2024 IEEE Intelligent Vehicles Symposium (IV)
June 2-5, 2024. Jeju Island, Korea

Characterizing Road Maps for Vehicle Endurance Testing with Machine Learning

Bijin Muthiyackal Abraham[1], Christian Drescher[2], Daniel Markert[1], Ana Perez Grassi[1] and Alejandro Masrur[1]

2024 IEEE Intelligent Vehicle Symposium (IV) | 979-8-3503-4881-1/24/$31.00 ©2024 IEEE | DOI: 10.1109/IV55156.2024.10588797

***Abstract*— To select optimal routes for vehicles endurance tests, it is necessary to have a road map characterized by events of interest. In this context, we define events as effects on the vehicle triggered by some proprieties of the routes. Such a characterization strongly relies on data from previous test drives. If new road maps are to be considered, e.g., in a different region, the route selection rather depends on the experience of engineers, which can lead to suboptimal decisions. To overcome this problem, we propose using the existing data from prior test drives to train a machine learning (ML) model, which then transfers this knowledge to unseen road maps. To this end, we formulate a sequential problem that can be solved with state-of-the-art ML architectures. Our experimental results based on real-world data show the potential of the proposed approach as we illustrate for the case of testing energy recuperation in electric vehicles.**

I. Introduction

Endurance testing is a critical part of the automotive development process. By testing a vehicle or a vehicle's component during extended use under real-world conditions, manufacturers can verify that their products meet high standards of durability, reliability and performance. In our experience, these tests can involve driving vehicles on public roads for extended periods of time, often hundreds of thousands of kilometers.

To ensure that a test project covers the stresses that vehicles will experience over their typical lifetime, test engineers are increasingly interested in extending the mere mileage accumulation with aggregated statistics of relevant events. In this context, events are defined as responses of the vehicle to certain features of the road, such as the occurrence of regenerative braking or kickdown (a downshift in automatic transmission during acceleration). To detect and measure such events, test vehicles are equipped with data loggers that collect data from sensors covering a range of attributes such as vehicle position, engine speed, energy recovery, etc. The collected data can be used to track the progress of the test project as well as to plan the route of future test drives (on the same roads) that maximizes the progress of a given test project. This involves solving the vehicle routing problem (VRP), which optimizes the chance of observing relevant event classes during the test. There is a large body of literature on VRPs and methods for solving them [1], e.g., using constraint programming [2]. As input, the VRP requires a road map characterized by the events of interest. For endurance testing, however, such data is limited to the roads and events covered in previous test drives. When much of the road map is unseen, i.e. not characterized by events, it is left up to experienced engineers to guess a suitable route for a test drive. This can, however, lead to suboptimal decisions that limit progress of the development project and negatively impact on the time to market.

[1]Department of Computer Science, Chemnitz University of Technology, Chemnitz, Germany
[2]Mercedes-Benz AG, Stuttgart, Germany

In contrast to relying on the engineers' subjective judgement, we introduce a data-driven approach to characterize road maps according to events of interest. We propose to train a machine learning (ML) model on available data from previous test drives and then apply it to regress the intensity of the learned events on unseen road maps. Finally, the regression results can be used to solve the VRP, which selects the optimal route on the given road map according to test-relevant events.

However, applying ML to event intensity regression using road maps is a non-trivial problem. The explicit graph structure of a road network makes this problem different from traditional ML tasks that consider vector-like or sequential data as input. One way to represent event intensity regression as an ML problem is to reduce it to an edge regression problem for graph neural networks [3]. Our work proposes a different approach by reformulating the event regression problem as a sequence problem. The contributions of this paper are threefold.

First, to represent the regression of event intensities as a sequence problem, we describe a route as a path through a graph (i.e., the road map). This graph consists of nodes and edges, where nodes are associated with real-world locations and edges between nodes denote roads connecting (adjacent) locations. In this way, a route is described by a sequence of segments.

Second, we propose a sampling technique to extract (low-level) features from routes, like distance or elevation, which allows us to use state-of-the-art ML architectures that operate on a vector-like representation. The resulting sequences of features from adjacent segments are used to train ML models as well as to regress their associated event intensities.

Third, we empirically evaluate our approach for the case of regressing values of energy recuperation caused by regenerative braking events. This is particularly relevant for testing electric vehicles, where the recovered energy can be used to extend the vehicle's range on the remaining charge. We train the models with real-world driving data collected during previous endurance tests. The presented experimental results demonstrate the applicability as well as some limitations of our approach.

The remainder of this paper is organized as follows:

Section II provides the necessary background. Section III introduces our approach. We report results of our experiments in Section IV. Then, we discuss our findings, including limitations and potential future work in Section V. Finally, the paper is concluded in Section VI.

II. Preliminaries

To address the challenge of regressing event intensities when data is limited, we present an approach based on road maps and ML techniques. The field of ML concerns algorithms that optimize the parameters of a (mathematical) model to fit data with respect to some cost function. This process is called learning or training an ML model. Supervised (machine) learning is the task of optimizing the parameters θ of a function $f_\theta(x)$ that maps input values x to a predicted values $\hat{y}$ based on sample pairs of input and ground-truth values (x, y) such that the dissimilarity (i.e., loss) between $\hat{y}$ and y is minimized. When the goal is to predict a continuous or numerical output value, the task is called regression.

In addition, when the input data consists of sequences of ordered observations (e.g., repeated measurements of a variable), and this order is crucial for finding and learning the underlying patterns to predict the output, the task is also referred to as a sequence problem. There are many types of models in the literature that can be used for supervised learning on sequence problems, including regression trees and artificial neural network.

A regression tree [4] is a directed tree-like model that can be constructed from data samples by recursively partitioning the samples according to the values of some selected attributes. Hence, the parameters θ of a regression tree model include the selection of attributes and their split points. Extremely randomized trees [5], for instance, randomize both attribute selection and split points in a strong manner. A regression tree model is applied to an unseen input by recursively following the partitioning (i.e., branches) that corresponds to the input's attribute values. The prediction made by the model is then the mean value of the partition where the recursion stops. A better generalization to unseen data can be achieved by combining multiple regression trees [6].

An artificial neural network is a directed acyclic graph-like model consisting of multiple layers of (computational) units called neurons. Each neuron in a layer has connections only to neurons in subsequent layers, providing information flow along the network. In the case of the first layer, each neuron takes one value of the input vector and gives it directly as output. The outputs of the neurons from all other layers (hidden layers and the last layer) are computed by a linear combination of their inputs passed through a nonlinear activation function. Finally, the outputs of the last layer provide the predicted values associated with the input data. The parameters θ of an artificial neural network include all the coefficients of the linear combinations computed by each neuron.

An artificial neural network can be trained by first defining its topology (or architecture), initializing θ with random values, and then adjusting θ based on pairs of sample inputs and their corresponding ground-truth values using a process called backpropagation [7]. Architectures that have been proven suitable for sequence problems often consider more specialized layers, including long short-term memory layers [8], time-dilated convolution layers [9], or self-attention layers [10].

Fig. 1. Illustration of a road map: Nodes correspond to geographic locations. Edges between nodes represent route segments and have varying distances or travel times associated.

In this work, the input data for the ML algorithms is extracted from road maps. Road maps encode unstructured data closely related to graphs. In practice, a road map is a visual representation of a real-world geographic area, showing the layout and arrangement of roads and landmarks, and other relevant data. We study road maps in terms of graphs representing a network of interconnected locations, as shown in Fig. 1. To this end, we apply basic definitions from graph theory, where a finite and directed graph with no loops or multiple edges is defined by finite sets of nodes and edges. A pair of consecutive connected nodes defines an edge and a sequence of successive connected nodes defines a path. In other words, a path provides a way to traverse the graph by moving from one node to another through connecting edges.

In summary, a road map can be conceptualized as a graph, where geographic locations are represented by nodes and route segments between locations are represented by edges with associated attributes (such as distance or travel time). Finally, a route is represented by a path in the graph.

In our application domain, the corresponding ground-truth values are obtained from previous test drives on some routes of a given road map. This means that the intensity of some event classes is only known for a few route segments in the road map and remains unknown for the others.

III. Event Intensity regression as a Sequence Problem

The goal of our work is to associate the route segments in a road map with intensity values of events of interest. To this end, we define a regression problem for supervised learning. Our work is based on the assumption that the intensity of events can be derived from structural features contained in road maps. For example, regenerative braking is typically observed when the speed limit is reduced or when approaching a sharp turn or an intersection. These low-level features can be extracted from the route data and an ML algorithm can abstract them to higher-level concepts such as energy recuperation.

However, state-of-the-art ML techniques are not directly applicable to network-like road maps. In fact, most models consider as input individual observations arranged in vectorial or sequential structures. For this reason, we propose first to fix the length of the explored routes according to the expected motion of the vehicle. These routes consist of sequences of nodes connected by route segments. Nodes are associated with attributes like geographic location including elevation and route segments with distance and travel time estimates.

These route attributes are the input of the regression model that estimates the event intensities as output. Many ML models are suitable for regression problems, but to generalize well to supervised learning tasks, the input data must be invariant to irrelevant transformations. In our application domain, road map data is invariant to translation and rotation. Thus, instead of using geographic locations of nodes to describe the motion of vehicles, we rather consider changes in absolute bearing and in elevation between nodes.

Since our regression problem involves data with inherent temporal dependencies, e.g., segments of a route are visited in a certain order, representing it as a sequence can better describe its nature. In turn, we can take advantage of ML architectures that are well suited for sequence problems, such as long short-term memory networks [8] or transformers [10].

Route segments have different lengths, which implies that their attributes are heterogeneous in terms of distance and travel time. However, ML models often require input data to be equidistant. For this reason, we resample the nodes along a route such that route segments represent fixed-length time intervals. Resampling involves interpolating the attribute values as well as the ground-truth values of consecutive route segments. Fig. 2 illustrates the resampling of a route, which results in a sequence of consecutive fixed-length segments. The resulting route features are then fed into the ML model in blocks, called *perceptive fields*, which are defined by a subset of consecutive route segments. The size of the perceptive field defines the portion of the route that is considered to regress the event intensity of one segment, which we call *target segment*.

The selection of the target segment within a perceptive field determines how much of the previous and the following route segments is included in the ML input data. The optimal ratio between previous and following route segments may vary depending on the event class to be regressed.

Fig. 2. Illustration of a route on a road map with a resampled mid-section. Square markers represent original nodes from the road map, while circular marker represent nodes resulting from resampling. Dashed edges represent original route segments, whereas solid edges represent resampled ones. In the depicted example, the perceptive field consists of 10 segments, where the target segment is the first segment in the perceptive field (thicker edge).

The routes on the road map are explored using a sliding window with the length of the perceptive field. Shifting this sliding window along the route results on sequences of regular segments. For each sequence the corresponding features are extracted and the event intensity of its target segment is estimated (with the possible exception of the segments at the beginning and at the end of the route, unless padding is applied).

The resampling frequency, the size of the perceptive field, and the position of the target segment within the perceptive field are hyperparameters of our regression sequence problem. An empirical analysis to derive their values is presented in the next section.

IV. Experiments

In this section, we report the results of our experiments to demonstrate the applicability and effectiveness of our proposed approach for the case of regressing energy recuperation derived from regenerative braking events. Our work is based on publicly available road maps and data collected from previous endurance test drives from the automotive domain. The data considered in our experiments stems from GPS traces and regenerative braking events of 512 hours of driving, day and night, on public roads, collected from three units of an electric vehicle model, recently introduced to the market.

Energy recovery has become a critical area of improvement, especially for energy-efficient electric motors, as the world transitions to sustainable mobility solutions. With the proliferation of energy recovery systems, regenerative braking events are growing in importance for electric vehicle endurance testing. These systems allow recovering energy that would normally be lost during deceleration or braking by converting the kinetic energy into electrical energy and storing it in a battery. This recovered energy can then be used to increase the range an electric vehicle can travel on the remaining charge.

Basically, our approach is based on the premise that the likelihood of regenerative braking events and the associated amount of energy recuperation is influenced by (static) route characteristics/attributes such as road profile (i.e., elevation changes), sharp turns, speed limit transitions, etc. as well dynamic attributes such as traffic and weather.

To consistently match GPS traces and regenerative braking events (from previous test drives) to route data, we used the match service from the Open Source Routing Machine (OSRM) project [11], which works on publicly available data from OpenStreetMap [12].

Further, to assign attributes to the route segments, we extracted travel time and distance data from OSRM. Additionally, we interpolated earth's surface elevation using publicly available data from Mapzen[1] [13]. Note that the elevation of the earth's surface and the elevation of the road are not always the same and may differ at tunnels and bridges where roads overcome geographical obstacles like rivers, mountains, etc. However, this can be compensated to some extent by preprocessing elevation data.[2] This is, however, out of the scope of this paper.

Route segments are then annotated with five low-level features: degrees turned left and right, elevation gained and lost, and driving speed. These features describe a three-dimensional motion and are invariant to absolute direction and absolute elevation.

Based on this, we applied state-of-the-art ML techniques to solve the sequence problem, which takes the five described features from the perceptive fields and outputs an estimate of the energy recuperation for each target segment. For training the model, the features are standardized to have a mean value of zero and a standard deviation of one, where we obtained ground-truth values from the energy recuperation measurements performed during the test drives.

Based on our initial tests, we define a baseline configuration consisting in resampling route data at one-second intervals, limiting the perceptive field to 20 seconds (at the expected speed of the vehicle), and predicting energy recuperation for the first route segment in the perceptive field (i.e., this is the target segment) yield satisfactory results.

We consider different perceptive field lengths of 5, 10, 20, 30, and 40 seconds (at the expected speed of the vehicle) to evaluate the impact on results. Then, to study the effect of the ratio between previous and following route segments on the regression results, we consider different locations of the target segment within the perceptive field. In particular, we divide the perceptive field in four same-length portions and select as target segment the first route segment (see again Fig. 2), the $\frac{1}{4}n^{\text{th}}$ segment, the $\frac{1}{2}n^{\text{th}}$ segment, the $\frac{3}{4}n^{\text{th}}$ segment, and the last segment of the perceptive field, where n denotes the size/length of the perceptive field, this time, given by the number of route segments in contains.

[1]https://mapzen.com/data/metro-extracts/

[2]For example, if there is an unbreachable elevation change for a vehicle on a particular road, it can be assumed that a tunnel or bridge has been built to keep elevation at an acceptable level.

Our experiments considered artificial neural networks using bi-directional long short-term memory layers (BiLSTM; [14]), time-dilated convolution layers (TCN: [9]), and multi-head self-attention layers (Transformer; [10]) as well as extremely randomized trees (ExtraTrees; [5]).

The BiLSTM setting considers an artificial neural network with two consecutive bidirectional long short-term memory layers, each producing sequences of 8-dimensional latent variables. The TCN setting uses eight time-dilated filters using dilations of 1, 2, 4, and 8 with no padding, and dropout of 50 percent on their output. The Transformer setting uses a learnable embedding and positional encoding of four dimensions, and stacks four transformer-style blocks of multi-head self-attention using four heads and two 1D-convolutions with eight/four dimensions each, applying layer normalization and dropout of 10 percent before adding residual connections between layers as described in [10]. BiLSTM, TCN, and Transformer were run for 10 epochs and a batch size of 256, using the Adam optimizer with the learning rate set to 0.001 and minimizing the mean squared error.

The ExtraTrees setting considers an ensemble of 100 extremely randomized trees, limiting tree-depth to 10, and using bootstrap-sampling.

The models are optimized to accurately estimate the value of energy recuperation during training. Our goal is to achieve sufficient generalization on the results, such that these serve as input to solve the VRP, which select the optimal route for testing energy recovery on an unseen road map. For this reason the evaluation of model effectiveness does not rely on the magnitude of prediction errors, but on the coefficient of determination R^2. R^2 represents the proportion of the total variability in the event data that is explained by the regression model by using only the proposed low-level features from the road map.

In order to assess how well a model generalizes to unseen road maps, we implemented our experiments using 10-fold cross-validation. To this end, we have divided the parts of the road map, in which there are at least 10 observations of the energy recuperation value into 10 subgraphs of roughly equal size. The route data extracted from each subgraph forms a cross-validation set. The distribution of energy recuperation values and route features will naturally vary between these road map partitions because they cover different geographic regions with different terrains, road types, and topologies. Therefore, we report predictive performance by collecting individual predictions from all cross-validation sets (as if predictions stem from a single experiment) instead of averaging performance metrics.

A summary of our experimental results is given in Fig. 3 and Fig. 4. In our baseline (hyperparameter) configuration (i.e., resampling route data at one second intervals, limiting the perceptive field to 20 seconds, and predicting energy recuperation value for the first route segment in the perceptive field), as well as in most of the other considered configurations, Transformer (43.6%) and TCN (42.5%) perform better than BiLSTM (36.7%) and ExtraTrees (32.6%). Although we have by far not exhausted the vast space of

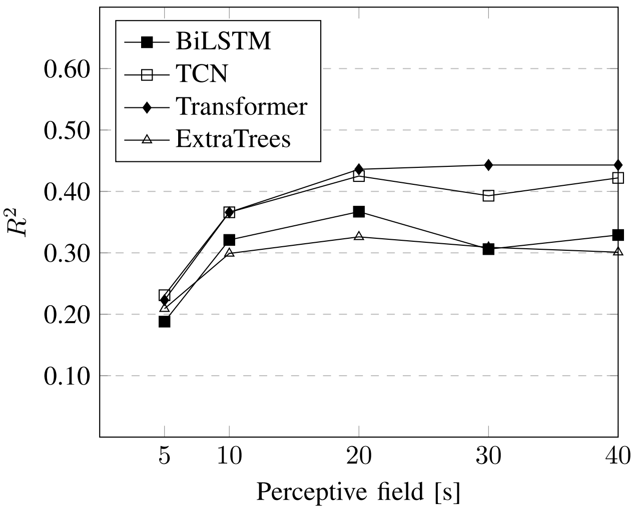

Fig. 3. Models performance vs. varying perceptive field length

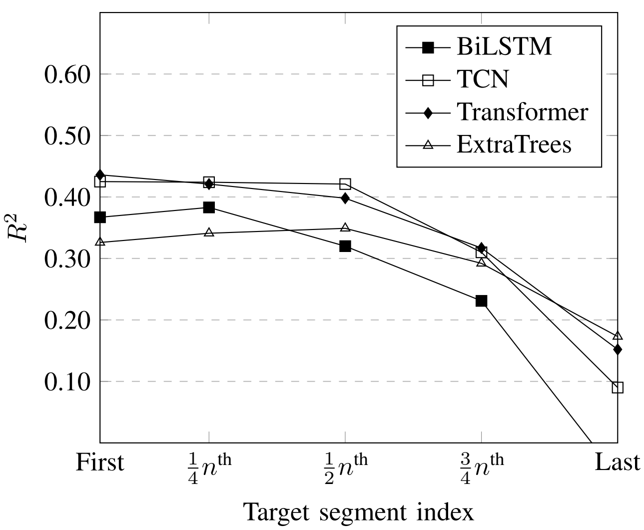

Fig. 4. Performance of models for varying target segment's positions

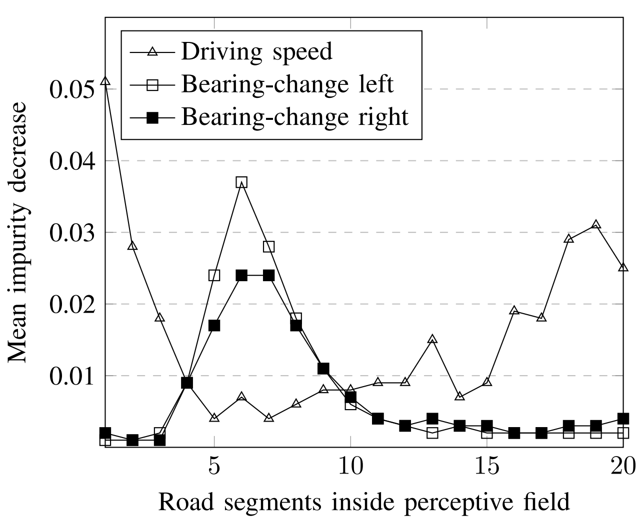

Fig. 5. Feature importance in baseline configuration using ExtraTrees

alternative settings and model-specific optimizations known in the literature, our results reflect the common understanding in the scientific community that TCN and Transformer are better suited for modeling sequence problems, even over LSTM, when there are long-range dependencies (due to the vanishing gradient problem inherent in LSTM).

A detailed analysis of model performance with varying perceptive field length (Fig. 3) shows that any performance gains saturate at 20 seconds for most models, with only negligible gains when using a larger perceptive field in the Transformer setting. Fixing the perceptive field size at 20 seconds and varying the target segment position within the sequence (Fig. 4) also yields no gains.

The best setting in our experiments is given by Transformer using a perceptive field length of 40 seconds when fixing the target segment to, i.e., predicting energy recuperation for, the first route segment within perceptive field. According to our results, this can explain 44.3% of the variability in energy recuperation by only using the proposed low-level features extracted from road map data.

Albeit of limited importance to our application scenario, we can identify the most relevant attributes contributing to predicting energy recuperation, e.g., by studying the mean impurity decrease of the low-level features within each regression tree in the ExtraTree setting. Using our baseline configuration, we observe that (1) bearing-changes during the route segments 4-10s beyond the target segment, and (2) driving speeds near the beginning and the end of the perceptive field appear indicative of energy recuperation at the target segment (Fig. 5).

Finally, there are also dynamic attributes that depend on traffic and weather, which cannot be predicted from the road map data alone. However, as mentioned before, our models were trained with road maps annotated with real-world measurement data from previous test drives. Such annotations implicitly contain traffic and weather information and, hence, the trained models learn to correlate this information with the low-level features we defined before. Even though a separate modeling of traffic and weather would certainly improve results, it can be expected that similar roads are affected by traffic and weather in a similar way, which the proposed approach accounts for to some extent.

V. Discussion

The experiments conducted in our study yield several key findings. First, our approach can predict energy recuperation on unseen routes to a reasonable extent, demonstrating applicability and practical relevance.While a relatively modest R^2 value of slightly above 40% suggests that the prediction model does not explain a large portion of the variability in the data, it is a quite reasonable result in our application scenario. Moreover, these results can be further improved as more data becomes available with future test drives.

Even now, regression results can be used as decision aid when selecting route segments based on events of interest, i.e., for which no data is available from previous test drives yet. This way, a considerable amount of engineering work can be saved, reducing testing costs. We illustrated this for the case of energy recovery during endurance tests. In particular, those route segments that are estimated to have significant event intensity may be relevant for VRP.

On the other hand, there are several factors that can cause regression errors. Some of these can potentially be included in related future work, but are beyond the scope of our benchmark application. First, as already mentioned, our approach does not directly account for dynamic attributes such as traffic (including traffic signals) and weather, but only indirectly. For individual test drives and some route segments, this can lead to significant differences between the observed and predicted event intensity.

Further, human factors such as the driving style were also neglected in this work. For example, a more dynamic driving style typically results in a higher number of regenerative braking events and higher energy recovery. Event intensity regressions may, hence, differ depending on the driver.

Data quality may also have affected model performance. Recall that our study considered the elevation of the earth's surface, but not the elevation of the road. While we assume that the data is consistent for most route segments, bridges or tunnels are exceptions. Even though this can be solved by preprocessing elevation data, there may still be a deviation from reality. A potential solution is to let test vehicles collect elevation data as well.

Another challenge is matching GPS data to the road map due to GPS inaccuracies caused by signal obstructions, atmospheric conditions, and receiver noise. In some cases, this can result in event measurements being assigned to the wrong route segments. Equipping test vehicle with more accurate positioning systems may solve this issue, however, it will again take a considerable amount of time to collect suitable data.

Finally, in addition to feature selection and data quality, the choice of hyperparameters can also affect model performance. For sequential problems, changes in sampling frequency can have a significant impact on the performance of an ML model. This has been demonstrated, for example, for crowd movement prediction using LSTMs [15] and for motor imagery classification [16] among others.

VI. Conclusions

The optimal routing of endurance tests drives is a laborious and time-consuming task that can in principle be automated by vehicle routing problem (VRP) solvers. In practice, however, much of the road map is not characterized by relevant data. In this work, we have presented a novel, data-driven method to characterize a road maps according to events of interest. Such a road map can then be used to solve the VRP, which selects the route that maximizes the observation of relevant events for the purpose of endurance test drives.

We propose the use of machine learning (ML) to abstract event intensity values from (low-level) route features. Using graph theory, we rearrange the information extracted from a road map and represent it as a sequence problem, so that state-of-the-art ML models (such as BiLSTM, TCN, Transformer, and ExtraTrees) can handle it directly. Experimental results for the case of regressing energy recuperation values demonstrate the practical applicability of the proposed approach. Moreover, our approach can be applied to other event classes, e.g., kickdown (downshifts in automatic transmission triggered by pressing down the throttle pedal), provided that sufficient training data is available.

As future work we propose to reformulate our approach for graph neural networks [3], including graph convolutional networks [17] and graph attention networks [18], which could provide a better alternative to learn graph structures. Future work could also include more low-level features like traffic signals from OpenStreetMap.

References

[1] P. Toth and D. Vigo, *Vehicle Routing: Problems, Methods, and Applications*. SIAM, 2014.

[2] P. Shaw, "Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems," in *International Conference on the Principles and Practice of Constraint Programming CP*, 1998, pp. 417–431.

[3] M. Gori, G. Monfardini, and F. Scarselli, "A New Model for Learning in Graph Domains," in *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, vol. 2. IEEE, 2005, pp. 729–734.

[4] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone, *Classification and Regression Trees*. Wadsworth and Brooks, 1984.

[5] P. Geurts, D. Ernst, and L. Wehenkel, "Extremely Randomized Trees," *Machine Learning*, vol. 36, no. 1, pp. 3–42, 2006.

[6] Z.-H. Zhou, *Ensemble Methods: Foundations and Algorithms*. Chapman & Hall/CRC, 2012.

[7] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning Representations by Back-propagating Errors," *Nature*, vol. 323, no. 6088, pp. 533–536, 198.

[8] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[9] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, "WaveNet:A Generative Model for Raw Audio," in *The 9th ISCA Speech Synthesis Workshop*, 2016, p. 125.

[10] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is All You Need," in *31st International Conference on Neural Information Processing Systems NIPS*, 2017, pp. 6000–6010.

[11] D. Luxen and C. Vetter, "Real-Time Routing with OpenStreetMap Data," in *19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems GIS*. ACM, 2011, pp. 513–516.

[12] OpenStreetMap contributors, "Planet Dump Retrieved from https://planet.osm.org," https://www.openstreetmap.org, 2017.

[13] Mapzen, "Terrain tiles on AWS," accessed on June 2nd 2023 from https://registry.opendata.aws/terrain-tiles.

[14] A. Graves, S. Fernández, and J. Schmidhuber, "Bidirectional LSTM Networks for Improved Phoneme Classification and Recognition," in *15th International Conference on Artificial Neural Networks ICANN*. Springer, 2005, pp. 799–804.

[15] J. Casals, M. Jerez, and S. Sotoca, "Modelling and Forecasting Time Series Sampled at Different Frequencies," *Journal of Forecasting*, vol. 28, pp. 316–342, 2009.

[16] T.-J. Luo, C.-L. Zhou, and F. Chao, "Exploring Spatial-Frequency-Sequential Relationships for Motor Imagery Classification with Recurrent Neural Network," *BMC Bioinformatics*, vol. 19, p. 344, 2018.

[17] T. N. Kipf and M. Welling, "Semi-Supervised Classification with Graph Convolutional Networks," in *5th International Conference on Learning Representations ICLR*, 2017.

[18] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph Attention Networks," in *6th International Conference on Learning Representations ICLR*, 2018.