

Approximating Optimisation Solutions for the Travelling Officer Problem with Neural Networks

Wei Shao
School of Science
RMIT University
Melbourne, Australia
wei.shao@rmit.edu.au

Jeffrey Chan
School of Science
RMIT University
Melbourne, Australia
jeffrey.chan@rmit.edu.au

Flora D. Salim
School of Science
RMIT University
Melbourne, Australia
flora.salim@rmit.edu.au

Abstract—Deep learning has been extended to a number of new domains with critical success, though some traditional orienteering problems such as the Travelling Salesman Problem (TSP) and its variants are not commonly solved using such techniques. Deep neural networks (DNNs) are a potentially promising and under-explored solution to solve these problems due to their powerful function approximation abilities, and their fast feed-forward computation. In this paper, we outline a method for converting an orienteering problem into a classification problem, and design a multi-layer deep learning network to approximate traditional optimisation solutions to this problem. We test the performance of the network on a real-world parking violation dataset, and conduct a generic study that empirically shows the critical architectural components that affect network performance for this problem.

Index Terms—deep learning, optimisation, function approximation

I. INTRODUCTION

The travelling salesman problem (TSP) is a well-known combinatorial optimisation problem in the artificial intelligence research community, with extensive real-world examples and applications [1]. As a classical, NP-hard problem, it has been studied extensively in both theory and applications [2]. As a result of its ubiquity, many efficient optimisation algorithms have been proposed for approximating solutions, such as ant colony optimisation, [3], greedy selection [4], and genetic algorithms [5]. However, in today’s era of big data, this classic formulation is not as suitable for many real-world applications which often include additional contextual and historical information [6]. Although traditional methods for solving the TSP can be used in such domains, we might hope to gain an advantage by developing new methods – that do take contextual and historical data into account – for solving such problems.

The travelling officer problem (TOP) is a variant of travelling salesman problem (TSP) [7] but provides a way to use contextual and historical data. Nowadays, parking violation has become a prominent challenge for administration in most big cities. Parking officers need to stick infringement notices on violating cars before they leave the parking zone, but this can be challenging for many reasons. Firstly, the majority of infringing vehicles leave within a short period. Secondly, many

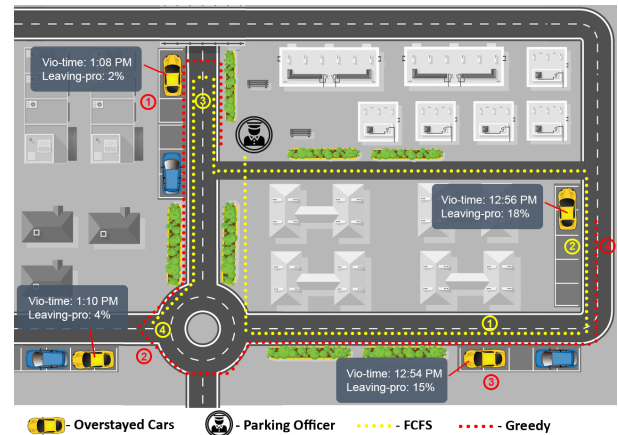


Fig. 1: The Travelling Officer Problem – a parking officer faces four parking violation in different directions. The parking officer knows the time-in-violation for each car, and estimates the probability of each car leaving. The parking officer needs to choose a path that catches as many of these violators as possible. Yellow line in the map shows the First-Come-First-Serve (FCFS) solution. The red line illustrates the probability-based greedy solution.

violation events occur at the same time in a large area. As shown in Figure 1, parking officers must balance the travelling time between the officer’s location relative to the infringing vehicles, and the probability that these cars will leave.

Shao et al. [7] previously defined this problem, and two heuristic solutions were demonstrated for generating paths using spatio-temporal data collected from on-ground sensors in parking spaces. The path generated by these optimisation methods (E.g. Red line in Figure 1) were shown to be better than the First-Come-First-Serve solutions (E.g. Yellow line in Figure 1), and achieved a higher return on parking fines when evaluated on the real-world parking dataset [7]. Despite this, the efficiency of these optimisation solutions cannot satisfy the requirements of real-time application. Given the dynamic nature of parking events [8]–[11] and short-term nature of parking violations, this makes application of more traditional methods difficult.

Deep learning models can achieve real-time performance

since their training and inference are broken up into two distinct sessions. Training can be done offline, and inference runs in constant time, making neural networks ideal for problems like the TOP where fast evaluation is desired. However, a significant portion of current classification approaches tend to solve supervised learning problems for which the solution is known and provided as a training dataset. In addition, the TOP is a typical optimisation problem, and deep learning models are not typically used to solve optimisation problems directly. In order to leverage deep learning models, the TOP problem needs to be transformed into a supervised classification problem in order to obtain good potential paths as labels.

To overcome the above challenges, we use solutions generated by optimisation methods as labels for training. We propose a spatio-temporal data segmentation approach to transform the optimisation problem into a classification problem, and design a deep feed-forward neural network to approximate optimisation solutions. Training deep neural networks to replace optimisation has many advantages in this case: the computationally expensive optimisation problem can be solved as part of the training session, and once this is done, the test session can roll out trajectories that approximate those of the original optimiser with simple feed-forward computation. Moreover, neural networks scale very well, allowing such a technique to take advantage of a huge amount of additional contextual and temporal information and explore the unclear structure of this data.

There are challenges in both the transformation and learning tasks. Traditional orienteering problems only focus on the spatial domain, usually in the form of a static 2D graph. The TOP also considers the temporal domain, which consists of many temporal views of this same spatial graph. We need a method that works with both spatial and temporal information. Secondly, classification methods need the input features and the corresponding labels; the TOP does not provide any existing solutions or features. It is non-trivial to integrate the optimisation solution with classification to solve this problem. Finally, it is unclear how to effectively use deep neural networks to approximate the optimisation solution in orienteering problems. As a result, the design of an appropriate deep learning architecture for this type of problem is needed.

This paper explores the use of deep learning techniques for the solution of the TOP, and makes the following contributions:

- We propose a framework to solve travelling officer problem incorporating optimisation approaches and deep neural networks.
- We propose a novel segmentation method to transform a spatio-temporal graph into a sequence of features.
- We explain the reason that DNN can approximate the greedy and other optimisation solutions.
- We validate our claim that the TOP can be solved using a combination of both optimisation solutions and neural network classifiers through extensive experiments with a large real-world dataset, including a comparison with traditional machine learning methods.

II. RELATED WORK

There is an extensive body of research in applying neural networks to TSP variants going back to [12], though the networks used in these studies typically fall into the category of Hopfield networks, and self-organising feature map networks (see [13]–[15]). Hopfield networks are fully recurrent, and memorise training examples by minimising an energy cost function. However, many recent advances in deep learning have been made with feed-forward neural networks, on the back of better optimisation algorithms and the ability to train on large datasets. Though Hopfield networks have been used for classification tasks, their performance is not as good as modern deep learning techniques. On combinatorial tasks, the number of neurons required by a Hopfield network scales with n^2 (where n is the number of nodes in the graph) which can be problematic for larger graphs. Our study is different from these previous works in two ways: firstly, we are approximating a solution to a more difficult, time-dependent version of the TSP (the TOP). Our problem is focused on maximising a temporally-dependent reward, rather than navigating a geographically-fixed set of nodes. Using different temporal views of the data to generate training samples can dramatically increase the size of the training set (this is inspired by [16]). Secondly, rather than using Hopfield networks, we propose re-framing the problem as a supervised learning task for classification. Under this framework, we use an optimisation algorithm to generate the training set for a classifier, which is then trained to generate a trajectory through a given graph.

Additionally, there numerous recent works in effectively combining optimisation with deep learning. Fischetti and Jo modelled deep neural networks as a 0-1 mixed integer linear program [17]. Galassi et al. used a deep neural net to learn the structure of a combinatorial problem, and mentioned that such research is still at an early stage [18]. Ma et al [19] propose a Graph Pointer Networks trained using reinforcement learning (RL) for solving the TSP. Bello et al. presents a framework to tackle combinatorial optimisation problems using neural networks and reinforcement learning [20]. Dai et al. presents an end-to-end machine learning framework for automatically designing greedy heuristics for hard combinatorial optimisation problems on graphs which inspires our work as well [21].

III. BACKGROUND

A. Overview of On-street Parking in City

The Melbourne Transportation Council set up in-ground sensor systems around CBD areas. For each car parking area, the sensor can detect parking space availability, and check its violation state with parking rules. The sensors report parking events to information centre periodically, and the system would send the message to the patrolling parking officer who supervises this area when the nearby parking cars are in violation state. Then the patrolling parking officer will go to check the cars in violation, and dispatched to issue an infringement notice.

B. Parking Sensor Data

The parking events data recently has been published online, which attracts researchers to study and analyse. The parking events data were recorded from October 1, 2011, to September 30, 2012 (12 months). A total number of 12,208,178 records were logged. Each record comprises the information of a parking event including area name, street name, street segment and some other parking information. It also provides the spatial information such as the latitude and longitude of the parking spaces.

Figure 2 shows the distribution of violations with the length of time in violation. The horizontal axis denotes the length of time from the beginning of the violation to the car leaving, and the vertical axis shows the total violation numbers within a month. Figure 2 shows that there are only a small part of cars in violation beyond 100 minutes, and most violations between five and 60 minutes. This is the key observation that we take into account in our spatio-temporal-based optimization approach. Using such observations we build a probabilistic model that can estimate the likelihood of a car in violation leaving which allows us to optimise the strategy that officers use to catch cars in violation.

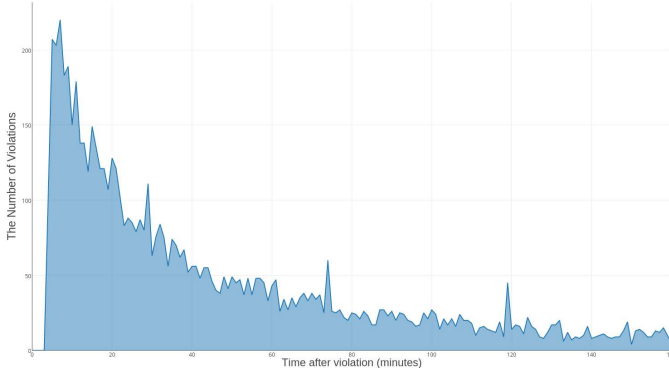


Fig. 2: The numbers of violation with total violation time

C. Travelling Officer Problem

The Travelling Officer Problem describes the problem of a parking officer traversing a fully connected graph to maximise a cumulative reward (in this case, parking violations). There is a time cost $C(u, w)$ associated with travelling from node $u \in \mathbf{v}$ to $w \in \mathbf{v}$ (or parking lot u to w , where \mathbf{v} denotes all nodes in the graph or all parking lots in the area) that is dependent on the officer's walking speed (we do not assume that the officer stops at intermediary nodes).

The officer must choose between chasing for the potential reward for catching a parking violation at a given node by considering the probability that the violation may no longer exist by the time he/she arrives, and the opportunity of cost saving from not travelling to other nodes containing parking violations. The solution of the TOP aims to find a path \mathbf{S} that maximises the number of valid nodes with time limits (e.g. working hours), and the time-varying state of each node. The

valid nodes denote the car at parking lot x_i is in a state of violation.

Let T be the total travelling budget, and $f_{j,t} \in \{0, 1\}$ denotes whether there is an infringement at node j at time t . We denote a solution $\mathbf{S} = \{(x_1, t_1), (x_2, t_2), \dots, (x_{|\mathbf{S}|}, t_{|\mathbf{S}|})\}$ as the path travelling over nodes, where $x_i \in \mathbf{v}$ and denote the i_{th} node in the path, and $t_i \in Time$ denote the time when officer arrives at the node x_i (note Time is whatever time unit/division you are using). Because in the TOP, t_i is deterministic from the path of nodes visited, we can infer t_i from just the visited nodes, and we will simplify our path \mathbf{s} to $x_1, x_2, \dots, x_{|\mathbf{S}|}$. Let R denote the infringement fine amount (assuming each infringement cost the same). In this paper, we assume the R is a constant value.

Then the TOP problem is to find a path \mathbf{S} that maximises the total return, satisfying the total travelling time budget. A formal definition of this problem is as follows:

$$\begin{aligned} & \arg \max_{\mathbf{S}} \sum_{x_i \in \mathbf{S}} f_{i,t} \cdot R \\ \text{s.t.} & \sum_{i=1}^{|\mathbf{S}|-1} C(x_i, x_{i+1}) \leq T \quad (1) \\ & t_k = \sum_{i=2}^k C(x_i, x_{i+1}), \text{ for } 1 < k < |\mathbf{S}| \quad (2) \\ & t_1 = 1 \quad (3) \end{aligned}$$

D. Heuristic Optimisation

Previously, Shao et al. [7] discussed two heuristic optimisation methods (greedy and ACO) to solve the Travelling Officer Problem; it was shown that both algorithms performed well at the task of collecting parking violation fines. In order to take advantage of temporal information in the TOP, the authors proposed a dynamic temporal probability model and integrated it with traditional optimisation methods. The proposed greedy algorithm can be formalised as a single function as follows:

$$\arg \max_i \exp \left(-\frac{\tau_i + \frac{d_i}{V}}{\alpha} \right), \quad (4)$$

where τ_i denotes the overstayed time of cars at node x_i , and d_i denotes the route distance between node x_i and the current position of the parking office. V is a constant to denote the speed of the parking officer, and α is a parameter which is set by historical data analysis. The proposed greedy algorithm seek for i_{th} nodes calculated by Eq. 4 as the next position for parking officer.

In the ACO algorithm, the ants decide the next node by the pheromone distribution left by previous ants. The probability of these choices was modelled as being proportional to $[\phi(x_{uw})]^\alpha \bullet [\eta(x_{uw})]^\beta$, where η is the probability of a node being invalid by the time the officer arrives. The greedy algorithm used a similar dynamic probability model to estimate the most promising node, and would greedily select the best one. The details of both algorithms are shown in [7].

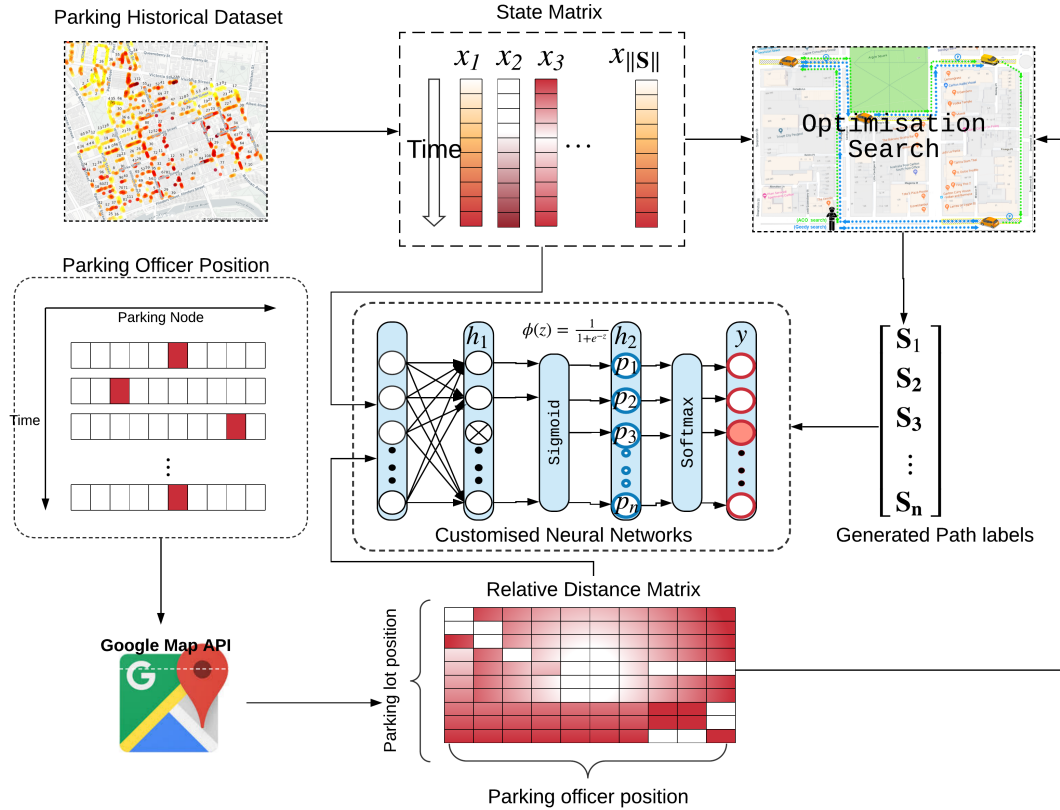


Fig. 3: Flowchart of the proposed model. The state matrix is extracted from parking event data. The officer position matrix is randomly generated, and a relative distance matrix which measures the distance between the parking officer and other parking lots is calculated using the Google Maps API. From there, the state matrix and relative distance matrix are fed into the optimiser to generate labels for the classifier. These labels are then used for training, and the classifier is used to roll out trajectories on the test set.

IV. METHODOLOGY

In this section, we propose a framework to transform the TOP – a spatio-temporal orienteering problem – into a classification problem that can be solved by a feed-forward deep neural network incorporating optimisation solutions. Figure 3 illustrates the overview of our framework. It primarily consists of three parts: spatio-temporal feature extraction, optimisation-based search, and a feed-forward neural network. Spatio-temporal feature extraction aims to extract features and build a training dataset from the public parking violation historical dataset. Classification labels are generated by using optimisation to find a path from any given map. The DNN then uses the state and relative distance as the input features, and uses the optimisation solution S as the ground truth label to learn the optimisation algorithm. We outline the details of each component in the following subsections.

A. Spatio-temporal Feature Extraction Method

A significant difference between traditional sub-path planning problems and the TOP is temporal dependence. Traditional sub-path problems generally consist of spatial data such as the location of the vertices, or the edges of the graph [22], [23]. For such static graphs, information is limited; however, time is continuous, meaning that even finer-grained slices can be taken to generate further data points when the temporal dimension is included. Each time frame becomes a static 2-D graph at timestamp t_i , corresponding to a data sample χ_i .

we extract a vector $\chi_i = \{\chi_{i,x_1}, \chi_{i,x_2}, \dots, \chi_{i,x_n}\}$ in Figure 4, where n is the number of parking lots (the number of nodes in the graph) from parking violation events at time t_i and slicing time frame every Δ_t . For example, at node x_j , there are p violation events $\{(\tau_{x_j,1}^a, \tau_{x_j,1}^e), (\tau_{x_j,2}^a, \tau_{x_j,2}^e), \dots, (\tau_{x_j,p}^a, \tau_{x_j,p}^e)\}$, where $\tau_{x_j,k}^a$ is

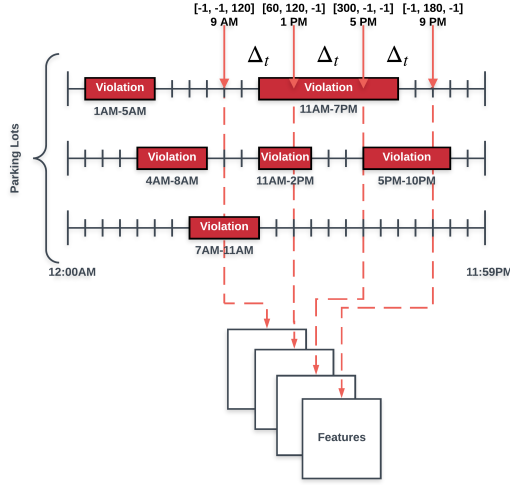


Fig. 4: Spatio-temporal feature extraction method: The bottom part represents the parking violation events in temporal domain. The red bar indicates the violation interval. The red arrows slices the state information of each node with a fix-length time step. Inf denotes that there is no violation in the specific parking lot.

the start time of the k_{th} violation interval at node x_j , and $\tau_{x_j,k}^e$ is the end time of the k_{th} violation interval at node x_j . For any time t_i , there only two cases: 1) t_i is located between one of violation events interval. That is, $\tau_{x_j,k}^a < t_i < \tau_{x_j,k}^e$ and $1 \leq k \leq p$, or 2) time t_i is not located in any violation events. In this first case, we set the value $\chi_{i,x_j} = t_i - \tau_{x_j,k}^a$. In the second case, We set the value as -1 if there is no violation for parking lot x_j at time t_i . The number of time lots is denoted as $m = \frac{T}{\Delta_t}$, where T is the maximum time constraint and Δ_t denotes the time step we choose for slicing the temporal domain.

Except for temporal features, we also extract spatial features. The first component in the second row of Figure 3 shows a officer position matrix, where position filled with red colour in each row denotes the current location of the parking officer at time t_i . For simplicity, the officer will not change the path between two nodes. Therefore, the possible position of officers are the same as the parking lots positions. The relative distance matrix shown in the bottom row in the Figure 3 store the distance between current parking officer location and other nodes $x_j \in \mathbf{v}$. Therefore, for any time t_i , we can get a relative distance vector $\mathbf{d} = \{d_1, d_1, \dots, d_n\}$, where d_j measure the route distance between the parking officer and node x_j .

B. Optimisation

The TOP is an NP-hard problem that can be solved by traditional optimisation methods under an objective function and constraints. As mentioned previously, optimisation solution \mathbf{S} is series of nodes at time t_i . However, a complete path should be broken into multiple decisions to transform an optimisation problem to a classification problem. That is, at any time t_i ,

we need to run the optimisation algorithm and only select the first node x_1 in the path \mathbf{S} . As a result, we can get a series of labels $\mathbf{l} = \{x_1^{t_1}, x_1^{t_2}, \dots, x_1^{t_m}\}$, where $x_1^{t_i}$ is the first node of path \mathbf{S} which is given by optimisation method at time t_i .

The classification problem aims to learn a categorical likelihood over a set of classes from a training set. In this case, if we regard each node x_i as a class label and use state matrix and relative distance matrix as the features for the training, the classifier should learn to choose the next node.

C. Neural Networks

In order to achieve similar performance with optimisation methods, it is not enough to use an existing general neural network because the general neural networks are sensitive to the hyper-parameters and architecture. Therefore, we propose a specific neural network which is designed with each operation of the greedy algorithm in TOP. Figure 5 shows the completed architecture of neural networks designed for replacing the greedy algorithm in TOP. The first layer is the input layer concatenated by two vectors: a relative distance vector \mathbf{d} , and the state vector χ_i which is calculated in Section IV-A. The first hidden layer h_1 aims to learn a linear combination of \mathbf{d} and χ_i . If we pay attention to Eq. 4, we can find that this layer can approximate the function $h_1(\mathbf{d}_i, \tau_i) = -\frac{\tau_i + \frac{d_i}{V}}{\alpha}$. Universal approximation theorem [24] states that a feed-forward network with a single hidden layer containing a finite number of neurons can approximate continuous functions on compact subsets of \mathbf{R}^n . This is a continuous function on a real value dataset. Therefore, we design a hidden layer h_1 to approximate this function. Then we use a sigmoid function which is also a non-linear activation function to approximate exponential function. This is also a non-linear function [25] [26]. Layer h_2 now consists probabilities p_i to denote the capture chance by parking officer for node x_i . Since $\arg \max$ is not a continuous function, we cannot use the hidden layer. Fortunately, softmax function is a perfect function to choose the max value from $\mathbf{p} = \{p_1, p_2, \dots, p_n\}$ [27]. Therefore, the output of the network becomes the next node that the officer should travel to.

For other components of the neural network, we used Adam [28] for optimisation, and early stopping for regularisation. We used dropout at each layer to prevent overfitting [29].

V. EXPERIMENTS AND RESULTS

In this section, we evaluate the performance of our proposed model, and compare it with traditional classification methods and optimisation-only solutions on a real-world dataset. The rules and assumptions are outlined in [7].

A. Dataset

We tested the proposed model on the Melbourne parking event dataset, published by the Melbourne City Council, and used previously in [30] [7] [31] [32]. A detailed description of this dataset is included in [7]. We took time slices throughout the week for training, and tested the performance of the classifier using different time slices to ensure that the test




Fig. 5: Architecture of greedy-like DNN. Input layer consists of relative distance vector and state vector for each parking node. Output layer is a one-hot vector denotes the next travelling node.

and train sets were drawn from the same distribution. For reward evaluation, we randomly chose a week’s worth of data from a year-long data set. For other experiments, we randomly selected a single day. For the rewards study, we chose 50 nodes from all vertices in the graph, and extracted sampling data at 10 second intervals. the dataset consists of all parking events in the CBD’s on-street car parking bays over a year.

To apply our model to the dataset, we used six attributes extracted from the dataset (as shown in Table I), and a list of attributes we defined from that dataset (Table II).

Distance is the most important attribute that we use to calculate the cost and find the solution in our model. Since we conducted experiments in real scenarios, it was important to measure the distance between two points on the map. Therefore, we extracted records that reflected accurate positions and driving distance (the distance calculated on the street path using Google Maps) between car parking bays.

B. Evaluation Metric

We use two criteria to evaluate the performance of our proposed method: rewards, and classification accuracy. The definition of rewards is given in [7]. It denotes how many cars in violation can be caught by parking officers. Since we use the optimisation solution as the ground truth, we also use the classification accuracy to measure the degree to which the classifier learns the optimisation algorithm.

C. Experimental settings

We applied both ACO and greedy algorithm which is used in Shao et al. [7] to the dataset and the DNN, Support vector machine (SVM) [33] and Random Forest (RF) [34] to learn the optimisation solutions.

D. Classification Model Comparison

In the first experiment, we evaluate all classifiers over a week by rewards and categorical accuracy. There are two sub-set of experiments. First experiment compares the greedy algorithms and classifiers that learn from the greedy algorithm.

Second experiments evaluates the ACO and classifiers that learn from ACO.

Figure 6 shows the weekly rewards obtained by optimisation solutions – greedy and ACO. It also shows achieved rewards from different classification methods learned from the optimisation solutions. the DNN outperformed the other techniques. Interestingly, we found that classification was more accurate on weekends compared to weekdays because weekends average a lower number of violations than weekdays, largely due to less stringent parking rules. Overall, The DNN achieves similar performance to the greedy on this problem as expected.




Fig. 6: Weekly rewards obtained through different algorithms.

E. Evaluation of Model Components

We also studied the effect of varying the parameter settings of the problem, as measured by classification accuracy and rewards achieved:

- Number of nodes: the number of nodes indicates the depth of the search space in orienteering problems. In this case, the number of nodes is also associated with the number of rewards, and the size of the training set;
- Minimum time step: we extract training samples from the dataset with a time step. For each time step t_i , we extract a temporal image from the dataset and add it to our training set.

We evaluated our classifiers on graphs with sizes varying from 10 to 50 nodes. Figure 7 shows that the classification accuracy drops significantly as the size of the graph is increased, which suggests that classification methods imitate optimisation solutions well in smaller search spaces. It is possible that in larger search spaces, the limited number of training examples is the limiting factor preventing the classifiers from being able to learn better an approximation of the optimisation routine. Though the rewards increase with the number of nodes in the problem, this is likely the result of greater potential rewards due to the presence of more nodes in the graph. Notably, the gap in performance between the optimisation solution and the classification solutions become larger for these larger problems. The DNN outperformed both the SVM and RF

TABLE I: Attribute list for parking events

Attribute Name	Description
Street Marker	The signs placed on the side of parking bays
Area Name	City area, - used for administrative purposes
Arrive Time	Time that the sensor detected a vehicle over it
Departure Time	Time that the sensor detected a vehicle is leaving
In Violation	Indicates that the parking event exceeded the legally permissible time
Sign	Parking sign in effect at the time of the parking event.

TABLE II: Other attributes used in experiments

Attribute Name	Description
Location	The longitude and latitude of the parking slot
Violation Time	The violation time based on signs
Violation Period	The length of time that car overstays




Fig. 7: The effect of the number of nodes n on the total reward obtained, and the categorical accuracy as seen on the test set.

solutions by increasing its total reward along with the graph size. In contrast, the SVM and RF performance dropped under the same conditions. Interestingly, we find that all classifiers achieve higher accuracy when they learn from greedy than ACO. This phenomenon may be caused by the complexity of the optimisation solution. This is planned for future study.

We also evaluated the model on data from a single day with varying time step sizes from 10 to 60 seconds. Figure 8 shows that smaller time steps resulted in better overall accuracy on the validation set, but this did not necessarily translate to better rewards. This is potentially because smaller time steps provided more training data. However, it does not suggest that testing accuracy is higher.




Fig. 8: The effect of the time step Δt on the total reward obtained, and the approximation accuracy as seen on the test dataset.

F. Computational Complexity Analysis

Finally, we evaluated the computational efficiency of both the optimisation method and the neural network, by varying the number of nodes from 10 to 50 and then measuring the execution time of the program. Figure III shows that the running time for DNN is much faster than the optimisation algorithm since we exclude the training time. We only consider the test session time as the running time of DNN because the training session can be done offline. That is, we can use historical data to train the DNN before using it. Optimisation methods cannot be applied to the real scenario before we know it. Therefore, it is fair to compare the testing time of DNN and running time of the optimisation methods.

It shows that testing time is significantly shorter than other algorithms and does not change with the number of nodes. Therefore, DNN-based model is much more efficient than optimisation solution in the real-world scenario.

VI. DISCUSSION AND FUTURE WORK

The neural network performed reasonably consistent across all days, though in general, it fared worse compared to the greedy algorithm during the week. Notably, we found that when the performance of the greedy algorithm was low, the gap between the neural network and the optimisation algorithm was very small.

In this paper, we only introduce one kind of DNN for the greedy algorithm which may not suitable for ACO and other optimisation algorithms such as greedy algorithm is one of the simplest optimisation algorithms. Our next goal is to design different neural networks architectures for different optimisation methods and generalise current DNN to more existing optimisation problems.

VII. CONCLUSION

A technique was shown for reformulating an orienteering problem as a classification problem, and using conventional optimisation to generate labels for training. We took finer time slices of the dataset to increase the amount of training data, and this was shown to improve accuracy. We design a neural network to approximate the greedy algorithm in TOP. We evaluated on a large real-world dataset by sampling at a different time interval to generate a distinct test set. It was shown that the DNN could be used to approximate the greedy algorithm in TOP.

TABLE III: Time consuming for different algorithms with different number of nodes. (seconds)

Methods	node 10	node 20	node30	node 40	node50
Greedy	7.12	15.71	21.78	28.05	33.96
ACO	112.16	217.60	232.41	427.52	528.08
DNN (Training Greedy)	77.12	152.6	244.49	349.57	443.17
DNN (Testing Greedy)	0.24	0.57	1.04	1.33	1.72
DNN (Training ACO)	74.89	152.42	232.60	323.39	405.49
DNN (Testing ACO)	0.27	0.34	0.44	0.49	0.53

REFERENCES

- [1] H. Sumita, Y. Yonebayashi, N. Kakimura, and K.-i. Kawarabayashi, "An improved approximation algorithm for the subpath planning problem and its generalization," in *26th International Joint Conference on Artificial Intelligence, IJCAI 2017*. International Joint Conferences on Artificial Intelligence, 2017.
- [2] S. Sahni and T. Gonzalez, "P-complete approximation problems," *Journal of the ACM (JACM)*, vol. 23, no. 3, pp. 555–565, 1976.
- [3] M. Dorigo and L. M. Gambardella, "Ant colony system: a cooperative learning approach to the traveling salesman problem," *IEEE Transactions on evolutionary computation*, vol. 1, no. 1, pp. 53–66, 1997.
- [4] J. Edmonds, "Matroids and the greedy algorithm," *Mathematical programming*, vol. 1, no. 1, pp. 127–136, 1971.
- [5] H. Braun, "On solving travelling salesman problems by genetic algorithms," in *International Conference on Parallel Problem Solving from Nature*. Springer, 1990, pp. 129–133.
- [6] Y. Mei, F. D. Salim, and X. Li, "Efficient meta-heuristics for the multi-objective time-dependent orienteering problem," *European Journal of Operational Research*, vol. 254, no. 2, pp. 443–457, 2016.
- [7] W. Shao, F. D. Salim, T. Gu, N.-T. Dinh, and J. Chan, "Travelling officer problem: Managing car parking violations efficiently using sensor data," *IEEE Internet of Things Journal*, 2017.
- [8] M. S. Rahaman, Y. Ren, M. Hamilton, and F. D. Salim, "Wait time prediction for airport taxis using weighted nearest neighbor regression," *IEEE Access*, vol. 6, pp. 74 660–74 672, 2018.
- [9] M. S. Rahaman, M. Hamilton, and F. D. Salim, "Queue context prediction using taxi driver knowledge," in *Proceedings of the Knowledge Capture Conference*. ACM, 2017, p. 35.
- [10] M. S. Rahaman, M. Hamilton, and F. D. Salim, "Predicting imbalanced taxi and passenger queue contexts in airport," in *PACIS*, 2017, p. 172.
- [11] S. Ahmed, M. S. Rahman, and M. S. Rahaman, "A blockchain-based architecture for integrated smart parking systems," in *2019 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. IEEE, 2019, pp. 177–182.
- [12] J. J. Hopfield and D. W. Tank, "Neural computation of decisions in optimization problems," *Biological cybernetics*, vol. 52, no. 3, pp. 141–152, 1985.
- [13] B. F. La Maire and V. M. Mladenov, "Comparison of neural networks for solving the travelling salesman problem," in *Neural Network Applications in Electrical Engineering (NEUREL), 2012 11th Symposium on*. IEEE, 2012, pp. 21–24.
- [14] S. Abdel-Moetty, "Traveling salesman problem using neural network techniques," in *Informatics and Systems (INFOS), 2010 The 7th International Conference on*. IEEE, 2010, pp. 1–6.
- [15] J.-Y. Potvin, *The traveling salesman problem: a neural network perspective*. Université de Montréal, Centre de recherche sur les transports, 1992.
- [16] Y. Liu, Y. Zheng, Y. Liang, S. Liu, and D. S. Rosenblum, "Urban water quality prediction based on multi-task multi-view learning," *25th International Joint Conference on Artificial Intelligence, IJCAI 2016*, 2016.
- [17] M. Fischetti and J. Jo, "Deep neural networks as 0-1 mixed integer linear programs: A feasibility study," *arXiv preprint arXiv:1712.06174*, 2017.
- [18] A. Galassi, M. Lombardi, P. Mello, and M. Milano, "Model agnostic solution of cpsps via deep learning: A preliminary study," in *International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research*. Springer, 2018, pp. 254–262.
- [19] Q. Ma, S. Ge, D. He, D. Thaker, and I. Drori, "Combinatorial optimization by graph pointer networks and hierarchical reinforcement learning," *arXiv preprint arXiv:1911.04936*, 2019.
- [20] I. Bello, H. Pham, Q. V. Le, M. Norouzi, and S. Bengio, "Neural combinatorial optimization with reinforcement learning," *arXiv preprint arXiv:1611.09940*, 2016.
- [21] E. Khalil, H. Dai, Y. Zhang, B. Dilkina, and L. Song, "Learning combinatorial optimization algorithms over graphs," in *Advances in Neural Information Processing Systems*, 2017, pp. 6348–6358.
- [22] M. S. Rahaman, Y. Mei, M. Hamilton, and F. D. Salim, "Capra: A contour-based accessible path routing algorithm," *Information Sciences*, vol. 385, pp. 157–173, 2017.
- [23] M. S. Rahaman, M. Hamilton, and F. D. Salim, "Coact: A framework for context-aware trip planning using active transport," in *2018 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. IEEE, 2018, pp. 645–650.
- [24] B. C. Csáji, "Approximation with artificial neural networks," *Faculty of Sciences, Eötvös Loránd University, Hungary*, vol. 24, p. 48, 2001.
- [25] Y. Ito, "Approximation of functions on a compact set by finite sums of a sigmoid function without scaling," *Neural Networks*, vol. 4, no. 6, pp. 817–826, 1991.
- [26] S. Ferrari and R. F. Stengel, "Smooth function approximation using neural networks," *IEEE Transactions on Neural Networks*, vol. 16, no. 1, pp. 24–38, 2005.
- [27] M. Tokic and G. Palm, "Value-difference based exploration: adaptive control between epsilon-greedy and softmax," in *Annual Conference on Artificial Intelligence*. Springer, 2011, pp. 335–346.
- [28] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [29] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 2014.
- [30] W. Shao, F. D. Salim, A. Song, and A. Bouguettaya, "Clustering big spatiotemporal-interval data," *IEEE Transactions on Big Data*, vol. 2, no. 3, pp. 190–203, 2016.
- [31] K. K. Qin, W. Shao, Y. Ren, J. Chan, and F. D. Salim, "Solving multiple travelling officers problem with population-based optimization algorithms," *Neural Computing and Applications*, pp. 1–27, 2019.
- [32] W. Shao, Y. Zhang, B. Guo, K. Qin, J. Chan, and F. D. Salim, "Parking availability prediction with long short term memory model," in *International Conference on Green, Pervasive, and Cloud Computing*. Springer, 2018, pp. 124–137.
- [33] A. J. Smola and B. Schölkopf, "A tutorial on support vector regression," *Statistics and computing*, vol. 14, no. 3, pp. 199–222, 2004.
- [34] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.