

Prediction of arrival times of freight traffic on US railroads using support vector regression

William Barbour*, Shankara Kuppa[†] and Daniel B. Work[‡]

June 14, 2017

Abstract

Variability of travel times on the United States freight rail network is high due to large network demands relative to infrastructure capacity, especially when traffic is heterogeneous. Variable runtimes pose significant operational challenges if the nature of runtime variability is not predictable. To address this issue, this article proposes a data-driven approach to predict *estimated times of arrival* (ETAs) of individual freight trains, based on the properties of the train, the properties of the network, and the properties of potentially conflicting traffic on the network. The ETA prediction problem from an origin to a destination is posed as a machine learning regression problem and solved using support vector regression trained and cross validated on over two years of detailed historical data for a 140 mile section of track located primarily in Tennessee, USA. The article presents the data used in this problem and details on feature engineering and construction for predictions made across the full route. It also highlights findings on the dominant sources of runtime variability and the most predictive factors for ETA. Improvement results for ETA exceed 21% over a baseline predictor method for predictions made at some locations and average over 14% across the study area.

1 Introduction

1.1 Motivation

The rail network in the United States has significant infrastructure capacity limitations that cause congestion of the rail traffic. Few rail corridors contain exclusively double (or more) track that allows simultaneous bi-directional traffic [1]. In comparison, the double and triple track railroads in Europe provide for double the train density of US rail networks [2]. Many US corridors contain a single track with short sections of double track known as *sidings*, where trains may meet or pass each other. These *movements* (i.e., meets, passes) are implemented in the railroad signaling system, but are directed by human dispatchers. Dispatchers are experienced with working on specific track corridors, but movements on sidings require planning and precise timing in order to achieve efficient operations [3, 4]. Freight volume is expected to increase in the US, so either infrastructure capacity must be increased or operational improvements must be made to increase capacity [5–7].

In addition to the track infrastructure constraints, there are numerous other factors that can contribute to variability of the runtime on a track segment. Traffic heterogeneity and the train priority differences directly influence both the runtime of trains and also the variability in the runtime [8, 9]. Physical characteristics of trains such as the length, tonnage, and power further influence the runtime due to track grade, track curvature and siding lengths [8]. The ability of a train to complete a trip and exit the *line of road* (i.e., the track segments connecting distant terminals) is also influenced by the degree of congestion in the arrival terminal. This is compounded by the possible actions required for the train in the terminal, such as refueling,

*W. Barbour is with the Department of Civil and Environmental Engineering, University of Illinois at Urbana-Champaign, Champaign, IL 61801, USA.

[†]S. Kuppa is with CSX Transportation, Jacksonville, FL 32202, USA.

[‡]D. Work is with the Department of Civil and Environmental Engineering and Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA. email: (dbwork@illinois.edu).

inspection, switching of cars, or crew change [8, 10]. Railroad operating strategies such as dynamically scheduled trains and maximizing train length are particularly vulnerable to delay [11, 12].

In the presence of runtime variability, ETAs are necessary in order to improve real-time decision making and the efficiency of the network [13, 14]. For example, future train schedules can be continually updated to provide new train plans to allow traffic to flow smoothly between terminals on the network [15]. Although there are many techniques available to derive optimal schedules (see [16] for a thorough review), the schedule may be very sensitive to delays when the network is near capacity. High capacity utilization leads to more complex dispatching where small delays are created, leading to larger deviations from the train plan; this is referred to as *knock-on delay* [1, 3, 17].

Highly variable runtimes increase operational uncertainty for the railroad and for other transportation systems that depend on them. On the rail network, propagation of delay to other trains is significant [18], and there are large direct costs incurred due to additional operating time alone [19]. Delays on the rail network can also influence non-rail transportation services. For example, surface street traffic and emergency vehicles, which conflict with rail freight traffic at grade crossings [20], can be significantly delayed if a grade crossing is occupied by a train for an extended period of time. If accurate, real-time ETAs are made available, revisions to the operating plan can be implemented, and surface street transportation services can be re-routed.

1.2 Problem statement and solution approach

The main focus of the present article is the prediction problem for ETAs on US freight railroads using real-time data. The estimation problem requires new ETAs to be produced as time elapses and the train progresses down the line of road. Each time the train reaches one of a number of fixed locations on the track, data is collected and a new estimated travel time to the destination is produced.

To produce the ETA estimate, a variety of routinely collected and maintained data sources available to freight railroads are used. This includes track geometry data (containing grade and curvature information, single and multi-track territory, length of sidings, etc.), historical runtime profiles of all trains, properties of all trains (such as length and tonnage), and crew records.

Several methodologies to produce ETAs are available, including microscopic simulation [21–23], analytical approaches [24], and data-driven techniques [25, 26]. Due to the complexity of the freight rail network (which limits the accuracy of analytical abstractions) and the difficulty to capture all delay inducing factors in a simulation based model (e.g., decisions made by human dispatchers, special cases involving priority elevation, unplanned maintenance, and weather), a data-driven approach is proposed in this article [27]. This approach is made possible through access to a large and comprehensive freight rail dataset also described in the article.

1.3 Related work

Several lines of research are related to the problem of ETA prediction. We briefly summarize the most closely related works, and direct the interested reader to the comprehensive reviews available in the works by Bonsra and Harbolovic [25] and Gorman [28].

The majority of freight trains operate according to a schedule that is constructed in an offline manner and robust to some random unplanned disturbances [3, 12]. When extreme disturbances cause the original schedule to deteriorate, online rescheduling measures must be implemented to account for the delay and to maintain robustness to further delay [14, 29, 30]. Numerous efforts are aimed at understanding and quantifying the causes of delay that influence scheduling, rescheduling, and predictability [1, 9, 31]. Delay is typically formulated in terms of deviation from a train schedule or historical performance, but it can be extended to arrival time prediction for individual trains [25].

Several works have proposed to empirically produce delay or runtime estimates using historical data for passenger rail networks. Kecman and Goverde [4] proposed an ETA prediction framework for passenger rail arrival time prediction using track occupancy data for conflict evaluation. Chapuis [32] used artificial neural networks to predict arrival times of frequent passenger trains using historical train and station delays. Compared to the proposed work, the ETAs were evaluated in the Netherlands and France, respectively, on high priority passenger traffic [33, 34], which also operates with higher punctuality compared to the freight or passenger traffic in the US [35]. Wang and Work [26] estimate passenger rail delays on the Amtrak passenger rail network in the US using vector regression techniques and only historical runtimes between

passenger stations. The regression problems are formulated in both a historical and online perspective, but the feature set for prediction is limited and does not contain any data on the freight traffic, which constitutes the majority of traffic on the shared line of road in the US. Online methods presented for passenger rail, accessible because of the data stream created by station arrivals and departures, have not been fully extended to freight rail. Additionally, magnitude of delay for passenger rail is typically on the order of minutes, while delay for non-priority freight rail traffic may exceed multiple hours.

The most closely related estimation works on freight trains are the works of Gorman [28] and Bonsra and Harbolovic [25]. In Gorman [28], an econometric analysis of free-running and congestion related factors are used to identify the primary causes of delay. The data was partitioned by geographic area and priority groupings. Congestion related factors, such as meets, passes, and overtakes that occur, are found to have the largest effect on delay. Bonsra and Harbolovic [25] predict runtimes for individual freight trains in an offline setting using regression. Prediction improvements were attained when estimated at the time of departure. The regression model used train and network parameters and a historical runtime averaging technique for evaluating model performance.

1.4 Outline and contributions

The main contribution of this article is to show how to pose the ETA prediction problem on a rail network as a sequence of machine learning regression problems, where one regression is performed for each origin-destination pair. This modeling choice is supported by results showing clear distinctions in the learned models across the route, indicating the advantage of model specificity. We provide practical insights by highlighting the datasets available to perform the prediction and describing some of the feature engineering required when the feature vectors change in time and space. We present a set of data features and several machine learning regression algorithms used to achieve more accurate ETAs than common statistical methods yield. Finally, the resulting models are discussed in detail with respect to their performance.

The remainder of the article is organized as follows. Section 2 presents the framework used to process and operate on the various data sources and the preliminaries for the machine learning regression. Section 3 discusses the datasets and the work that is necessary to process the data for use in the machine learning framework as constructed data features. Section 4 details the model trials that are conducted as well as the means by which to evaluate them. Section 5 describes the process for tuning and evaluating a single model, which is then extended to models across the full testing route; results are given for models using select feature combinations. We provide a summary and discuss future lines of research in Section 6.

2 Framework and Problem Formulation

This section briefly describes the machine learning framework used for ETA estimation. It reviews general machine learning terminology and parameters specific to the algorithms, both used later during analysis and discussion.

2.1 ETA Machine learning framework

The problem of predicting an estimated time of arrival for a train from an origin point to a destination point on the rail network is posed as a supervised machine learning regression problem. The goal of the regression problem is to predict the true runtime $y(i) \in \mathbb{R}^1$ of a train i given the properties of train i , the network, and other traffic on the network, which are contained in the feature vector $x(i) \in \mathbb{R}^n$. Given a dataset of m trains with true runtimes $Y = [y(1), y(2), y(3), \dots, y(m)]^T \in \mathbb{R}^m$ and corresponding feature vectors $X = [x(1), x(2), x(3), \dots, x(m)]$, where $X \in \mathbb{R}^{n \times m}$, the machine learning regression problem is to find a mapping $f_w: \mathbb{R}^n \rightarrow \mathbb{R}^1$ parameterized by w such that $f_w(x(i))$ is an accurate predictor of $y(i)$. In general, supervised machine learning regression uses a set of *training data* $\{X_{tr}, Y_{tr}\}$ (where the subscript *tr* is used to indicate the training data) to learn the function f_w , by minimizing prediction error between $f_w(X_{tr})^1$ and Y_{tr} over the m records in the training data.

¹With a slight abuse of notation, we overload the function f_w to also operate on the entire dataset X_{tr} .

The machine learning model f_w must generalize (i.e., make good predictions on data that has not been used to train the model), in order to maintain high accuracy on new data and to avoid *overfitting* the training data. To test the degree of generalization, the accuracy of the prediction is assessed on hold out *test data* $\{X_{te}, Y_{te}\}$, which is not used to train the model.

The central difficulty of posing the train ETA prediction problem into the framework above stems from the fact that many of the features used for prediction change in time and in space as the train moves towards the destination. For example, the amount of traffic on the line of road will change as trains enter and leave the line of road. The number of available sidings on a route in single track territory also changes across the route and as other trains occupy or vacate sidings. If a single model is used for all origin-destination predictions in the network, it may be difficult to predict area-specific delays (e.g., due to local dispatching decisions and route characteristics) that may not occur throughout the network; results pertaining to origin-destination specificity are discussed in Section 5.2. Moreover, because some features change over time (as described above), while others may not (e.g., train priority), construction of an unbiased training dataset is a nontrivial engineering task. For example, one cannot simply create a new training data point each time a single property of a train changes (e.g., corresponding to a new feature vector) without biasing the training data, since the feature vector still corresponds to a single train trip.

2.2 Generating ETAs on a freight rail network

To address these difficulties, we propose to build a distinct regression model for each origin-destination pair for which predictions are required. Because the models are independent, each model can be trained using all trips that pass between the corresponding origin-destination pair by constructing features according to the state of the train and network at the time the train reaches the origin node. Localized and geography-specific performance effects may be captured in the individual models without explicitly constructing them in the feature vector. For example, longer travel times will be observed for heavy or underpowered trains in areas of high track grade. Feature construction can also vary between models (e.g., in dimensionality) since each uses a custom dataset built for the origin-destination pair.

The primary disadvantage of building a model for each origin-destination pair in the network is the number of models required. In a rail network with k nodes, at minimum k^2 origin-destination predictors are required (possibly more if multiple paths exist between each origin-destination pair). In contrast to road networks where the number of nodes and the number of viable paths between any two node pairs may be large, rail networks have fewer nodes and less path redundancy. In practice, few locations are relevant destination points from a given origin because a single route (excluding small deviations for sidings) is typically used to connect two points on the network. Therefore, the number of origin-destination paths for which predictions are required is tractable. In the area of study in this work, there exist 35 points that can serve as origin points; there are only four practical destination points from each origin point, which results in at most 148 predictors as opposed to $35^2 = 1225$.

In order to map spatiotemporal train data to the network topology, infrastructure data can be reconstructed into a directed graph format, $G = (V, E)$ where V is a set of vertices and E is a set of directed edges. Vertices are points where the track merges and splits (e.g., endpoints of sidings). Data on passing trains is recorded at *OS-points*, which are fixed locations $v \subset V$. Directed edges represent track segments across which trains travel between OS-points and a direction of travel. This alignment between track infrastructure and the constructed graph is shown in Figure 1. In the figure, OS-points are denoted $a_0, a_1, a_2,$ and a_3 with corresponding graph vertices $v_0, v_1, v_2,$ and v_3 . All tracks are delineated between these OS-points with multiple tracks such as the main line and siding between a_1 and a_2 remaining distinct. Pairs of directed edges representing each delineated track segment allow different runtimes and feature values in each direction of travel, which is necessary when properties such as grade are considered [25]. This results in two directed graph edges for a single track and four or more edges for an area of multiple tracks such as the siding between vertices v_1 and v_2 . In this formulation, all trains can be routed on the graph across their unique path considering track usage (e.g., siding track versus main line track). Data can be gathered on the behavior of trains for each directed edge with respect to speed, grade, train occupancy, and other location/direction specific attributes. Also, features that consider estimates of the positions of multiple trains and track topology can be mined from this data.

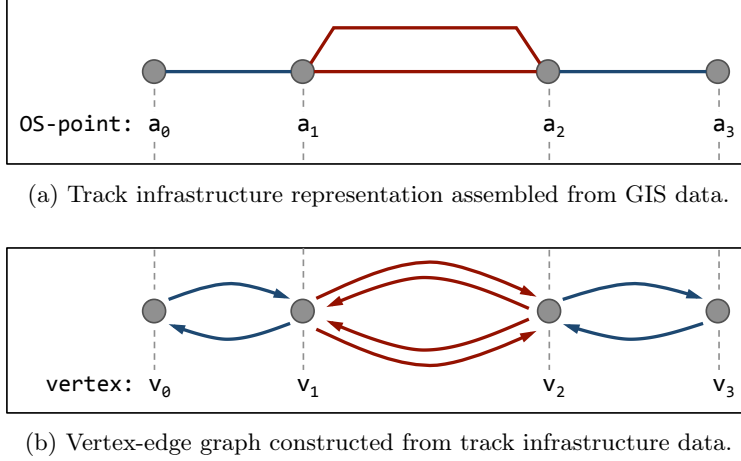


Figure 1: Graph vertices align with OS-points and each track segment between them is represented by a graph edge in each direction. Double track areas and sidings, therefore, are represented by four graph edges.

2.3 Regression via support vector machines

The regression problem of predicting an ETA from a vector of features is proposed to be solved with a *support vector regression* (SVR) machine, introduced in [38]. Support vector regression is a popular machine learning algorithm grounded in statistical learning theory and for which training is efficient due to the convexity of the training problem. The model parameters in linear SVR are straightforward to interpret, which can be invaluable in the application of the algorithm. Additionally, the SVR formulation provides for extension to nonlinear regression via kernel functions. Other applicable algorithms include linear ridge regression [39], elastic net regression [40], and kernel ridge regression [41], to name a few.

The training step in a generalized regression problem may be written as:

$$\min_w L(f_w(X) - Y) + \|w\|, \quad (1)$$

where L is a loss function measuring the quality of the predicted output, $f_w(X)$, relative to the true output, Y , and the feature weights w are penalized via a norm to avoid overfitting the training data. This characteristic is informally referred to as model *flatness* [42].

SVR is a special case of (1) and uses a two-norm on w and in the simple case assumes an affine predictor of the form $f_w(x) = w^T x + b$, where $w \in \mathbb{R}^n$ and the offset $b \in \mathbb{R}$. SVR uses an ε -insensitive loss function $|\cdot|_\varepsilon$, which penalizes prediction residuals $r = y - f_w(x)$ larger than a threshold defined by ε . The loss function is constructed as [43]:

$$|r|_\varepsilon = \begin{cases} 0 & \text{if } |r| \leq \varepsilon \\ |r| - \varepsilon & \text{otherwise.} \end{cases} \quad (2)$$

The ε -insensitive loss function quantifies the distance between a prediction and the band created by $y \pm \varepsilon$. For a vector of residuals, the sum of the element ε -insensitive losses are used as the loss function.

The training step in SVR can be reformulated as computing the weights w and offset b by solving the following convex optimization problem:

$$\begin{aligned} & \underset{w, b, \xi, \xi^*}{\text{minimize}} && \frac{1}{2} \|w\|_2^2 + C \sum_{i=1}^m (\xi(i) + \xi^*(i)) \\ & \text{subject to} && y(i) - w^T x(i) - b \leq \varepsilon + \xi(i), \\ & && w^T x(i) + b - y(i) \leq \varepsilon + \xi^*(i), \\ & && \xi(i), \xi^*(i) \geq 0, \end{aligned} \quad (3)$$

where $\xi, \xi^* \in \mathbb{R}^m$ are variables introduced to rewrite the ε -insensitive loss (2) as linear inequalities. The total ε -insensitive loss (i.e., accuracy of model fit) is balanced against model flatness by a scalar factor C .

The optimization problem (3) can be solved via the dual problem, yielding the optimal dual variables $\alpha, \alpha^* \in \mathbb{R}^m$. These dual variables are related to the feature weights such that $w = \sum_{i=1}^m (\alpha(i) - \alpha^*(i))x(i)$. This results in a predictor of the form:

$$f(x) = \sum_{i=1}^m (\alpha(i) - \alpha^*(i))x(i)^T x + b, \quad (4)$$

see [36] for a comprehensive description.

When the ETAs in the training data are not linearly related to the features, an alternative strategy is to transform the training data into a much higher dimensional feature vector denoted by $\Phi(x)$, where $\Phi: \mathbb{R}^n \rightarrow \mathbb{R}^N$ with $N \gg n$, which can then be used for regression. The new predictor becomes:

$$f(x) = \sum_{k=1}^m (\alpha(k) - \alpha^*(k))\Phi(x(k))^T \Phi(x) + b. \quad (5)$$

Interestingly, it is not necessary to explicitly define the mapping to the high dimensional space, since only the inner product $\Phi^T \Phi$ is needed in the regression function. The inner product can instead be defined through a kernel function $K(x(k), x(j)) = \Phi(x(k))^T \Phi(x(j))$. The use of the kernel function directly in (5) is known as the *kernel trick* [44] in machine learning. In the present work, we adopt the *radial basis function* (RBF) kernel [45] of the form:

$$K(x(k), x(j)) = \exp\left(-\frac{1}{2\sigma^2} \|x(k) - x(j)\|_2^2\right), \quad (6)$$

where σ is a parameter controlling the decay rate of the kernel, effectively limiting the influence that any single observation may have on the trained model.

3 Feature construction and data cleaning steps

This section discusses the process of combining and mining datasets to be used in feature construction. The data used in this work is described first, before describing the features that are calculated from the datasets which are subsequently used to train the machine learning ETA algorithm. Due to the proprietary nature of the data, some values are reported in relative terms.

3.1 Description of raw data

This work uses a series of datasets describing the rail network and operations from December 1, 2014 through January 31, 2017 inclusive. It consists of freight train movement, train car operations, crew, and locomotive data in the CSX Transportation network, extracted from dispatching, operations, and signaling data.

The movement data consists of records generated at OS-points between terminals. The data includes the track on which the train was reported and the time at which the train triggered the OS-point. This dataset also contains information about track mileage covered, direction of travel, and the next destination at which the train is scheduled to stop.

The train car operations data details the actions performed on the train once it enters a terminal from the line of road. This includes the switching operations (i.e., picking up and setting out rail cars) that are referred to as train *work*, inspections, refueling, and crew changes that are scheduled to occur and may incur delay in getting track space in the terminal. The planned work schedule, as well as adherence to the schedule, is reported in this data. Changes in physical train characteristics (e.g., total number of cars, length, tonnage) are inferred based on the work reported on the train.

Crew data contains information about the crew assigned to the train, the originating location, the time at which they were called on duty, and the time at which the crew must legally go off duty (i.e., 12 hours after going on duty). The time between a crew going on duty and the departure of the train is non-negligible and is referred to as *on duty time to departure* (ODTOD). Crew information is important because the maximum crew on-duty requirement must always be satisfied, even at the large expense of stopping a train and transporting a replacement crew to finish the trip.

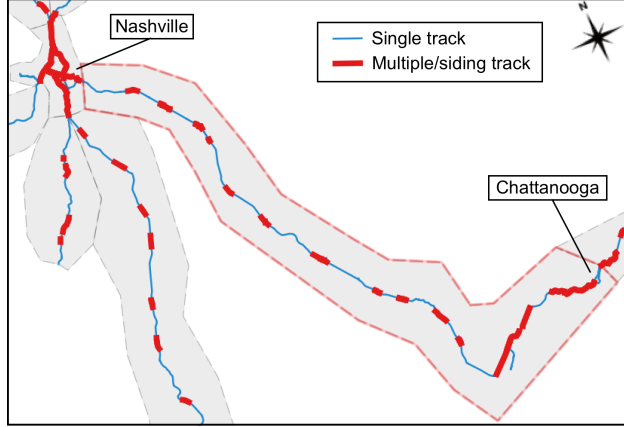


Figure 2: GIS network view depicting a portion of the Nashville division, with multi-track segments shown in bold red lines and single-track segments in thin blue lines. The primary study route is bounded with the dashed line.

Locomotive assignment data indicates the equipment and total locomotive power available on each train, which can be important for predicting delays in regions with high grades.

This work also uses GIS data describing the physical infrastructure of the network, which includes individual tracks, switches, mileposts, and terminals. All locations referenced in the movement, work, crew, and locomotive data map to physical infrastructure locations, such as track mileposts and control points. Reconciling these GIS data sources is necessary to gather data on the number of tracks and siding locations and lengths on each route and build the network graph described in Section 2.2. Figure 2 depicts this data, along with the distinction between single track sections (shown by the thin blue lines) and multiple track sections or sidings (shown by the bold red lines) for a portion of the rail network. The primary study area is bounded by the dashed line.

3.2 Data cleaning and standardization tasks

A variety of data cleaning and data transformation tasks are necessary to organize the input for any prediction algorithm. With over 150,000 trips in the division in a two year period, the decision was made to neglect trips with data completeness issues or data errors instead of devising a scheme to impute values; this resulted in the discarding of approximately 10% of trains. Errors consist of fields that contain missing data, or fields that contain illogical values. Examples include non-physical train lengths, or an arrival time prior to the train departure time.

The GIS data is examined to ensure proper connectivity and accuracy before being transformed into the network graph. Common errors encountered include duplicate geometries, disconnected track components, and minor mislabeling of infrastructure components. Many errors are automatically identified and resolved, while some errors require manual correction.

The detection and resolution methods for each of these data fields are summarized in Table 1. Each was implemented at the time of data mining and feature construction, so that an origin-destination dataset is clean at the time of model training.

3.3 Handling of recrewed trains

In the process of early prediction efforts and data exploration efforts, a dominant source of runtime variability was discovered. Specifically, it was found that *recrewed* trains (i.e., a train that did not reach its destination before the crew reached its maximum on-duty time and needed a relief crew) define the dominant source of runtime variability on the study route.

To further investigate the impact of recreds on train variability, all trains were ex post facto labeled as either recrewed or non-recrewed. Less than 10% of the trains on the route were recrewed. The two classes

Table 1: Data cleaning and standardization steps including detection and fix.

Data field	Data error	Data correction
Train arrival time	arrival time \leq departure time	discard train
Train length	length ≤ 0	discard train
Train tonnage	tonnage ≤ 0	discard train
Train horsepower	horsepower ≤ 0	discard train
Crew assignment	no crew assigned to train	discard train
Crew on duty time	on duty time \geq train departure time	discard train
Track segments	polylines connect spatially, but endpoint ID mismatch	detect spatial connection, resolve endpoint ID mismatch
Duplicate GIS elements	identical geometry encoded strings	check connectivity of each, keep most connected element

(recrewed and non-recrewed trains) were separated and descriptive statistics were calculated for each class at each of the 35 OS-points, which are spread at irregular intervals across the route depicted in Figure 2. The standard deviation of runtimes was used to quantify the runtime variability of trains in each class as well as the variability of all trains in the dataset (not separated on recrew). As shown in Figure 3, the runtime variability of the recrewed trains is several times larger than that of the non-recrewed trains across all OS-points; runtime variability is expressed as a relative value to protect proprietary operational properties in the data. Despite recrewed trains representing less than 10% of the trips, they represent 53% of the variability within the dataset of all trains, when averaged across the full route.

Recrewed trains introduce high variability in runtime and their runtimes are not predictable by features inside the scope of the train and network state features (e.g., location and status of potential relief crews is a significant factor and is not in our dataset). It is likely, however, that the circumstances leading to a recrew will enable its preemptive classification and could be captured with the available data. For the scope of this project, recrewed trains were historically identified and removed from the training data as part of data cleaning and standardization steps.

3.4 List of calculated features

This section lists and discusses the features that are generated from the raw data in Section 3.1 and used to train machine learning algorithms in ETA prediction. A summary of the implemented scalar features appears in Table 2. These features include six train characteristics, two features that capture the state of the crew on each train, and multiple scalar features quantifying the network characteristics and traffic. Numerous series of features describe the traffic state of the network in the geographic vicinity of prediction. These all depend on segmentation of the track between OS-points. The network traffic state is described for these segments in terms of occupancy, direction, and priority; these are summarized in Table 3. All of the features chosen for exploration were based on extensive discussions with operations research personnel from CSX Transportation.

The priority of a train is determined by its cargo (e.g., bulk, merchandise, automotive, intermodal), and its type of service (e.g., local, yard, road). These combinations are categorized and ranked in different levels of granularity. At the highest resolution, all trains are placed into one of twenty priority classes. The relative priority ranking is understood to be non-linear based on its construction. The high resolution ranking was aggregated to a medium-resolution ranking using five priority classes and a low-resolution ranking of three priority classes. For example, scheduled merchandise trains have significantly higher priority than bulk/unit trains (e.g., loaded coal train), but in medium- and low-resolution classifications, the two types will get the same priority designation.

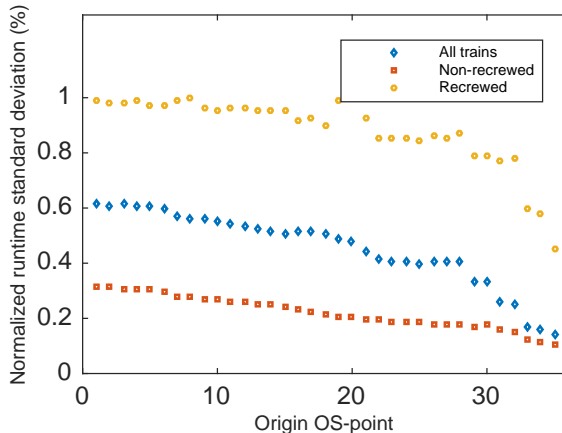


Figure 3: Comparison of variability in runtime, between recrewed trains, non-recrewed trains, and all trains, for each origin OS point. Standard deviation is normalized by the largest variance OS-point, OS-point #21. OS-point IDs increase from Nashville (1) to Chattanooga (42).

The physical train characteristics such as train length and train tonnage were calculated by examining the work data which contains the train dimensions after the most recent work was completed. The distributions of these parameters are shown in Figures 5a and 5b and demonstrate that there is a preferred maximum threshold for train length, but that train tonnage is subject to a significant tail of very heavy trains. Train length factors into train performance in conjunction with the number of sidings in which each may fit. While the majority of sidings are sufficiently long for any train, certain sidings can not accommodate lengthier trains. This disparity is reflected in Figure 5c with a smaller number of trains fitting in 50-70% of sidings and the majority able to use 80-100% of remaining sidings.

The most recent crew change can be identified in the crew data, which is used to calculate the maximum time the current crew may continue to operate the train, called *crew time remaining*. The distribution of crew time remaining is shown in Figure 5d. Ideally, this parameter should be maximized when setting out on a route to give the best chance of not needing another train crew and this is reflected by the increasing frequency of longer remaining crew times; it should be noted that all trains in this distribution were not recrewed. The practical importance of this feature is in terms of the difference between crew time remaining and expected train runtime (constant), called *slack time*. If the crew time remaining is less than the typical train runtime, then a train may be expedited in order to avoid a recrew. In terms of feature construction, crew time remaining and slack time reduce to the same values under min-max normalization applied to each element in the feature vector used for the model input X .

A total of six scalar quantifying measures of traffic are constructed, each based on the line of road between the current location of a train and the destination. First, within the remaining line of road, all other trains including local trains are counted, and then categorized based on the direction of travel for each (e.g., same direction, subscript ω , or opposite direction, subscript ψ , relative to the train being predicted). A second parameter categorizes directional traffic counts (two counts) by their priority, as higher (subscript α) or lower/equal priority (subscript β) relative to the train being predicted, resulting in another four features.

The exact traffic state of the origin-destination route and surrounding area is described further using numerous feature series, the components of which correspond geographically to track segments. The segments refer to the connections between OS-points, and span the origin-destination route and distances around the origin and destination limited to 50% of the origin-destination route length. The dimension of each feature series is equal to the number of track segments that are considered. In our study area of approximately 140 miles, the 35 OS-points delineate the segments on the origin-destination route. OS-points within 70 miles of the origin and 70 miles of the destination not included on the origin-destination route delineate the track segments around the origin and around the destination. These series features are enumerated and explained in Table 3. Segment occupancy is a series of binary features, each of which is non-zero when another train is present in a track segment at the time that a prediction is made for a train at the origin point. Likewise,

Table 2: Summary of implemented scalar features.

Feature	Notation	Description
Train length	λ_i	The total length of locomotives and cars of train i .
Train tonnage	μ_i	The total mass of locomotives and cars.
Train horsepower per ton	η_i	Total horsepower of locomotives divided by train tonnage, μ_i .
Train priority (high-resolution)	$\rho_{20,i}$	Priority ranking on a 1-20 scale.
Train priority (medium-resolution)	$\rho_{5,i}$	Five priority classes are constructed by aggregating the high-resolution priority ranking.
Train priority (low-resolution)	$\rho_{3,i}$	Three priority classes are constructed by aggregating the high-resolution priority ranking.
Crew time remaining	γ_i	Amount of time remaining that the current train crew can legally work.
On duty time to departure	θ_i	Time between crew on duty time and train departure.
Full traffic count	τ_i	Count of trains on the remaining line of road.
Directional traffic count	$\tau_{\omega,i}, \tau_{\psi,i}$	Count of trains on the remaining line of road, divided by direction of travel (i.e., in the same direction, ω , or in the opposite direction of travel, ψ).
Prioritized directional traffic count	$\tau_{\omega,\alpha,i}, \tau_{\omega,\beta,i},$ $\tau_{\psi,\alpha,i}, \tau_{\psi,\beta,i}$	Count of train on the remaining line of road, divided by both direction of travel and priority relative to that of the train being predicted (i.e., lower or equal priority, β , or higher priority, α).
Available sidings	π_i	Count of sidings on route with length greater than that of train i .

trains that are present on these segments can be described with respect to their relevant properties, namely direction of travel, priority, and relative priority. These features are strictly zero for unoccupied segments and may have positive scalar values for segments occupied by other trains. This process of describing the traffic state via network segments and traffic properties is illustrated in Figure 4. Predictions are made at the origin OS-point a_0 and OS-points delineating segments are labeled a_0 through a_l (for origin-destination route, only). An example traffic scenario for the moment at which a prediction is made for train Z_0 at the origin is shown with trains Z_1, Z_2, Z_3 . The relevant features (i.e., direction and priority) of each train are listed, and are mapped to the track segments corresponding to the location of each train.

4 Model implementation and evaluation

This section describes a set of model experiments and a metric to assess the machine learning methods described above. The feature sets used in the models are composed of the features described previously in Section 3.4.

Table 3: List of calculated and implemented series features, which correspond in dimension to the segmentation of the remaining route, and details for each.

Feature	Notation	Description
Track segment occupancy	$O_i = [O_{1,i}, \dots, O_{l,i}]$	Vector of elements denoting whether a segment, numbered 1 through l on the origin-destination route, is occupied by another train; non-zero when occupied.
Occupying train direction	$D_i = [D_{1,i}, \dots, D_{l,i}]$	Denotes the direction (same or opposite) of a train occupying a track segment numbered 1 through l on the origin-destination route; zero if segment is unoccupied.
Occupying train priority	$P_i = [P_{1,i}, \dots, P_{l,i}]$	Assigns high-resolution train priority values to a train occupying a track segment on the origin-destination route. Higher value for high priority train; zero if no train.
Relative priority of occupying train	$R_i = [R_{1,i}, \dots, R_{l,i}]$	Non-zero when a train occupying a track segment on the origin-destination route has higher priority than the train for which the ETA is being predicted. Reflects likelihood of meet/pass delay.
Track segment occupancy around origin point	$G_i = [G_{-1,i}, \dots, G_{-h,i}]$	Indicates track segment occupancy for segments -1 through $-h$ around the origin point, but not included in the primary route (segments 0 to l); captures trains that may enter the primary route and pass/overtake.
Track segment occupancy around destination point	$E_i = [E_{l+1,i}, \dots, E_{l+p,i}]$	Indicates track segment occupancy for segments $l+1$ through $l+p$ around the destination point, but not included in the primary route; captures trains that may enter the primary route and conflict during the trip.

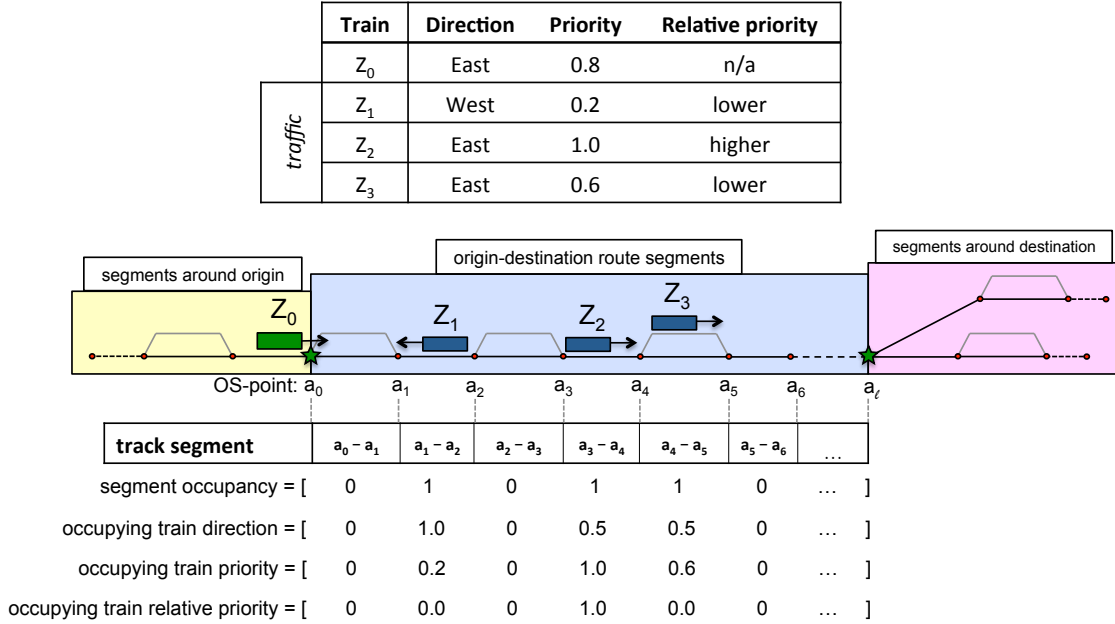


Figure 4: Segment-wise series features are calculated for the area around an origin point, on the origin-destination route, and around the destination. Each is segmented by OS-points, a_0 through a_l on the origin-destination route in this case. The occupancy series feature is constructed, followed by series based on properties of the occupying train, if present.

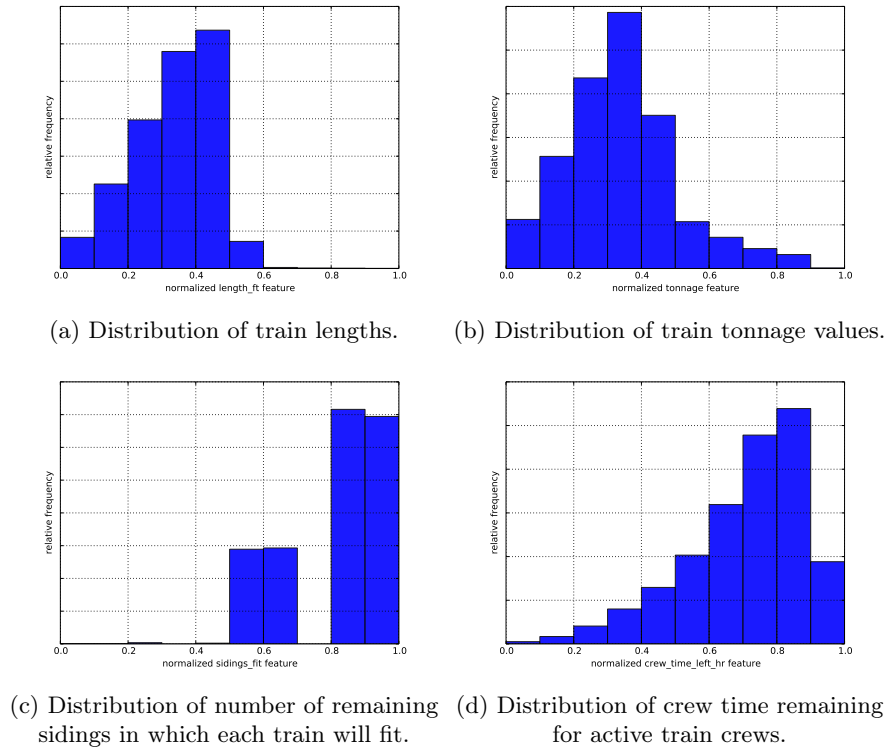


Figure 5: Frequency distributions of various scalar features in the training dataset taken at OS-point #1, outside of Nashville.

4.1 Description of models

Numerical experiments were performed with concentration on a single route, shown by the dashed area in Figure 2, in the Nashville division of the CSX Transportation network which contains a mix of single and double track segments, highly heterogeneous traffic, and high volume relative to capacity. The route represents one of the most challenging segments on which to estimate ETAs within the CSX Transportation network. Without loss of generality of the methods, the present analysis is restricted to common train types with sufficient trips in the two year dataset and includes the automotive, merchandise, and intermodal trains. These train types have differing priorities, and consequently have distinct behaviors in meet/pass movements and when delays occur. The dataset for trains running the full study route in the correct direction of travel initially contains over 10,000 trips. When the dataset is filtered by train type, recreated trains are removed, local trains and trains with intermediate work are eliminated, and data errors and incomplete records are removed, there are still approximately 4,200 trips.

The selected route is composed of 35 points along the 140 mile route for which an ETA to the destination must be produced. For each of the 35 ETA problems, a total of five models are implemented and compared. The models include the baseline median predictor algorithm as well as four SVR-based algorithms. Many combinations of algorithm type and feature set were explored, and the presented models are representative of the model type and performance. For example, the various priority features were each evaluated for predictive performance by performing single-feature model trials and $\rho_{5,i}$ was found to be the most informative.

The exact model configurations are as follows:

- Model 0: baseline median predictor where $f(x(i)) = \text{median}\{y(i) \mid y(i) \in Y_{tr}\}$
- Model 1: linear SVR with all scalar features (length, tonnage, hp/ton, priority, crew time, ODTOD, traffic counts, and sidings fit); the feature vector is constructed as: $x(i) = [\lambda_i, \mu_i, \eta_i, \rho_{5,i}, \gamma_i, \theta_i, \tau_i, \tau_{\omega,i}, \tau_{\psi,i}, \tau_{\omega,\alpha,i}, \tau_{\omega,\beta,i}, \tau_{\psi,\alpha,i}, \tau_{\psi,\beta,i}, \pi_i,]$, where $x(i) \in \mathbb{R}^{14}$.
- Model 2: linear SVR with all scalar features plus track segment occupancy series feature $x(i) = [\lambda_i, \mu_i, \eta_i, \rho_{5,i}, \gamma_i, \theta_i, \tau_i, \tau_{\omega,i}, \tau_{\psi,i}, \tau_{\omega,\alpha,i}, \tau_{\omega,\beta,i}, \tau_{\psi,\alpha,i}, \tau_{\psi,\beta,i}, \pi_i, [O_{1,i}, \dots, O_{l,i}]]$, where $x(i) \in \mathbb{R}^{14+l}$.
- Model 3: linear SVR with all scalar features and all series features (occupancy, direction, priority, and relative priority on origin-destination route; occupancy around origin point; occupancy around destination point) $x(i) = [\lambda_i, \mu_i, \eta_i, \rho_{5,i}, \gamma_i, \theta_i, \tau_i, \tau_{\omega,i}, \tau_{\psi,i}, \tau_{\omega,\alpha,i}, \tau_{\omega,\beta,i}, \tau_{\psi,\alpha,i}, \tau_{\psi,\beta,i}, \pi_i, [O_{1,i}, \dots, O_{l,i}], [D_{1,i}, \dots, D_{l,i}], [Q_{1,i}, \dots, Q_{l,i}], [R_{1,i}, \dots, R_{l,i}], [G_{-1,i}, \dots, G_{-h,i}], [E_{l+1,i}, \dots, E_{l+p,i}]]$, where $x(i) \in \mathbb{R}^{14+4l+h+p}$.
- Model 4: RBF kernel SVR with all scalar features and all series features $x(i) = [\lambda_i, \mu_i, \eta_i, \rho_{5,i}, \gamma_i, \theta_i, \tau_i, \tau_{\omega,i}, \tau_{\psi,i}, \tau_{\omega,\alpha,i}, \tau_{\omega,\beta,i}, \tau_{\psi,\alpha,i}, \tau_{\psi,\beta,i}, \pi_i, [O_{1,i}, \dots, O_{l,i}], [D_{1,i}, \dots, D_{l,i}], [Q_{1,i}, \dots, Q_{l,i}], [R_{1,i}, \dots, R_{l,i}], [G_{-1,i}, \dots, G_{-h,i}], [E_{l+1,i}, \dots, E_{l+p,i}]]$, where $x(i) \in \mathbb{R}^{14+4l+h+p}$.

4.2 Model evaluation

The error metric used to evaluate a given model is *mean absolute error* (MAE), defined as:

$$MAE = \frac{1}{p} \sum_{i=1}^p |f(x(i)) - y(i)|, \quad (7)$$

where $f(x(i))$ and $y(i)$ correspond to the predicted runtime and true runtime of train i , respectively, and p denotes the number of records in the testing dataset. It follows that numerically low MAE *scores* are better than high scores. Performance of each model is compared to that of the historical median predictor, Model 0. The improvement for each model is given as the reduction in the MAE relative to the historical median predictor.

The machine learning algorithms and benchmark predictors are implemented in Python and leverage the scikit-learn package [46]. The data processing steps are completed once for each origin-destination pair and the feature set is stored in a database. Model trials are performed by loading the feature set, selecting the desired data, normalizing features, and training the model and testing the performance via cross-validation. Data processing and model testing both occur on a dedicated, modern workstation computer with 4 GHz

quad-core processor with eight virtual cores and dynamic overclocking, 64 GB DDR4 RAM, and NVMe solid state drives. Building the feature set is the most time consuming step in the process, due primarily to the size of the database of raw data. This process can take up to 15 minutes per origin-destination pair but is only required once, and adding more features and data does not require reprocessing of previous data. Model training is accomplished in 5.0 to 60.0 seconds depending on the model parameters, with prediction on all test data taking no more than 3.0 seconds.

5 Results

This section first presents the process for choosing hyper-parameters C and ε in (3) for a single origin-destination model with scalar features. We then analyze the series linear of models trained with scalar features on the full 35-OS-point route and the differences between them, specifically with respect to feature weights. Performance results are then shown for each model in Section 4.1 across the route, which demonstrate the impact of feature richness and the nonlinear kernel.

5.1 Choosing hyper parameters for a single model

For each origin-destination model, the SVR parameters, C and ε are chosen as a result of the training and testing process. This is performed using a dataset containing approximately 4200 trips using a five-fold cross validation with an 80/20 training/testing data split. On each fold of the cross validation, the parameter space is explored using a grid search that explores all combinations of parameters within the bounds of each. The results are aggregated across the five folds and the optimal parameter combination is chosen based on the mean minimum testing error. The trained model must be checked for suitability such that it generalizes well to testing data. This check is done using validation curves for C and ε and a learning curve for the amount of data used to train the model.

The interpretation of these parameters as well as analysis of training and testing behavior kept the search space limited. The ε parameter is directly related to residual values between $f(x(i))$ and $y(i)$ and, therefore, can be limited to a search space proportional to the normalized spread of the true outputs $y(i)$. The C parameter penalizes the model training error, summed across all observations $x(i)$, relative to the model flatness, given by the two norm of feature weights. We normalize C such that it is scaled by the number of features and inversely scaled by the number of observations in the training dataset. This maintains the impact of the parameter across models with different dimensionalities.

A validation curve explores training and testing scores across a range of a model parameter, with other parameters fixed. Using a fixed C value of 1.0, the validation curve for the ε parameter in Figure 6a shows relatively little effect of ε value on training and testing scores. Mean training scores are plotted in the solid orange line, with the narrow shading showing a single standard deviation above and below the mean of training scores in cross-validation. The mean testing scores and distribution of scores are plotted with the dashed purple line and shading. Within the acceptable parameter range, high values nor low values make an appreciable effect on testing MAE. Approaching $\varepsilon = 0.0$, the ε -insensitive loss disappears and the algorithm converges to linear regression. When ε is set too high, the minimization of prediction errors focuses only on observations with exceptionally high residual prediction values. In comparison, the value of C has significantly higher impact on model score. The validation curve is also shown in Figure 6b with ε fixed at 0.1 and plotted with common normalized MAE score for comparison. Small values of C emphasize model flatness, but result in poor training and testing scores because model complexity is low. Large values of C achieve low training MAE but generalize poorly to the testing data because of overfitting.

Validation curves show sensitivity for individual parameters. The optimal parameters are chosen simultaneously by evaluating the model on the grid space of all parameter combinations ($C \in [10^{-5}, 10^4]$, $\varepsilon \in [0.0, 0.3]$). For origin-destination model at OS-point #1, the optimal values that minimize MAE are found to be $C = 0.75$ and $\varepsilon = 0.05$; the difference in training and testing scores is less than 3% of training error, which is clearly avoiding overfitting.

The learning curve for a model shows the convergence of training and testing performance by increasing the amount of data available to build the model. The curve is analyzed after hyper-parameters have been chosen. The learning curve shows the MAE score against the number of training examples available to the model. With smaller amounts of data available, training scores will be improved, but at the expense of the

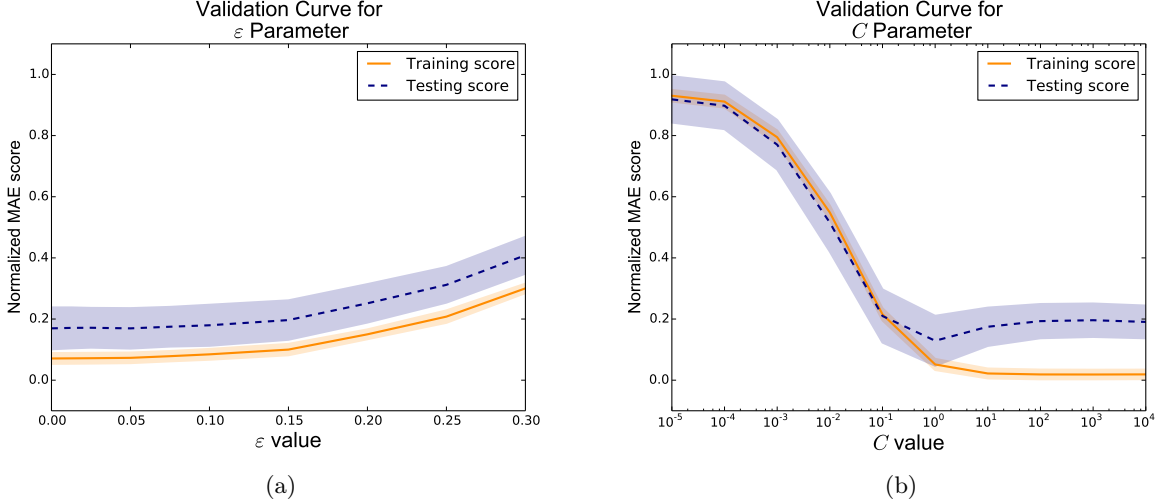


Figure 6: The validation curves for the ϵ and C parameters on a single origin-destination model are plotted with a common MAE score axis, which is min-max normalized across both parameters. The sensitivity of the model to ϵ is relatively low compared to that of C .

model generalizing poorly to testing data. The learning curve in Figure 7 indicates that the model is trained on a sufficient amount of data because the curves converge before the full training dataset is used. The mean training and testing scores are denoted by the lines with the shaded areas showing one standard deviation in each direction between cross-validation folds. The model training and testing occur at 10%, 32.5%, 55%, 77.5% and 100% of available data. The disparity in training and testing scores begins at over 40% of training score with only 10% of available data and decreases to less than 3% with 100% of available data.

For any SVR model with a linear kernel, the feature weights can be interpreted meaningfully in both magnitude and sign. The feature weights were recorded at all cross validation folds and for each origin-destination model to assess the relative impact of each. The feature weights are normalized by absolute sum within each model, such that $\sum_{j=1}^n |w_j| = 1$, to allow comparison between models and are reported in absolute terms for ranking in terms of absolute impact. The feature weight rankings for the origin-destination model at OS-point #1 (the beginning of the route in Nashville) are shown in Table 4. For this model, the effect of train priority is the dominant feature; this is supported by the statistically distinct runtimes between priority classes at this distance from the destination. Crew time remaining has a large impact because it can affect the runtime of a train at this distance if the train experienced significant delay leaving its last terminal. Tonnage also play a large role due to the lower overall performance of the train during acceleration and deceleration. Features with particularly low impact include traffic counts separated by direction and priority, which is an overly simplistic view of the traffic state on long routes. Horsepower per ton also has less of an impact because few trains are under powered in an area such as this with significant terrain.

5.2 Model training across route

The hyper-parameter selection process and model evaluation is replicated for origin-destination predictions made at each of the 35 OS-points on the full route. Because the C parameter is normalized by the training data size and feature dimension, and the ϵ parameter demonstrated little sensitivity, the optimal set of hyper-parameters found by the selection process varied little across the route.

The main finding is that reported feature weights show significant variability across the route. This supports the idea that dispatching techniques have fundamental differences based on relative location of the train to a terminal point and that unique origin-destination models capture some of this nuance. All scalar feature weights are shown at each prediction point in Figure 8. The means of each feature at each location are represented by the lines and min-max ranges for each are shown by the shaded area around the lines. In prediction of the full route, at OS-point #1, the factor most heavily impacting train runtime is priority. Other factors certainly play into the dispatching decisions made for the route, but do not appear to have strong




Figure 7: The learning curve for a single origin-destination model shows convergence of the training and testing error scores given increasing availability of observations in the full dataset. The MAE score values are min-max normalized.

Table 4: Average feature weights for 5-fold cross validation on origin-destination prediction at OS-point #1 (Nashville). Exact model output weights are normalized by absolute sum. Features are ordered by highest mean absolute weight.

Feature	Mean absolute weight
Priority, P_5	0.346
Crew time remaining, C	0.137
Tonnage, M	0.110
Traffic opposite direction lower/equal priority, $T_{U,L}$	0.089
Available sidings, S	0.067
Total traffic, T	0.055
Traffic opposite direction, T_U	0.050
Length, L	0.047
Traffic same direction, T_W	0.040
Horsepower per ton, Q	0.019
Traffic opposite direction higher priority, $T_{U,H}$	0.011
Traffic same direction higher priority, $T_{W,H}$	0.011
Traffic same direction lower/equal priority, $T_{W,L}$	0.011
On duty time to departure, θ	0.008

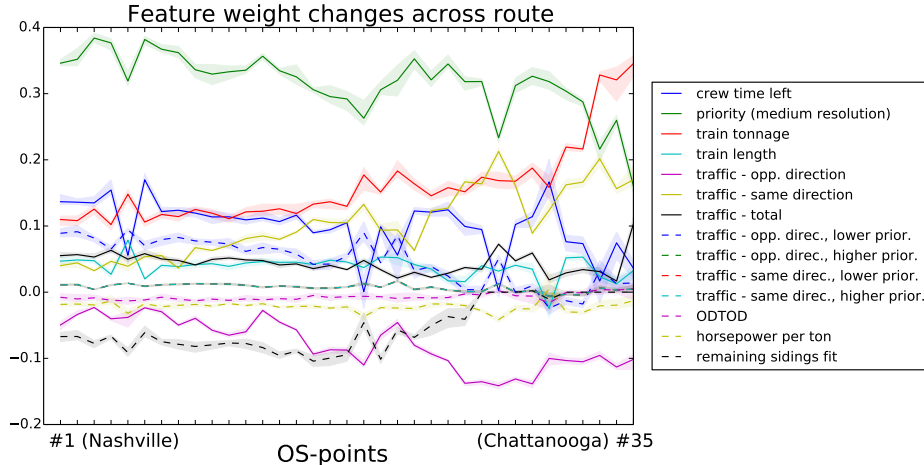


Figure 8: Feature weight change of scalar features in origin-destination SVR models across the route, Nashville (1) to Chattanooga (35). Feature importance is measured by absolute magnitude.

relationships far from the destination. Closer to the destination, traffic counts and train tonnage are driving factors due to decisions around the yard and natural choke point in the network. The changing importance of train characteristics supports the domain understanding that many factors contribute to train ETA and the impact of these factors is not constant. Along the route, feature weight experience some sharp changes, but some of these can be explained by certain characteristics of the route. For example, the sharp dip at OS-point #27 in the weight of priority and crew time remaining, along with the sharp increase in weight on traffic in the same direction, is located on the most significant hill on the route. At this location, a separate helper locomotive attaches to and assists some trains in climbing the hill. Availability of this locomotive is a driving factor in runtime from this location and its presence seems to be captured indirectly in the feature set by the number of trains in the same direction without being explicitly defined as its own feature. This separation of origin-destination models is, after all, intended to encode these geographic features in the data via specificity.

This specificity serves as a justification for the separation of origin-destination predictions to distinct models, as opposed to a unified model predicting ETAs for all origins and destinations. Feature weights may be learned optimally for a single model, instead of forcing feature weighting conditioned upon location information. It is possible that some features may change weight predictably relative to the distance from the destination point such that a distance normalization in feature construction could account for weight change. However, the results of feature weights across the route show that many of the trends are inconsistent and would not be captured fully using the technique.

5.3 Performance comparison of SVR models

Performance of each model detailed in Section 4.1 is evaluated across all OS-points on the route. The four non-baseline models demonstrate increasingly levels of model complexity based on the richness of features included in each, but are not a comprehensive list of all feature combinations that were tested.

The model results across the full route are shown in Figure 9 in terms of relative reduction in MAE over the historical median predictor (Model 0). A features set with only scalar features (Model 1), as explored in the choice of hyper-parameters and examination of feature weights in Sections 5.1 and 5.2, represents the largest incremental performance gain for every origin-destination predictor. Inclusion of the segment-wise occupancy feature series (Model 2) and inclusion of all segment-wise features series (Model 3) each attain small MAE improvements in addition to improvements gained by the scalar features. The RBF kernel, however, does not provide any substantive performance gain over the fully featured linear model. The γ parameter in the RBF kernel was chosen by exploring the range $\gamma \in [10^{-4}, 10^3]$ in a grid search alongside the ϵ and C parameters. This lack of performance gain is likely due to the high variability of the data. Even in non-linear regression, the relationship between inputs, X , and output, Y , is not strongly defined.




Figure 9: Improvement in MAE over baseline historical median predictor for each model at all OS-points between Nashville (1) and Chattanooga (35).

The scalar feature set (Model 1) contains information representing basic relationships and understandings about how the rail network functions. Unsurprisingly, groups of trains such as those that are heavy and low priority generally run slower than the lighter high priority trains. A few counts indicating the amount of traffic on the route will be roughly indicative of the total amount of delay due to meets and passes. These sorts of relationships can be determined by a linear model given the information in the scalar feature set. This information would be crucial to any ETA prediction, but it does not provide a complete picture.

Based on the relative performance improvements of the SVR models over the baseline, it is evident that the addition of the segment-wise network traffic state features (Models 2 and 3) show clear advantages by providing high-resolution information compared to the simple counts of network traffic (Model 1). The larger gains are achieved by the segment-wise occupancy feature series, but the additional information on the direction and priority of the traffic improve the model performance by better informing on the likely meets and passes that will occur. For instance, the mere presence of another train on the route will be somewhat likely to increase runtime, but if this train is traveling in the direction opposing trains on the origin-destination route and has high priority then it may be highly likely to increase runtime.

Model performance varies somewhat across the origin OS-points due to the distribution of route delay. The spacing of sidings and the likelihood of each to be used varies due to their length and topological characteristics of the route. Overall, the relative performance of the models to each other is consistent across the route. Prediction performance decreases closer to the destination. We expect this is due to the more unpredictable nature of operations close to rail yards. The factors that affect the exact arrival of a train when it gets close are not necessarily present in the data (e.g., ability of the yard to accept more trains, availability of the next train crew)

These route results are summarized in Table 5 in terms of mean, maximum, and minimum percent improvement over the baseline. The minimum improvement values are consistently observed for models close to the destination point and the maximum improvement is consistently observed near the beginning of the route.

Table 5: Comparison of SVR model performance to baseline predictor, summarized for the 35 OS-points on the full route.

Predictor	Mean % improvement	Maximum %	Minimum %
Model 0 (Baseline)	0.0%	0.0%	0.0%
Model 1	9.4%	14.2%	4.0%
Model 2	12.2%	19.9%	4.6%
Model 3	14.0%	21.6%	6.2%
Model 4	14.3%	21.8%	7.0%

6 Conclusion

This article presents a data driven approach to predict ETAs on freight rail networks. The ETA generation problem is posed as a series of independent origin-destination ETA prediction problems to capture location specificity and avoid bias in the training data of a single general model due to time varying features. The problem is tractable for rail networks due to relatively low network complexity that reduces the number of relevant origin-destination pairs. This approach is shown to demonstrate specificity with respect to distinct feature weights (i.e., relative importance) between origin-destination models. This nuance would not be captured by a unified model.

ETA prediction models are constructed for each relevant origin-destination pair using combinations of features and linear and nonlinear SVR algorithms. Compared to naive prediction based on historical median runtimes, average improvements of 14% and maximum improvements of over 21% are achieved by the best performing SVR models.

Based on these findings, our future research steps include the following. Due to the large variance caused by reworks, we are interested in developing a data-driven classifier to preemptively classify trips that are likely to be reworked. For the trains that are not likely to be reworked, further improvements in the ETA accuracy are possible with the construction of additional targeted traffic features constructed from the output of a meet-pass planner. We are also interested in building models of the state of the origin or departure yard, which may create delays that cascade onto the line of road.

Acknowledgment

The authors acknowledge funding by the Roadway Safety Institute, the University Transportation Center for USDOT Region 5, which includes Minnesota, Illinois, Indiana, Michigan, Ohio, and Wisconsin. Financial support was provided by the United States Department of Transportation’s Office of the Assistant Secretary for Research and Technology (OST-R).

References

- [1] Pavankumar Murali, Maged Dessouky, Fernando Ordóñez, and Kurt Palmer. A delay estimation technique for single and double-track railroads. *Transportation Research Part E: Logistics and Transportation Review*, 46(4):483–495, 2010.
- [2] Oliver Wyman. Assessment of european railways: Characteristics and crew-related safety, June 2016.
- [3] Michiel JCM Vromans, Rommert Dekker, and Leo G Kroon. Reliability and heterogeneity of railway services. *European Journal of Operational Research*, 172(2):647–665, 2006.
- [4] Pavle Kecman and Rob MP Goverde. An online railway traffic prediction model. In *RailCopenhagen2013: 5th International Conference on Railway Operations Modelling and Analysis, Copenhagen, Denmark, 13-15 May 2013*. International Association of Railway Operations Research (IAROR), 2013.
- [5] Cambridge Systematics. *National rail freight infrastructure capacity and investment study*. Cambridge Systematics, 2007.
- [6] Brian A Weatherford, Henry H Willis, David S Ortiz, Louis T Mariano, J Enrique Froemel, and Sara A Daly. *The State of US Railroads: A Review of Capacity and Performance Data*. Rand Corporation, 2008.
- [7] Association of American Railroads. Class I railroad statistics, February 2013.
- [8] Mark Dinger, Yung-Cheng Lai, and Christopher Barkan. Impact of train type heterogeneity on single-track railway capacity. *Transportation Research Record: Journal of the Transportation Research Board*, (2117):41–49, 2009.
- [9] Mark Dinger, Amanda Koenig, Sam Sogin, and Christopher PL Barkan. Determining the causes of train delay. In *AREMA Annual Conference Proceedings*, 2010.
- [10] Andrew Higgins, E Kozan, and L Ferreira. Modelling delay risks associated with train schedules. *Transportation Planning and Technology*, 19(2):89–108, 1995.
- [11] Quan Lu, Maged Dessouky, and Robert C Leachman. Modeling train movements through complex rail networks. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 14(1):48–75, 2004.
- [12] Shi Mu and Maged Dessouky. Scheduling freight trains traveling on complex networks. *Transportation Research Part B: Methodological*, 45(7):1103–1123, 2011.
- [13] Julie H Hertenstein and Robert S Kaplan. *Burlington Northern: the ARES decision (a)*. Harvard Business School, 1991.
- [14] Susan F Hallowell and Patrick T Harker. Predicting on-time performance in scheduled railroad operations: methodology and application to train scheduling. *Transportation Research Part A: Policy and Practice*, 32(4):279–295, 1998.
- [15] David R Kraay and Patrick T Harker. Real-time scheduling of freight railroads. *Transportation Research Part B: Methodological*, 29(3):213–229, 1995.
- [16] Rob MP Goverde. *Punctuality of railway operations and timetable stability analysis*. TU Delft, Delft University of Technology, 2005.
- [17] Rob MP Goverde and Lingyun Meng. Advanced monitoring and management information of railway operations. *Journal of Rail Transport Planning & Management*, 1(2):69–79, 2011.
- [18] Andrea D’Ariano and Marco Pranzo. An advanced real-time train dispatching system for minimizing the propagation of delays in a dispatching area under severe disturbances. *Networks and Spatial Economics*, 9(1):63–84, 2009.

- [19] Alexander H Lovett, C Tyler Dick, and Christopher PL Barkan. Determining freight train delay costs on railroad lines in North America. *Proceedings of RailTokyo 2015*, 2015.
- [20] R Estes and Laurence Rilett. Advanced prediction of train arrival and crossing times at highway-railroad grade crossings. *Transportation Research Record: Journal of the Transportation Research Board*, (1708):68–76, 2000.
- [21] ER Petersen and AJ Taylor. A structured model for rail line simulation and optimization. *Transportation Science*, 16(2):192–206, 1982.
- [22] İsmail Şahin. Railway traffic control and train scheduling based on inter-train conflict management. *Transportation Research Part B: Methodological*, 33(7):511–534, 1999.
- [23] Marin Marinov and José Viegas. A mesoscopic simulation modelling methodology for analyzing and evaluating freight train operations in a rail network. *Simulation Modelling Practice and Theory*, 19(1):516–539, 2011.
- [24] Arjang A Assad. Models for rail transportation. *Transportation Research Part A: General*, 14(3):205–220, 1980.
- [25] Kunal Baldev Bonsra and Joseph Harbolovic. Estimation of run times in a freight rail transportation network. Master’s thesis, Massachusetts Institute of Technology, 2012.
- [26] Ren Wang and Daniel B Work. Data driven approaches for passenger train delay estimation. In *Intelligent Transportation Systems (ITSC), 2015 IEEE 18th International Conference on*, pages 535–540. IEEE, 2015.
- [27] Hongfei Li, Dhaivat Parikh, Qing He, Buyue Qian, Zhiguo Li, Dongping Fang, and Arun Hampapur. Improving rail network velocity: A machine learning approach to predictive maintenance. *Transportation Research Part C: Emerging Technologies*, 45:17–26, 2014.
- [28] Michael F Gorman. Statistical estimation of railroad congestion delay. *Transportation Research Part E: Logistics and Transportation Review*, 45(3):446–456, 2009.
- [29] Andrea D’Ariano, Marco Pranzo, and Ingo A Hansen. Conflict resolution and train speed coordination for solving real-time timetable perturbations. *IEEE Transactions on Intelligent Transportation Systems*, 8(2):208–222, 2007.
- [30] Harshad Khadilkar, Shripad Salsingkar, and Sudhim Kumar Sinha. A machine learning approach for scheduling railway networks. In *Proceedings of the 7th International Conference on Railway Operations Modelling and Analysis, RailLille2017*. International Association of Railway Operations Research (IAROR), April 2017.
- [31] Bintong Chen and Patrick T Harker. Two moments estimation of the delay on single-track rail lines with scheduled traffic. *Transportation Science*, 24(4):261–275, 1990.
- [32] Xavier Chapuis. Arrival time prediction using neural networks. In *Proceedings of the 7th International Conference on Railway Operations Modelling and Analysis, RailLille2017*. International Association of Railway Operations Research (IAROR), April 2017.
- [33] Francisco Manuel Bastos Andrade Furtado, Francisco Manuel Bastos, et al. US and European freight railways: The differences that matter. In *Journal of the Transportation Research Forum*, volume 52, pages 65–84. Transportation Research Forum, 2013.
- [34] Hamed Pouryousef, Pasi Lautala, and Thomas White. Railroad capacity tools and methodologies in the US and Europe. *Journal of Modern Transportation*, 23(1):30–42, 2015.
- [35] Amtrak. Route on-time performance, May 2016.

- [36] Bernhard Scholkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2001.
- [37] Alex Smola and S.V.N. Vishwanathan. *Introduction to Machine Learning*. Cambridge University Press, Cambridge, United Kingdom, 2008.
- [38] Harris Drucker, Christopher JC Burges, Linda Kaufman, Alex Smola, Vladimir Vapnik, et al. Support vector regression machines. *Advances in neural information processing systems*, 9:155–161, 1997.
- [39] Arthur E Hoerl and Robert W Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- [40] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.
- [41] Craig Saunders, Alexander Gammerman, and Volodya Vovk. Ridge regression learning algorithm in dual variables. In *ICML*, volume 98, pages 515–521, 1998.
- [42] Debasish Basak, Srimanta Pal, and Dipak Chandra Patranabis. Support vector regression. *Neural Information Processing-Letters and Reviews*, 11(10):203–224, 2007.
- [43] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
- [44] Christopher Burges. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2(2):121–167, 1998.
- [45] Bernhard E Boser, Isabelle M Guyon, and Vladimir N Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 144–152. ACM, 1992.
- [46] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.