

Simulation-Driven Railway Delay Prediction: An Imitation Learning Approach

Clément Elliker¹, Jesse Read¹, Sonia Vanier¹, Albert Bifet^{2,3}

¹LIX (École Polytechnique, IP Paris, CNRS)

²AI Institute, University of Waikato

³LTCI, Télécom Paris, IP Paris

clement.elliker@polytechnique.edu, jesse.read@polytechnique.edu, sonia.vanier@polytechnique.edu,
albert.bifet@telecom-paris.fr

Abstract

Reliable prediction of train delays is essential for enhancing the robustness and efficiency of railway transportation systems. In this work, we reframe delay forecasting as a stochastic simulation task, modeling state-transition dynamics through imitation learning. We introduce Drift-Corrected Imitation Learning (DCIL), a novel self-supervised algorithm that extends DAgger by incorporating distance-based drift correction, thereby mitigating covariate shift during rollouts without requiring access to an external oracle or adversarial schemes. Our approach synthesizes the dynamical fidelity of event-driven models with the representational capacity of data-driven methods, enabling uncertainty-aware forecasting via Monte Carlo simulation. We evaluate DCIL using a comprehensive real-world dataset from INFRABEL, the Belgian railway infrastructure manager, which encompasses over three million train movements. Our results, focused on predictions up to 30 minutes ahead, demonstrate superior predictive performance of DCIL over traditional regression models and behavioral cloning on deep learning architectures, highlighting its effectiveness in capturing the sequential and uncertain nature of delay propagation in large-scale networks.

Code — <https://github.com/orailix/rail-delay-simulator>

1 Introduction

Railway networks are vital infrastructure supporting sustainable, large-scale mobility globally, facilitating billions of passenger journeys annually. The extensive reliance on rail transport directly reflects service quality; hence, transport providers prioritize reliable, efficient, and user-friendly operations to fulfill passenger expectations. As a result, accurate delay prediction has become a critical research area, enabling commuters to anticipate disruptions and allowing operational personnel to proactively manage service impacts.

Following (Rößler et al. 2021), we distinguish primary delays, which stem from operational problems such as rolling stock failures, signal malfunctions, or severe weather, from secondary delays, the knock-on effects that propagate these initial disruptions as late-running trains interfere with downstream slot allocations. Due to the nature of accessible data, the delay prediction literature focuses mostly on modeling

secondary delays, whose dynamics are governed by intricate spatiotemporal dependencies among interconnected services and infrastructure elements.

As described in (Spanninger et al. 2022), these approaches can be grouped into two categories. The first, event-driven approaches, capture the interdependence between arrival and departure events using stochastic models (Graph Models (Goverde 2010), Markov Chains (Şahin 2017)), offering interpretability, uncertainty quantification and modest data requirements. The second, data-driven approaches, cast delay prediction as a supervised regression problem, allowing them to learn complex traffic dynamics from historical data using machine learning models (Linear Regression and Tree-based methods (Kecman and Goverde 2015), Neural Networks (Oneto et al. 2018), Transformers (Arthaud, Lecoeur, and Pierre 2024)). Yet event-driven dynamical models insufficiently represent complex interactions, while data-driven one-shot regressors may undervalue the inherent temporal dependencies of successive events.

To leverage the advantages of both methodologies, we propose to frame delay prediction as a stochastic simulation problem: a policy is trained by imitation learning to reproduce the state-transition dynamics $p(s_{t+1}|s_t)$ observed in historical data. During roll-out, the policy stochastically predicts following states, enabling uncertainty quantification through Monte Carlo simulation. We combine the sequential nature of event-driven models with the representational power of data-driven methods to capture traffic complexity.

We propose *Drift-Corrected Imitation Learning* (DCIL), a self-supervised extension of DAgger’s dataset-aggregation approach (Ross, Gordon, and Bagnell 2011). Rather than relying on an expensive external oracle, DCIL applies distance-based drift correction during roll-outs by evaluating each candidate action’s induced next-state and choosing the one that minimizes a distance $\psi(\cdot, \cdot)$ back toward the expert’s subsequent state. This strategy mitigates covariate shift without resorting to complex adversarial (Ho and Ermon 2016) or inverse reinforcement learning (Ng, Russell et al. 2000) frameworks, outperforming Behavioral Cloning (Torabi, Warnell, and Stone 2018).

Our evaluation uses an extensive real-world dataset derived from publicly accessible operational logs from INFRABEL, the Belgian railway infrastructure manager. This dataset encompasses over three million train operations

across three years, covering 682 stations and diverse service types (regional, intercity, high-speed) under various conditions (peak, off-peak, weekday, weekend, disruptions). We benchmark simulation-based methods against regression across multiple architectures. Strikingly, we find that a 1.4-million-parameter Multi-Layer Perceptron trained with DCIL outperforms a 19-million-parameter Transformer trained with conventional regression on our task.

Our key contributions to delay forecasting are threefold:

1. We model delay prediction in a stochastic simulation framework, focusing learning on short-term dynamics;
2. We propose DCIL, a novel self-supervised imitation learning algorithm that effectively addresses covariate shift through drift correction;
3. We conduct an extensive evaluation on a large-scale, real-world open dataset for predictions up to 30 minutes ahead, demonstrating substantial improvements in predictive accuracy over regression approaches.

2 Problem Definition

Let train i be present in the rail network at time t . We define its feature vector $s_t^{(i)} \in \mathbb{R}^d$ where d depends on the chosen feature-encoding configuration, and $s_t^{(i)}$ contains the following components:

- train type;
- past and scheduled future stations and lines;
- station roles (*departure, arrival, passage*);
- theoretical passage times at stations;
- realised delays at past stations;
- temporal context (hour of day, day of week).

Stations are embedded using the first eight non-trivial eigenvectors of the normalised graph Laplacian of the rail-network graph (Belkin and Niyogi 2003), then per-node (row-wise) L2-normalised to yield eight-dimensional spectral coordinates that compactly encode topology-preserving neighborhoods. Line embeddings are obtained by averaging the embeddings of their constituent stations.

Let n denote the number of trains present at time t . Then, we define the full network **state** at time t as

$$s_t = (s_t^{(i)})_{i=1}^n$$

i.e. all per-train states at time t .

Let the **itinerary** of train i be defined as a sequence of m stations (l_1, \dots, l_m) . We denote the *scheduled* passage time at station l_j as $\tau_{l_j}^{(i)} \in \mathbb{R}$ and the *actual* passage time by $\hat{\tau}_{l_j}^{(i)} \in \mathbb{R}$. The resulting delay is

$$d_{l_j}^{(i)} = \hat{\tau}_{l_j}^{(i)} - \tau_{l_j}^{(i)}, \quad d_{l_j}^{(i)} \in \mathbb{R},$$

so $d_{l_j}^{(i)} > 0$ indicates lateness, $d_{l_j}^{(i)} = 0$ on-time operation, and $d_{l_j}^{(i)} < 0$ earliness. For stations not yet visited, $d_{l_j}^{(i)}$ constitutes the quantity we want to predict.

Given the current network state s_t , our goal is, for every train i and every future station l , to model the conditional distribution

$$p(d_l^{(i)} \mid s_t), \quad (1)$$

or, in a point-prediction setting, some aspect of this distribution, such as its mean or median.

3 Background and Related Work

Regression for train delay prediction

As noted in the introduction, train-delay research divides into two broad streams: event-driven approaches, which model the interdependence between arrival and departure events, and data-driven approaches, which learn predictive models directly from observations (Spanninger et al. 2022).

Recent years have seen a surge in the latter, with studies exploring deep-learning architectures—including Convolutional Neural Networks and bidirectional LSTMs (Guo et al. 2022), Graph Neural Networks (Heglund et al. 2020), and Transformers (Arthaud, Lecoeur, and Pierre 2024). Although these models outperform simple baselines they are compared to, they are evaluated on private datasets, so results cannot be compared across studies.

To our knowledge, (Yang et al. 2024) presents the first large-scale head-to-head comparison of data-driven methods for train-delay prediction. They find that transformer-based architectures perform best on an open-source dataset covering two Chinese high-speed routes.

Datasets Only a handful of refereed papers rely on *public* delay data. The best-curated example is the HSR-DELAY corpus (3399 trains, 727 stations, Oct 2019–Jan 2020) released in *Scientific Data* (Zhang et al. 2022). However, the dataset spans only four months on a limited number of trains, and the follow-up study of (Yang et al. 2024) evaluates just two lines. (Dekker et al. 2022) model delay using diffusion on the Belgian open-data feed (same dataset as our paper), but do not predict delays for individual trains; instead, they simulate the evolution of delay across clusters of stations over time. (Lapamonpinyo, Derrible, and Corman 2022) scrape the U.S. Amtrak API for a single corridor. To our knowledge, no train delay prediction study yet covers an entire *mixed-traffic* national network over multiple years, leaving generalisation beyond high-speed or single-corridor settings largely untested.

Simulation-based prediction

Past work on railway network simulation relies on hand-crafted microscopic, mesoscopic, and macroscopic simulators such as RailSys (Bendfeldt, Mohr, and Muller 2000), OpenTrack (Nash and Huerlimann 2004), PETER (Koelemeijer et al. 2000), or PROTON (Sipilä 2023) that encode operating rules and dispatcher heuristics instead of learning from historical data. Macroscopic simulators neglect precise train dynamics but scale well to large networks, whereas microscopic ones demand extensive calibration and significant compute, and mesoscopic models target only critical sections (Tiong and Palmqvist 2023). Consequently, forecasts

from these rule-driven tools often fail to match real-world delay patterns, especially on busy networks.

The work on rail-network simulation reflects that of the wider time-series forecasting literature. Train delay is deterministically linked to train trajectory (i.e., delay at t can be derived from state s_t), with trajectory prediction being a specific case of time-series forecasting (cf. Eq. (1)).

A traditional approach to time series forecasting, often used in the context of signal processing, is to implement a dynamical model $p(s_{t+1} | s_t)$ whose parameters have a physical meaning (speed, acceleration, etc.), e.g. (Prevost, Desbiens, and Gagnon 2007; Martino et al. 2017).

Specifying reliable dynamical models with an expert is not always feasible (particularly with complicating external factors such as human behaviour involved), but such models can instead be learned in a data-driven (and often, model-agnostic) approach via machine-learning; essentially learning $p(s_{t+1} | s_t)$ or similar from training pairs $\{(s_t, s_{t+1})\}_t$.

Data-driven simulation approaches have shown their efficiency in predicting future events in various domains. In weather forecasting, (Price et al. 2025) introduced GenCast, a conditional diffusion model trained on decades of reanalysis data that generates 15-day global ensemble forecasts in minutes and surpasses the leading physics-based system (ECMWF-ENS) on 97% of evaluation targets. In road-traffic modeling, (Kuefler et al. 2017) applied Generative Adversarial Imitation Learning to train a driver-behavior simulator that reproduces realistic lane changes, speed profiles and collision-free trajectories, showing that imitation-learning-based simulation can successfully predict future traffic events.

Imitation Learning

Imitation learning seeks a policy π_θ that reproduces expert behaviour from demonstration trajectories $\tau = (s_0, a_0, \dots, s_T, a_T)$. Four families dominate the literature:

Behavioral Cloning (BC). Treats imitation as supervised learning: fits $\pi_\theta(s)$ to expert actions a on recorded (s, a) pairs. Fast and data-efficient but prone to covariate shift once the learner visits states absent from the dataset.

Dataset Aggregation (Dagger). Iteratively executes the current policy, queries the expert on the visited states, and augments the training set. This feedback loop curbs covariate shift but can be expensive when expert labels are costly.

Inverse Reinforcement Learning (IRL). Alternates between (i) fitting a reward function under which the expert is (near-)optimal and (ii) training the policy via RL using learned rewards. This bi-level loop generalises beyond the demonstration manifold but is computationally heavy.

Generative Adversarial Imitation Learning (GAIL). Derived from IRL, GAIL sets up a GAN-style game: a discriminator distinguishes expert from learner state-action pairs, and its negative log-probability serves as the reward fed to the policy optimiser (e.g. TRPO/PPO). This removes the expensive bi-level loop, but inherits typical GAN instabilities such as mode collapse.




Figure 1: Illustration of actions

4 Delay Prediction via Macroscopic Simulation

We frame the railway network as a Markov Decision Process (MDP). Let $s_t = (s_t^{(1)}, \dots, s_t^{(n)})$ be the state of the rail network at time t , defined in Section 2. Let \mathcal{A} denote the action space, with $a_t = (a_t^{(1)}, \dots, a_t^{(n)})$ the actions describing the movements of trains on the network at time t . As illustrated in Figure 1, action $a^{(i)} \in \{0, 1, 2\}$ moves train i by that many scheduled stations during the next time step, with the cap at 2 ensuring a finite discrete action space. In this work, we set $\Delta t = 30$ seconds, so s_{t+1} represents a full *snapshot* of the network 30 seconds after s_t . No hand-crafted block-constraint rules are enforced; instead, the policy approximates them implicitly by imitation, as a precise specification is impractical. Actions and states remain discrete to fit available data, a continuous version would require GPS.

The environment dynamics, applied per train, are defined by the function $\phi : \mathcal{S}^{(i)} \times \mathcal{A}^{(i)} \rightarrow \mathcal{S}^{(i)}$, so that $s_{t+1}^{(i)} = \phi(s_t^{(i)}, a_t^{(i)})$. For simplicity, we still denote the transformation as $\phi(s, a)$, mapping a set of train-action pairs $(s, a) = ((s^{(1)}, \dots, s^{(n)}), (a^{(1)}, \dots, a^{(n)}))$ to their next state $(\phi(s^{(1)}, a^{(1)}), \dots, \phi(s^{(n)}, a^{(n)}))$. As a result, we obtain

$$p_\pi(s' | s) = \sum_{a \in \mathcal{A}} \pi(a | s) \mathbb{1}[\phi(s, a) = s']$$

whose sum has at most one non-zero term since ϕ is injective for a fixed s : by construction, distinct actions map to distinct states, with $\pi(a | s) = \prod_{i=1}^n \pi_i(a^{(i)} | s)$.

As delay can be deduced from states, we rewrite delay prediction as prediction of future states:

$$p_\pi(d_t^{(i)} | s_0) = f(p_\pi(s_1, \dots, s_T | s_0))$$

where f retrieves the delay from the predicted states.

Via Markov assumption:

$$p_\pi(s_1, \dots, s_T | s_0) = \prod_{t=1}^T p_\pi(s_t | s_{t-1})$$

Exact evaluation is infeasible, so we approximate the distribution with Monte Carlo rollouts of policy π . In this work, we obtain the delay point forecast by retrieving the median of the empirical distribution from the sampled trajectories.

In summary, we have reduced delay prediction to learning a policy π that approximates $p(a | s)$. The following Section presents the process of learning said policy.

5 Drift-Corrected Imitation Learning (DCIL)

We assume an (implicit) expert policy π_* has generated a logged dataset of train trajectories,

$$\mathcal{D} = \left\{ \left((s_0, a_0), \dots, (s_T, a_T) \right)_{n=1}^{|\mathcal{D}|} \quad a_t \sim \pi_*(\cdot | s_t). \right.$$

The term ‘‘expert’’ is purely notional: it refers to the behavioural patterns encoded in historical train-movement data rather than to a conscious decision-maker. Each recorded trajectory is treated as if it were sampled from π_* .

Our goal is to learn a policy π_θ via imitation learning that closely matches the expert action distribution $\pi_*(a_t | s_t)$.

The standard method within Imitation Learning is Behavioural Cloning (BC): it trains a parametric policy π_θ by minimising the cross-entropy between the predicted action distribution $\pi_\theta(\cdot | s)$ and the one-hot expert labels from the demonstration set \mathcal{D} . However, this method is notoriously prone to covariate shift: prediction errors push the policy into unseen states, where its performance quickly degrades. In order to overcome this limitation in our setting, we propose Drift-Corrected Imitation Learning (DCIL), a self-supervised extension of DAgger that injects simulator-generated corrective labels during training, counteracting covariate shift without requiring any additional expert queries or complex adversarial schemes.

Starting from a state s_0 drawn from a demonstration trajectory $(s_0, \dots, s_T) \in \mathcal{D}$, we roll out our policy π_θ for T steps to obtain a *policy trajectory* (s_0, s'_1, \dots, s'_T) . For each intermediate state s'_t we construct a synthetic ‘‘expert’’ action a_t^* by selecting, for each train i in the network, the action $a_t^{*(i)}$ that brings it closest to its next ground truth state $s_{t+1}^{*(i)}$:

$$a_t^{*(i)} = \arg \min_{a \in \mathcal{A}^{(i)}} \psi(s_{t+1}^{(i)}, \phi(s'_t, a)),$$

where $\psi : \mathcal{S}^{(i)} \times \mathcal{S}^{(i)} \rightarrow \mathbb{R}_{>0}$ is a distance in state space. In this work, ψ computes the number of stations separating the two individual states $s^{(i)}$ and $s'^{(i)}$ along the itinerary of train i . At each step t , we choose for each train the action that minimizes the *future* distance, i.e., the action $a^{*(i)}$ is 0 when the train is at the right station or ahead, 1 when it is one station behind, and 2 when it is 2 or more stations behind.

This algorithm is off-policy. As a result, we store synthetic state-action pairs in a replay buffer and use them for multiple epochs to improve sample efficiency.

Because synthetic labels become less reliable as the roll-out drifts farther from the expert trajectory, we down-weight their influence according to the distance between the current step t policy states $s_t^{(i)}$ and the next expert state $s_{t+1}^{(i)}$. One SGD step on the cross-entropy loss uses the scaled gradient

$$\sum_i \frac{1}{1 + \alpha \psi(s_{t+1}^{(i)}, s_t^{(i)})^\beta} \nabla_\theta \ell(\pi_\theta(s'_t)^{(i)}, a_t^{*(i)}) \quad (2)$$

with hyperparameters $\alpha > 0$ and $\beta \geq 1$. Thus, trajectories very close to the expert receive nearly full weight, while poorly matching ones contribute only marginally.

The full training loop, described in Algorithm 1, starts with initializing π_θ and B . Then, for E epochs, it (i) rolls out

Algorithm 1: Drift-Corrected Imitation Learning (DCIL)

Require: demonstrations \mathcal{D} ; horizon T ; buffer capacity C ; new samples per epoch n_s ; epochs E ; mini-batch size m ; learning rate η ; weighting factors α, β

- 1: initialise policy parameters θ
- 2: initialise empty replay buffer B
- 3: **for** $e = 1$ **to** E **do**
- 4: $k \leftarrow 0$
- 5: **while** $k < n_s$ **do**
- 6: sample $(s_0, \dots, s_T) \sim \mathcal{D}$
- 7: roll out $(s_0, s'_1, \dots, s'_T) \sim p_{\pi_\theta}$ from s_0
- 8: **for** $t = 0, \dots, T - 1$ **and** $k < n_s$ **do**
- 9: $a_t^* : a_t^{*(i)} = \arg \min_{a \in \mathcal{A}^{(i)}} \psi(s_{t+1}^{(i)}, \phi(s'_t, a))$
- 10: $w_t : w_t^{(i)} = \frac{1}{1 + \alpha \psi(s_{t+1}^{(i)}, s_t^{(i)})^\beta}$
- 11: push (s'_t, a_t^*, w_t) into B ; **if** $|B| > C$ **then** discard oldest
- 12: $k \leftarrow k + 1$
- 13: **end for**
- 14: **end while**
- 15: **for** each mini-batch $M \subset B$ of size m **do**
- 16: $g \leftarrow \nabla_\theta \sum_{(s', a^*, w) \in M} \sum_i w^{(i)} \ell(\pi_\theta(s')^{(i)}, a^{*(i)})$
- 17: $\theta \leftarrow \text{Adam}(\theta, g, \eta)$
- 18: **end for**
- 19: **end for**

the current policy to inject n_s synthetic samples, discarding the oldest to keep the buffer at capacity C , and (ii) performs Adam updates on mini-batches drawn from the entire buffer B , where each gradient is down-weighted according to the distance-based weight in Eq. (2).

6 Experimental Evaluation

Data

Our empirical evaluation uses raw operational logs provided by INFRABEL, the Belgian railway infrastructure manager, covering a three-year period from **1 January 2022** to **31 December 2024**. Every 30 seconds, we build a network-wide *snapshot* that contains the vector encoding of each train, as per Section 2. Network density varies from fewer than 10 trains during late-night service to more than 400 at peak hour. To augment the schedule context, we insert each train at a *placeholder* station 5 min before its planned departure and keep it at a *placeholder* station for 5 min after its final observed stop. Records with missing or incoherent timestamps are discarded (less than 1% of trains). Finally, we uniformly subsample 10% of the *snapshots*, yielding **255 k snapshots** with **51 million** train instances in total.

We adopt a strict temporal split to avoid information leakage and match industrial settings with a full calendar year for testing to maximise seasonal and operational diversity:

- **Training:** 2022-01-01 \rightarrow 2023-09-30
- **Validation:** 2023-10-01 \rightarrow 2023-12-31
- **Test:** 2024-01-01 \rightarrow 2024-12-31

Evaluating a single *snapshot* with the simulation-based approach requires sampling 50 trajectories and simulating 66 time steps; on an NVIDIA A100 this takes about 1.5 seconds of wall-clock time. To keep evaluation tractable, we down-sample the test split to 800 *snapshots* (≈ 200 trains per *snapshot*), yielding **1.6 million** arrival-time predictions (≈ 10 per train) that are produced in roughly 20 min end-to-end. Test snapshots are uniformly drawn, but busy periods produce more predictions, effectively emphasizing peaks.

Reproducibility. All experiments are fully reproducible; code and appendix are available on GitHub (link on the first page).

Method-Specific Details

In this work, we are interested in multi-station delay forecasting: i.e., predicting delay for the next n stations. This section provides implementation details for the evaluated methods (regression, BC, DCIL) and architectures (XGBoost, MLP, Transformer).

Regression target Following Section 2, we model delay via the conditional distribution $p(d_l^{(i)} | s_t)$. For a train i travelling on itinerary l and currently located at station index j , we need a forecast for the next n stations. Rather than predicting absolute delays, we regress on the difference to the last known delay

$$\Delta d_{l_{j+1}:l_{j+n}}^{(i)} := (d_{l_{j+1}}^{(i)} - d_{l_j}^{(i)}, \dots, d_{l_{j+n}}^{(i)} - d_{l_j}^{(i)}).$$

Our model outputs a point estimate $\Delta \hat{d}_{l_{j+1}:l_{j+n}}^{(i)}$ conditionally on s_t ; we then recover absolute delays by adding the last known delay $d_{l_j}^{(i)}$, yielding the prediction $\hat{d}_{l_{j+1}:l_{j+n}}^{(i)}$.

Ensuring fair comparison To guarantee that regression and simulation models are evaluated on **exactly the same targets**, we restrict the predictive horizon to $H = 30$ min. In the final metrics, we therefore retain only station events whose observed arrival time lies within H of the reference *snapshot* s_0 . Because the regression baselines output a fixed-length vector, we set the output length to the next $K = 15$ stations—a compromise that covers almost all instances within 30 min while avoiding an unnecessarily large head that would penalise regression. The simulator is forced to respect the same limit and produces at most K predictions per train. Should a simulated train fail to reach one of the K future stations within the simulated steps (e.g. it keeps choosing STAY: $a = 0$), we copy forward its most recent known delay and, if necessary, enlarge it so that the predicted arrival time never precedes the last simulated time step. To allow the simulator to make mistakes by predicting arrival times beyond the predictive horizon, we simulate 10% more time steps.

Method-dependent station/line horizon The raw features of Section 2 are computed for the previous five visited stations and lines for *all* models. To capture information about upcoming stations and lines, we include the next five scheduled stations/lines when the input feeds the simulator and the next fifteen for a regression model.

Mitigating Stalled Policies in BC During roll-out, we observed that the BC learned policy occasionally *stalls*: it assigns almost all probability mass to the STAY action ($a = 0$) at every step, so the train never advances. We keep for each train a running floor μ , defined as the largest value of $\pi(a=1)$ since that train last advanced. At every step, we clamp $\pi(a=1) \leftarrow \max(\pi(a=1), \mu)$ and then renormalise the action distribution; this simple trick reduces the fraction of stalled trains and improves long-horizon accuracy. This isn’t necessary for DCIL.

Architecture-Specific Details

To highlight the contrast between *regression* and *simulation* paradigms, we benchmark three architectures of ascending expressive power: a classical gradient-boosted tree ensemble (**XGBoost**), a lightweight **multi-layer perceptron (MLP)**, and an encoder-only **Transformer**.

Transformer. Following (Arthaud, Lecoer, and Pierre 2024), each input token represents one train, i.e. $s_t^{(i)}$ defined in Section 2. Through the attention mechanism, the model can propagate information between trains, predicting: (i) delay for regression and (ii) probability over the action space for simulation, conditionally to the state of the network.

MLP/XGBoost Due to the fixed-length input nature of both the multi-layer perceptron (MLP) and XGBoost, we cannot feed them the variable-size set of trains present in a network *snapshot* the way the Transformer does. As a result, the model inputs are train-specific rather than snapshot-specific.

In order to give the model network-related information, we build features that we concatenate with those introduced in Section 2. For *each* train, we compute, at the five radii $r \in \{0.1, 0.3, 0.6, 1.0, 2.0\}$ (chosen ad hoc to capture different ranges of network interactions),

1. the *count* of neighbouring trains lying within Euclidean distance r in the embedding space, min-max normalised
2. the *mean* of the past delays of trains in the neighborhood

These ten scalars provide a compact summary of the local traffic context; appending them to the base feature vector yields the fixed-length representation required by both the MLP and XGBoost models.

Hyperparameter tuning Table 1 and Table 2 are produced with the following two-phase protocol. First, for each architecture we carry out a *model-specific, multi-stage grid search* over a subset of hyper-parameters, using mean absolute error (MAE) on the validation set as the selection metric. Second, the best configuration is **run** ten times with different random seeds; each run is retrained on the combined train and validation sets and evaluated on the held-out test set. These runs are compiled using mean \pm standard deviation across metrics.

7 Results and Discussion

Table 1 reports the mean and standard deviations of test-set errors from 10 runs for all three model families—Transformer, MLP and XGBoost—under the three

Model	Method	MAE	RMSE
Transformer	Regression	60.53 ± 0.33	116.21 ± 4.06
	BC	56.17 ± 1.02	106.78 ± 2.82
	DCIL	52.24 ± 0.31	96.34 ± 0.58
MLP	Regression	64.34 ± 0.84	112.20 ± 2.37
	BC	59.10 ± 0.67	107.29 ± 1.39
	DCIL	57.76 ± 0.41	105.23 ± 0.73
XGBoost	Regression	64.48 ± 0.02	117.59 ± 0.19
	BC	66.77 ± 0.04	120.42 ± 0.19

Table 1. Test-set errors (mean ± std), all predictions.

training methods: *Regression*, *Behavioural Cloning* (BC, the imitation-learning baseline), and *Drift-Corrected Imitation Learning* (DCIL, our proposed imitation-learning method). BC and DCIL predictions are retrieved using the median of an ensemble of 50 sampled trajectories. Since we rescale predictions and labels back to their original units, evaluation metrics such as Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) are measured in seconds. Lower values indicate better performance for both metrics.

For Transformer, moving from pure Regression to the simulation baseline BC already yields a clear improvement: MAE drops by 7.2% and RMSE by 8.1%. DCIL amplifies these gains, ultimately reducing MAE by 13.7% and RMSE by 17.1%, respectively, relative to Regression, and by roughly half of that again relative to BC. The accompanying lower standard deviations hint at a stabler optimisation compared to BC when trajectory-level feedback is used. Regression yields surprisingly worse results for RMSE, hinting at a lack of robustness to outliers.

The same pattern holds for the MLP. BC narrows the gap to Regression (-8.1% MAE and -4.4% RMSE), while DCIL delivers the best absolute scores and the lowest variance (-10.2% MAE and -6.2% RMSE versus Regression). Thus, even a lightweight neural network benefits from simulation. Noticeably, the gap between BC and DCIL is smaller.

In contrast, the tree-based XGBoost does not profit from the imitation-learning signal. Switching from Regression to BC actually *increases* the errors by 3–5% across the board. DCIL was not evaluated here because the iterative training scheme is not compatible with XGBoost.

With DCIL, the Transformer achieves the lowest overall errors (MAE = 52.24, RMSE = 96.34). Interestingly, even the simpler MLP architecture (1.4M parameters), when trained with imitation learning, surpasses the Transformer (19M parameters) trained purely by regression: MLP-BC attains an MAE of 59.10 s and an RMSE of 107.29 s, while MLP-DCIL improves further to 57.76 s and 105.23 s, respectively—both lower than the Transformer-Regression baseline (MAE = 60.53 s, RMSE = 116.21 s).

Error by predictive horizon

Table 2 breaks the MAE down into six 5-minute predictive horizon bins. Predictions in the 0-5 minutes bin correspond to events where the *observed* arrival time is within 0 to 5 minutes of the *snapshot*. As a result, the trends observed in

the aggregate metrics can be more sharply analysed.

For Transformers, BC reduces MAE by 21.2% at 0–5 min to 2.9% at 25–30 min relative to Regression; DCIL further cuts the error by 27.2% and 8%, respectively, delivering the lowest values in every horizon. These gains indicate that DCIL has an advantage over BC, most notably in longer horizon predictions, highlighting covariate shift mitigation.

With MLP models, BC reduces MAE by 23.4% at 0–5 min to 2.8% at 25–30 min compared with Regression; DCIL brings the gains to 26.3% and 4.3%, respectively, mirroring the Transformer’s pattern.

Using XGBoost models, BC reduces MAE by 10.7% at 0–5 min to 1% at 5-10 min, but increases it by 2.5% at 10-15 min to 9.8% at 25–30 min relative to Regression. It performs better for short-range predictions and worse for the rest. Within the simulation setting, XGBoost’s errors grow faster with horizon than the deep models’, indicating a tree-specific sensitivity to distribution shift rather than a limitation of the simulation framework. In particular, the XGBoost policy is not explicitly trained as a stochastic policy or calibrated probabilistic model, so treating its outputs as action probabilities can amplify compounding errors under simulation.

DCIL with the Transformer remains the overall winner, yielding the best MAE across all predictive horizons. Notably, MLP-BC beats Transformer regression up until 20 minutes, while MLP-DCIL beats it all the way to 25 minutes, despite a much simpler model complexity. Taken together with the aggregate results, these horizon-wise findings confirm that simulation is more effective than regression for 30-minute delay prediction, and that DCIL effectively reduces the effect of covariate shift.

Discussion

As noted in Section 3, imitation learning is not limited to behavioural cloning. Generative Adversarial Imitation Learning (GAIL) is the standard non-BC alternative. Yet all of our GAIL runs with Transformer models collapsed: despite using PPO and applying stability tricks such as label smoothing, policy BC pretraining, gradient-norm clipping and reward shaping—the policy converged to a single constant action prediction, producing high validation MAE. We suspect the discriminator’s reward couples token (train) contexts in the Transformer, creating noisy, non-local gradients on individual actions. Decoupling state–action pairs could stabilise learning, but would discard key inter-train dependencies. This failure illustrates the complexity of adversarial imitation learning and motivates simpler, more stable training schemes such as DCIL that still counteract covariate shift.

Experimental results outline the superiority of simulation for train delay prediction compared with regression for deep learning methods. Most notably, simpler neural networks trained with an imitation learning scheme have better results than a Transformer using regression despite having $\approx 14\times$ fewer parameters and a less expressive architecture. This highlights that the training objective, rather than raw model capacity, is the primary lever for accuracy on this task.

Additionally, DCIL’s advantage over BC is more pronounced for Transformers than for MLP. This hints that

Model	Method	MAE ₀₋₅	MAE ₅₋₁₀	MAE ₁₀₋₁₅	MAE ₁₅₋₂₀	MAE ₂₀₋₂₅	MAE ₂₅₋₃₀
Transformer	Regression	30.63 ± 0.22	45.42 ± 0.27	58.67 ± 0.29	70.25 ± 0.42	80.27 ± 0.46	90.35 ± 0.53
	BC	24.14 ± 0.50	40.24 ± 0.86	54.39 ± 1.12	66.51 ± 1.28	77.17 ± 1.42	87.75 ± 1.30
	DCIL	22.31 ± 0.30	36.88 ± 0.40	50.02 ± 0.43	61.53 ± 0.44	72.11 ± 0.37	83.38 ± 0.51
MLP	Regression	32.54 ± 0.87	48.48 ± 0.79	62.64 ± 0.94	74.73 ± 1.01	85.28 ± 1.02	95.40 ± 0.82
	BC	24.93 ± 0.52	42.36 ± 0.75	57.24 ± 0.90	70.04 ± 0.91	81.29 ± 0.94	92.71 ± 1.31
	DCIL	23.98 ± 0.42	41.06 ± 0.62	55.86 ± 0.65	68.55 ± 0.55	79.73 ± 0.37	91.29 ± 0.54
XGBoost	Regression	31.31 ± 0.03	48.13 ± 0.03	62.84 ± 0.04	75.75 ± 0.03	86.25 ± 0.03	95.94 ± 0.03
	BC	27.95 ± 0.04	47.65 ± 0.06	64.38 ± 0.09	79.10 ± 0.07	92.28 ± 0.08	105.31 ± 0.09

Table 2. Test-set MAE (mean ± std) for different predictive horizons (5-minute bins)

DCIL scales with model expressiveness: the richer inductive bias of self-attention allows the Transformer to exploit the trajectory-level gradients that DCIL provides, whereas the lower-capacity MLP already harvests most of the benefit from simply matching expert actions step-by-step.

A natural next step is to test whether the simulation-based objectives retain their edge when labelled data are scarce. Because DCIL supplies a trajectory-level reward that is generated on-the-fly by the simulator, it can augment each real sequence with many synthetic roll-outs, effectively multiplying the learning signal. We therefore hypothesise that, under progressive down-sampling of the training set, the gap between DCIL and both BC and regression will widen—especially for the Transformer, whose larger capacity typically demands more data.

Table 2 shows that the absolute MAE difference between Regression and BC narrows as the forecast horizon grows (e.g. 6.5 s at 0–5 min vs. 2.6 s at 25–30 min). In contrast, the Regression–DCIL gap stays roughly constant (≈ 8 s), indicating that DCIL controls horizon-dependent drift more effectively than BC. Investigating whether the same trend holds for horizons longer than 30 minutes is an interesting direction for future work.

Uncertainty Quantification

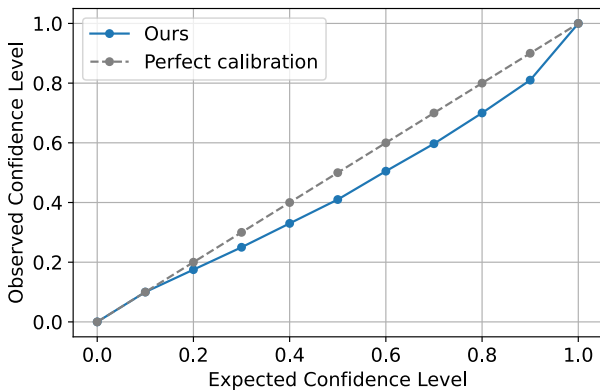


Figure 2: Calibration curve.

Following the procedure presented in (Kuleshov, Fenner, and Ermon 2018), we produce the calibration plot Fig. 2;

prediction intervals are formed from rollout percentiles, and their reliability is evaluated via percentile-based PIT calibration. The coverage curve lies consistently below the diagonal, indicating that the model is moderately overconfident: a nominal 80% prediction interval, for instance, contains the ground-truth delay only about 70% of the time. We attribute this under-coverage to difficulty modelling extreme-delay events—rare but operationally important—which the ensemble rarely samples, causing mass to concentrate in the right-most bin and leaving earlier bins under-populated. We hypothesise that this may, in part, stem from inputs that only weakly encode the precursors of extreme delays.

8 Conclusions and Future Work

In this paper, we framed probabilistic delay prediction as a sequential simulation task and trained an autoregressive policy via imitation learning using Behavioural Cloning and our novel approach, *Drift-Corrected Imitation Learning* (DCIL), a self-supervised extension of DAGger that curbs the covariate shift that plagues Behavioural Cloning without relying on an external oracle or complex adversarial schemes.

Through extensive evaluation on large-scale real-world data from INFRABEL, covering over three million train movements spanning three years, DCIL demonstrated superior performance compared with traditional regression methods and Behavioural Cloning on deep learning architectures for all predictive horizons.

Future work could refine the simulation by incorporating higher-resolution representations of the railway network, leveraging GPS-based data for added detail. A second line of inquiry involves investigating DCIL’s applicability to other imitation learning scenarios beyond delay prediction, assessing its generalisability and robustness across diverse domains. Furthermore, examining the theoretical and practical connections between DCIL and reinforcement learning—particularly model-based reinforcement learning—could provide valuable insights. Additionally, investigating the integration of DCIL within model-predictive control frameworks, specifically regarding performance and scalability over extended prediction horizons, presents a promising avenue for enhancing operational decision-making in complex, dynamic environments.

References

- Arthaud, F.; Lecoeur, G.; and Pierre, A. 2024. Transformers à Grande Vitesse: Massively parallel real-time predictions of train delay propagation. *Journal of Rail Transport Planning & Management*, 29: 100418.
- Belkin, M.; and Niyogi, P. 2003. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6): 1373–1396.
- Bendfeldt, J.-P.; Mohr, U.; and Muller, L. 2000. RailSys, a system to plan future railway needs. *WIT Transactions on The Built Environment*, 50.
- Dekker, M. M.; Medvedev, A. N.; Rombouts, J.; Siudem, G.; and Tupikina, L. 2022. Modelling railway delay propagation as diffusion-like spreading. *EPJ Data Science*, 11(1): 44.
- Goverde, R. M. 2010. A delay propagation algorithm for large-scale railway traffic networks. *Transportation Research Part C: Emerging Technologies*, 18(3): 269–287.
- Guo, J.; Wang, W.; Tang, Y.; Zhang, Y.; and Zhuge, H. 2022. A CNN-Bi-LSTM parallel network approach for train travel time prediction. *Knowledge-Based Systems*, 256: 109796.
- Heglund, J. S.; Taleongpong, P.; Hu, S.; and Tran, H. T. 2020. Railway delay prediction with spatial-temporal graph convolutional networks. In *2020 IEEE 23rd international conference on intelligent transportation systems (ITSC)*, 1–6. IEEE.
- Ho, J.; and Ermon, S. 2016. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29.
- Kecman, P.; and Goverde, R. M. 2015. Predictive modelling of running and dwell times in railway traffic. *Public Transport*, 7: 295–319.
- Koelemeijer, G. S.; Goverde, R.; Van Egmond, R.; et al. 2000. PETER, a performance evaluator for railway timetables. *WIT Transactions on the Built Environment*, 50.
- Kuefler, A.; Morton, J.; Wheeler, T.; and Kochenderfer, M. 2017. Imitating driver behavior with generative adversarial networks. In *2017 IEEE intelligent vehicles symposium (IV)*, 204–211. IEEE.
- Kuleshov, V.; Fenner, N.; and Ermon, S. 2018. Accurate uncertainties for deep learning using calibrated regression. In *International conference on machine learning*, 2796–2804. PMLR.
- Lapamponyio, P.; Derrible, S.; and Corman, F. 2022. Real-time passenger train delay prediction using machine learning: A case study with amtrak passenger train routes. *IEEE Open Journal of Intelligent Transportation Systems*, 3: 539–550.
- Martino, L.; Read, J.; Elvira, V.; and Louzada, F. 2017. Cooperative Parallel Particle Filters for online model selection and applications to Urban Mobility. *Digital Signal Processing*, 60(January): 172–185.
- Nash, A.; and Huerlimann, D. 2004. Railroad simulation using OpenTrack. *WIT Transactions on The Built Environment*, 74.
- Ng, A. Y.; Russell, S.; et al. 2000. Algorithms for inverse reinforcement learning. In *Icml*, volume 1, 2.
- Oneto, L.; Fumeo, E.; Clerico, G.; Canepa, R.; Papa, F.; Dambra, C.; Mazzino, N.; and Anguita, D. 2018. Train delay prediction systems: a big data analytics perspective. *Big data research*, 11: 54–64.
- Prevost, C. G.; Desbiens, A.; and Gagnon, E. 2007. Extended Kalman filter for state estimation and trajectory prediction of a moving object detected by an unmanned aerial vehicle. In *2007 American control conference*, 1805–1810. IEEE.
- Price, I.; Sanchez-Gonzalez, A.; Alet, F.; Andersson, T. R.; El-Kadi, A.; Masters, D.; Ewalds, T.; Stott, J.; Mohamed, S.; Battaglia, P.; et al. 2025. Probabilistic weather forecasting with machine learning. *Nature*, 637(8044): 84–90.
- Ross, S.; Gordon, G.; and Bagnell, D. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 627–635. JMLR Workshop and Conference Proceedings.
- Röbber, D.; Reisch, J.; Hauck, F.; and Kliewer, N. 2021. Discerning primary and secondary delays in railway networks using explainable AI. *Transportation Research Procedia*, 52: 171–178.
- Şahin, İ. 2017. Markov chain model for delay distribution in train schedules: Assessing the effectiveness of time allowances. *Journal of rail transport planning & management*, 7(3): 101–113.
- Sipilä, H. 2023. Simulations with PROTON and RailSys: Use of a macroscopic and microscopic railway simulation tool in Swedish applications.
- Spanninger, T.; Trivella, A.; Büchel, B.; and Corman, F. 2022. A review of train delay prediction approaches. *Journal of Rail Transport Planning & Management*, 22: 100312.
- Tiong, K. Y.; and Palmqvist, C.-W. 2023. Quantitative methods for train delay propagation research. *Transportation Research Procedia*, 72: 80–86.
- Torabi, F.; Warnell, G.; and Stone, P. 2018. Behavioral cloning from observation. *arXiv preprint arXiv:1805.01954*.
- Yang, S.; Guo, M.; Wu, J.; Chen, Y.; Huang, G.; and Zeng, B. 2024. Hybrid Inverted Transformer-CNN Model for Train Primary-Delay Recovery Time Prediction. In *Proceedings of the 4th Asia-Pacific Artificial Intelligence and Big Data Forum*, 1103–1110.
- Zhang, D.; Peng, Y.; Xu, Y.; Du, C.; Zhang, Y.; Wang, N.; Chong, Y.; Wang, H.; Wu, D.; Liu, J.; et al. 2022. A high-speed railway network dataset from train operation records and weather data. *Scientific data*, 9(1): 244.