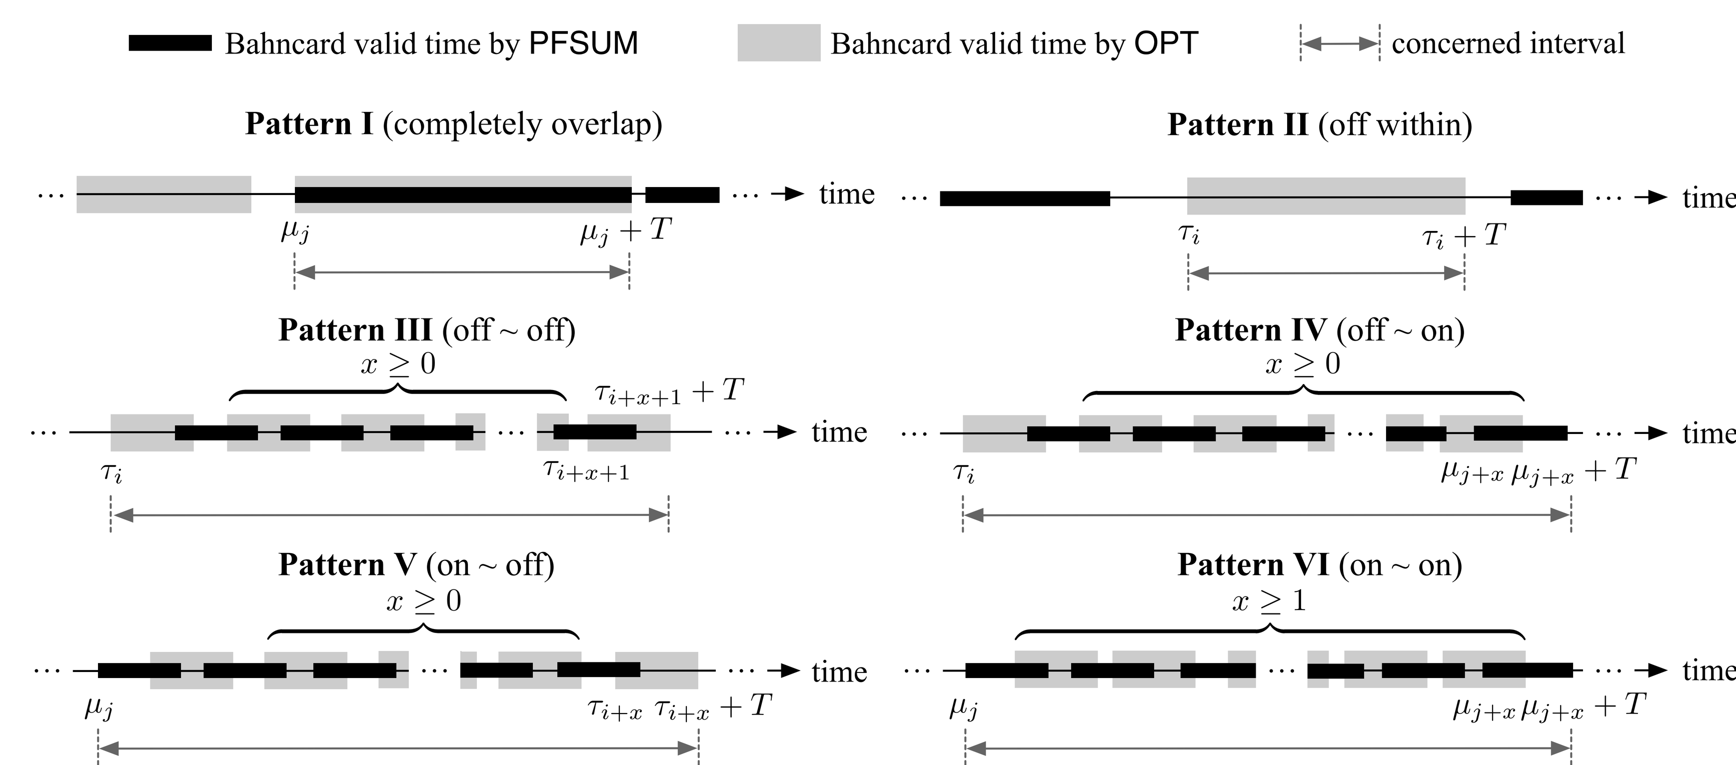

Figure 1. All the 6 patterns of concerned time intervals in which either PFSUM or OPT has a Bahncard. In Patterns III to VI, x is the number of Bahncards purchased by OPT in an on phase and expiring in the next on phase. x can be any non-negative integers.

We define η as the maximum prediction error among all the predictions:

$$\eta := \max_{(t,p) \text{ is a regular request}} \left| \hat{c}\big(\sigma; [t, t+T)\big) - c\big(\sigma; [t, t+T)\big) \right|.$$

The competitive ratio of PFSUM with any prediction errors is shown in (1).

Experimental Results

We model the inter-request time of occasional travelers using an exponential distribution with a mean of 2 days. The predictions are generated by adding noise to the original instance.

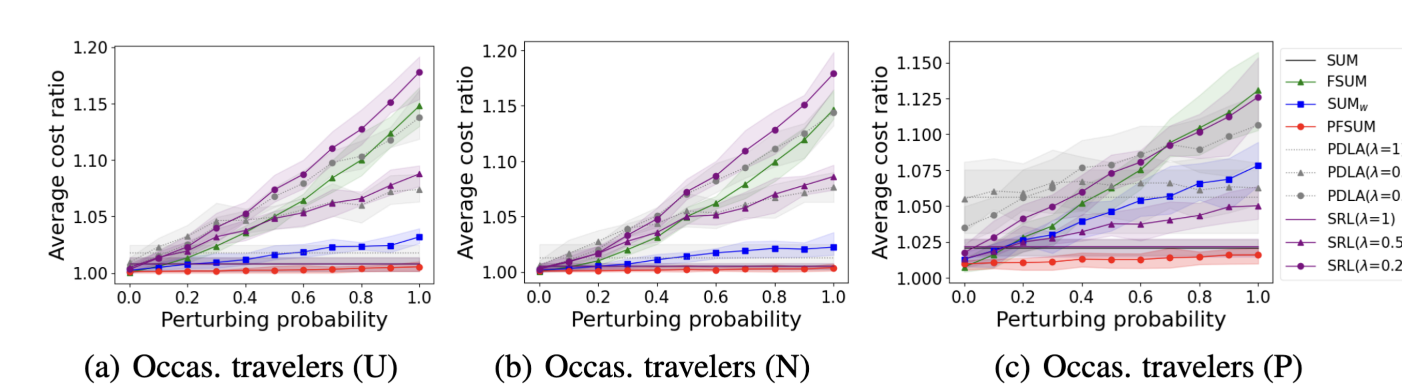

Figure 2. The cost ratios for commuters ($\beta = 0.8$, $T = 10$, $C = 100$). "U", "N" and "P" represent Uniform, Normal and Pareto ticket price distributions respectively.

PFSUM outperforms all the baselines in both consistency and robustness.

 hliangzhao@zju.edu.cn