Confidence-Aware Learning-Augmented Algorithms for the Bahncard Problem

Ziyi Han
The Chinese University of Hong Kong

Bo Sun
University of Ottawa

Xuchuang Wang
The Chinese University of Hong Kong

Mohammad Hajiesmai
University of Massachusetts Amherst

John C.S. Lui
The Chinese University of Hong Kong

ABSTRACT

Learning-augmented algorithms integrate predictions into online decision-making while maintaining worst-case guarantees. In this work, we revisit the learning-augmented algorithm design for the classical Bahncard problem, and propose a confidence-aware algorithm that can dynamically adjust its trust on predictions. We first identify subtle issues in existing analysis, including restrictive assumptions on interval dominance and missing edge cases in the interaction between the online algorithm and the optimal solution. To address these issues, we develop a refined pattern-based analysis and derive improved competitive-ratio bounds that explicitly capture the trade-off between prediction accuracy and robustness. Our results demonstrate that incorporating confidence into prediction-driven decisions leads to strictly better performance in favorable regimes, while preserving strong worst-case guarantees.

1. INTRODUCTION

The Bahncard problem [2] originates from a real-world pricing scheme used by the German railway system. A Bahncard is a discount pass that grants its holder a substantial reduction on train ticket prices for a fixed period, at the cost of an upfront payment. For a typical traveler, deciding when to purchase such a pass is inherently an online decision problem. While buying the Bahncard reduces future travel costs, the traveler does not know in advance how often he will travel or how much each trip will cost. Abstracting this scenario, the Bahncard problem captures a fundamental trade-off between an upfront investment and uncertain future benefits. The objective is to minimize the total cost over time, balancing immediate expenditures against potential long-term savings under incomplete information. This problem generalizes the classical ski-rental problem and captures a richer class of settings in which the benefit of a purchase extends over a finite time horizon. Beyond transportation, similar trade-offs arise in a variety of problems in modern systems, such as cloud resource reservation and subscription-based services, and energy management.

Recently, learning-augmented algorithms have emerged as a framework for incorporating predictive information into online decision-making. They aim to improve performance under typical inputs while maintaining provable worst-case guarantees. Bamas et al. [1] propose a primal-dual learning-augmented algorithm for the Bahncard problem, assuming access to a prediction of the optimal solution in the form of a complete set of purchase times. However, such predictions effectively require full knowledge of future requests, which is generally unavailable in practice and thus limits the applicability of the approach. Recent work [3] considers a more practical prediction model, where the algorithm is provided with short-term predictions of future cost within a fixed time window. By combining historical observations with predicted future costs, it achieves improved trade-offs between consistency and robustness. However, there exists some subtle issues in its current analysis. Particularly, it relies on implicit assumptions about the structure of certain intervals when comparing the online algorithm with the optimal solution. We prove that these assumptions do not always hold, and that there exist edge cases in which the competitive ratio deviates from the claimed bounds. Beyond analysis issues, the algorithm effectively treats the predictions as fully reliable, without accounting for their reliability or trust level. Incorporating such information, when available, can lead to more adaptive decisions and improved performance.

In this paper, we revisit the Bahncard problem in the learning-augmented setting. We consider a general model with short-term predictions of future demand and develop a confidence-aware algorithm that adapts its decision strategy according to the reliability of predictions. We further introduce a refined pattern-based analysis that resolves gaps in prior work and yields competitive-ratio guarantees under arbitrary prediction error.

2. MAIN RESULT

Problem Formulation. The Bahncard problem is defined by three parameters (C, β, T), where C denotes the cost of purchasing a Bahncard, $\beta \in (0, 1)$ is the discount factor applied to ticket prices, and T is the validity period of the card. Travel demands arrive over time randomly, where each demand is associated with an arrival time t and a ticket price p. If no Bahncard is active at time t, the demand is served at its full price p; otherwise, a discounted price βp is charged. The objective is to minimize the total cost incurred over time, including both payments for travel demands and the cost of purchasing Bahncards.

A central quantity in this problem is the break-even threshold $\gamma := C/(1-\beta)$, which characterizes that purchasing a Bahncard at time t is beneficial if the total (full) price of demands over the interval $[t, t+T)$ exceeds γ.

Prediction Model. We consider a learning-augmented

Learning-Augmented Algorithms: Theory and Applications (LATA) 2026
Ann Arbor, MI, USA

setting where, for each travel demand arriving at time t, the algorithm is given a prediction of the total (regular) ticket price over the future interval $[t, t+T)$. Let η denote the maximum prediction error, defined as the largest deviation between the predicted and actual total price over any such interval. In addition, we introduce a trust parameter $\lambda \in (0,1]$ to capture the reliability of the predictions. Specifically, λ reflects the level of confidence in the predicted values: smaller values correspond to higher trust in the predictions.

Algorithm Design. To motivate our design, we first recall two representative approaches. Fleischer [2] first studied the Bahncard problem and proposed an optimal $(2-\beta)$-competitive deterministic algorithm, referred to as SUM. The SUM algorithm purchases a Bahncard only when the total regular price accumulated over the past interval $(t-T, t]$ reaches the break-even threshold γ. This strategy is robust, as it relies solely on observed costs, and its theoretical guarantee is based on the fact that the cost of each Bahncard can be charged to the accumulated ticket price in the competitive analysis. However, it does not exploit any information about future demand. Building on this idea, PFSUM incorporates predictive information by requiring that both the accumulated past price and the predicted future price are sufficiently large before making a purchase decision. This approach achieves both consistency and robustness guarantees. However, it treats the prediction in a binary manner and does not adapt to varying levels of prediction reliability.

To address this limitation, we propose a confidence-aware algorithm that explicitly accounts for prediction reliability. At each decision point, the algorithm first checks whether the predicted future price exceeds the break-even threshold γ. Based on this outcome, it adaptively adjusts the required threshold on the past accumulated price. Specifically, when the predicted future price is large, i.e., at least γ, the algorithm adopts a more aggressive strategy and purchases a Bahncard once the past price exceeds a reduced threshold $\lambda\gamma$. In contrast, when the predicted future price is small, the algorithm becomes more conservative and requires the past price to exceed an inflated threshold γ/λ before purchasing. The confidence parameter λ controls this asymmetric adjustment. Smaller values of λ lower the threshold $\lambda\gamma$ and increase the threshold γ/λ, leading to stronger reliance on predictions. Larger values of λ reduce this gap and recover a more conservative behavior closer to the classical SUM algorithm. This adaptive thresholding mechanism allows the algorithm to smoothly interpolate between prediction-driven decisions and robust worst-case behavior, thereby achieving improved performance when predictions are accurate while maintaining strong guarantees under uncertainty.

Theoretical Bound. Our proposed algorithm can attain the following trade-off between consistency and robustness.

Theorem 1. *Given a confidence parameter $\lambda \in (0,1]$, the learning-augmented algorithm for Bahncard problem is $\frac{2-\beta+\lambda}{1+\beta\lambda}$-consistent and $\max\left\{\frac{4\lambda-3\beta\lambda+1}{\lambda+\beta}, \frac{\lambda+1-\beta}{\lambda}\right\}$-robust.*

To prove this result, we partition the entire time horizon into maximal contiguous intervals during which at least one of our algorithm and the optimal solution holds a valid Bahncard. Within each such interval, we analyze the interaction between the two solutions.

As illustrated in Fig. 1, there are six possible overlap patterns, where each pattern is characterized by the relative

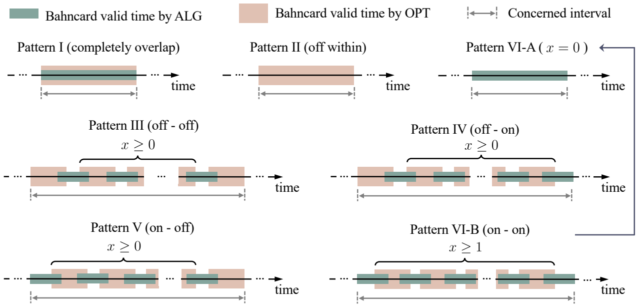

Figure 1: All patterns of concerned time intervals in which either our algorithm or the optimal solution has a Bahncard.

positions of the on-phases of the two solutions, and x denotes the number of Bahncards purchased by the optimal solution within the interval. We analyze each pattern separately and then combine the resulting bounds to derive the overall competitive ratio of our algorithm:

$$CR_{ALG} = \begin{cases} \max\{\min\{\frac{4\lambda-3\beta\lambda+1}{\lambda+\beta}, \frac{(4-3\beta+\lambda)\gamma+\eta}{(1+\beta\lambda)\gamma+\beta\eta}\}, \frac{\lambda+1-\beta}{\lambda}\}, \\ \qquad\qquad \text{if } \gamma N_B > \sum_{\text{on phase}} c, \\ \min\{\frac{3\lambda-2\beta\lambda+1}{\lambda+\beta}, \frac{(3-2\beta+\lambda)\gamma+\eta}{(1+\beta\lambda)\gamma+\beta\eta}\}, \\ \qquad\qquad \text{if } \gamma N_{A+B} > \sum_{\text{on phase}} c, \\ \min\{\frac{2\lambda-\beta\lambda+1}{\lambda+\beta}, \frac{(2-\beta+\lambda)\gamma+\eta}{(1+\beta\lambda)\gamma+\beta\eta}\}, \quad \text{else.} \end{cases} \tag{1}$$

where $N_B = N_{VI-B} - N_I - N_{II} - N_{III} - N_{IV} - N_V$ and $N_{A+B} = N_{VI-A} + N_{VI-B} - N_I - N_{II} - N_{III}$. Here, N_i denotes the number of occurrences of Pattern i in the decomposition.

3. REFERENCES

[1] E. Bamas, A. Maggiori, and O. Svensson. The primal-dual method for learning augmented algorithms. *Advances in Neural Information Processing Systems*, 33:20083–20094, 2020.

[2] R. Fleischer. On the bahncard problem. *Theoretical Computer Science*, 268(1):161–174, 2001.

[3] H. Zhao, X. Tang, P. Chen, and S. Deng. Learning-augmented algorithms for the bahncard problem. *Advances in Neural Information Processing Systems*, 37:113234–113281, 2024.