Long-Distance Electric Vehicle Navigation using a Combinatorial Semi-Bandit Approach

Niklas Åkerblom
Volvo Car Corporation
Gothenburg, Sweden
Chalmers University of Technology
Gothenburg, Sweden
niklas.akerblom@chalmers.se

Morteza Haghir Chehreghani
Chalmers University of Technology
Gothenburg, Sweden
morteza.chehreghani@chalmers.se

Abstract

In this work, we address the problem of long-distance navigation for battery electric vehicles (BEVs), where one or more charging sessions are required to reach the intended destination. We consider the availability and performance of the charging stations to be unknown and stochastic, and develop a combinatorial semi-bandit framework for exploring the road network to learn the parameters of the queue time and charging power distributions. Within this framework, we first outline a method for transforming the road network graph into a graph of feasible paths between charging stations to handle the constrained combinatorial optimization problem in an efficient way. Then, for the feasibility graph, we use a Bayesian approach to model the stochastic edge weights, utilizing conjugate priors for the one-parameter exponential and two-parameter gamma distributions, the latter of which is novel to multi-armed bandit literature. Finally, we apply combinatorial versions of Thompson Sampling, BayesUCB and Epsilon-greedy to the problem. We demonstrate the performance of our framework on long-distance navigation problem instances in large-scale country-sized road networks, with simulation experiments in Norway, Sweden and Finland.

1 Introduction

In the coming years, it will be crucial for society to shift the transport sector (personal and commercial) towards electrification, to reach global targets on reduced greenhouse gas emissions. *Range anxiety* is still a major obstacle to the widespread adoption of battery electric vehicles (BEVs) for transportation. This phenomenon can be characterized as the fear that (potential or current) BEV drivers might feel about exceeding the electric range of their vehicle before reaching either their destination or a charging station, thus being stranded with an empty battery.

There is another, perhaps less commonly discussed, but potentially more relevant issue called *charging anxiety*. Whereas the range of typical electric vehicles, while still shorter than that of combustion engine vehicles, has been steadily increasing, the act of charging the battery is still, often, a cumbersome task. Even at locations with fast chargers, various factors may severely impact the total travel time of a particular trip. At the highest possible charging power, charging a BEV battery from nearly empty to close to maximum capacity may take more than 30 minutes. However, maximum power might not always be provided, in practice. Furthermore, queues to charging stations may appear due to the (relatively) long charging times and few charging locations. Issues like these might diminish public trust in BEVs as viable alternatives to combustion engine vehicles.

In this work, we attempt to mitigate any charging anxiety arising due to the aforementioned factors by developing an online self-learning algorithmic framework for navigation of BEVs, capable of taking

16th European Workshop on Reinforcement Learning (EWRL 2023).

such charging issues into account. We view the task as a *sequential decision-making problem under uncertainty* and model it as a *combinatorial semi-bandit problem* to address the trade-off between exploring charging stations to learn more information about them and exploiting previously collected knowledge to select charging stations that are likely to be good. Within this framework, we employ a Bayesian approach, with conjugate priors novel to bandit literature. To our knowledge, our work is the first study that addresses the challenging real-world problem of charging station selection in the partial information setting using multi-armed bandit (MAB) methods. Thereby, our work provides a novel framework to develop and investigate advanced multi-armed (combinatorial) bandit methods. As mentioned, our work is also one of the few large-scale real-world applications of multi-armed bandits, specifically in the combinatorial setting. To achieve such scalability, we transform the road graph into a feasibility graph, where feasible paths between charging stations are pre-computed to improve run-time efficiency. Such a transformation of the problem instances for the purpose of computational efficiency is novel to the MAB community.

2 Related work

Several works have studied shortest path algorithms to address the problem of energy efficient navigation, e.g., [5, 39] focusing on minimizing energy consumption, and [7], studying how to minimize travel time while ensuring that battery energy is not fully depleted (utilizing charging stations, if necessary). The authors of [41] outline a Dynamic Programming (DP) approach for adaptive routing of electric vehicles, and model charging station availability and queue / waiting times, but assume that all distributions are known in advance. A similar work [20] formulates charging station selection as a stochastic search problem, addressed with a DP-based approach. BEV navigation problems (including charging station selection) have been modelled as detailed reinforcement learning problems [31, 36], but often with high computational demands making them infeasible for long-distance navigation. To our knowledge, this problem has not been studied in the partial information setting using multi-armed bandits, and our work is the first contribution of this kind. Prior works on multi-armed bandits for energy-efficient navigation, e.g., in [2, 3], are significantly simplified without considering travel time and charging. Dealing with charging, in particular in a computationally efficient way, requires considerably more sophisticated models and methods, beyond the existing applications of (combinatorial) multi-armed bandit methods, as will be demonstrated in this paper.

In general, the multi-armed bandit problem is a versatile way of describing how to utilize limited resources to balance exploration of an environment to gain new knowledge and usage of previously collected knowledge to increase long-term reward. Thompson Sampling [42] is an early algorithm attempting to address this trade-off, which has recently increased in popularity due to demonstrated experimental performance [19, 11] and proven theoretical performance guarantees [1, 29, 9, 37].

Another type of method commonly used for sequential decision-making problems is the Upper Confidence Bound (UCB) [6] algorithm. Like Thompson Sampling, this method has been adapted to many different settings, including combinatorial optimization problems [12]. UCB methods have also been used for MAB problems with sub-exponential rewards [25], e.g., for selection of bike rental companies with exponential service times, but not in combinatorial settings as far as we are aware.

3 Model

In this section, we describe how to represent the road network as a graph and how to transform it into a graph of feasible paths between charging stations to allow for computationally efficient charging station selection. Furthermore, we outline an approach for probabilistic modelling of the queue time and charging power of each charging station.

3.1 Road network graph

We model the *underlying* road network using a directed and weighted graph $\mathcal{G}^{\text{road}}\left(\mathcal{V}^{\text{road}}, \mathcal{E}^{\text{road}}, \boldsymbol{\tau}^{\text{road}}\right)$. Each vertex $u \in \mathcal{V}^{\text{road}}$ corresponds to *either* an intersection or some other important location of the road network (e.g., a charging station). Each directed edge $e \in \mathcal{E}^{\text{road}}$ represents a road segment from a location $u \in \mathcal{V}^{\text{road}}$ to another location $u' \in \mathcal{V}^{\text{road}}$, which we may also indicate by writing $e = (u, u')$. We further denote the travel time of each road segment $e \in \mathcal{E}^{\text{road}}$ as τ_e^{road}, and the

vector of all edge travel times as $\boldsymbol{\tau}^{\text{road}}$. Additionally, each road segment $e \in \mathcal{E}^{\text{road}}$ has an associated energy consumption $\varepsilon_e^{\text{road}}$, which is the energy needed by a given vehicle to traverse the complete road segment. A convention we follow throughout the paper is to indicate vectors with bold symbols, with individual elements indexed by edge or vertex subscripts. Furthermore, we let $\mathcal{V}^{\text{charge}} \subseteq \mathcal{V}^{\text{road}}$ be the set of locations which contain charging stations. Each charging location $u \in \mathcal{V}^{\text{charge}}$ is associated with a maximum charging power $\varrho_u^{\max}$ which the charging station is able to provide. We also assume that there is a corresponding value for the minimum charging power $\varrho_u^{\min}$ available at each station. The actual charging power is denoted by $\varrho_u^{\text{charge}}$.

A connected sequence $\boldsymbol{p}$ of edges (or vertices, equivalently) is called a *path* through the graph. Given a source vertex $u^{\text{src}} \in \mathcal{V}^{\text{road}}$ and a target vertex $u^{\text{trg}} \in \mathcal{V}^{\text{road}}$ (both assumed to be fixed and known), we denote the set of all paths starting in u^{src} and ending in u^{trg} as $\mathcal{P}^{\text{road}}_{(u^{\text{src}},u^{\text{trg}})}$. Assuming that we aim to find a path which minimizes the total travel time, we let, for each edge $e \in \mathcal{E}^{\text{road}}$, the *edge weight* be τ_e^{road}. Then, the *shortest path problem*, given u^{src}, u^{trg} and $\mathcal{G}^{\text{road}}$ is defined as

$$\boldsymbol{p}^* = \arg \min_{\boldsymbol{p} \in \mathcal{P}^{\text{road}}_{(u^{\text{src}},u^{\text{trg}})}} \left(\sum_{e \in \boldsymbol{p}} \tau_e^{\text{road}} \right), \tag{1}$$

which may be addressed by one of several classical methods, e.g., Dijkstra's algorithm [15], the Bellman-Ford algorithm [40, 16, 8] or the A* algorithm [23]. The A* algorithm, in particular, can be described as a *best-first search* method (see e.g., [14]), where a provided heuristic function is used to guide the algorithm towards promising solutions. An admissible heuristic function should be able to provide an underestimate of the total weight of any path between a pair of given vertices. For a road network graph with travel time edge weights, such as $\mathcal{G}^{\text{road}}$, we can use a function which calculates a travel time value based on the maximum allowed speed in the road network and the beeline distance between the two vertices. When a good heuristic function is used, the A* algorithm is computationally more efficient than Dijkstra's algorithm, while still guaranteeing that the optimal path is found.

3.2 Construction of feasibility graph

The model described in the previous section is sufficient for many applications. Since fossil fuel stations are ubiquitous in most road networks, and since the time required for refueling is typically negligible, the model can be used for navigation of combustion engine vehicles without significant modifications. For BEVs, however, charging can take more than 30 minutes, and multiple charging sessions may be required for longer trips. These factors, combined with the relative sparsity of the charging infrastructure, means that charging should not be disregarded in the navigation problem.

The time spent on charging depends on the amount of energy needed and the charging power provided. Furthermore, queues may occur if all charging stations at a particular location are occupied at the same time. In this work, for simplicity, we assume that each charging session has to fully charge the battery. In principle, it is possible (and may be time optimal) with partial charging, but this significantly increases the computational complexity of the problem. We also assume that the paths between charging stations should be chosen to minimize travel time, even if there are alternative paths with less energy consumption. For clarity, throughout this work when the battery is stated to be either empty or fully charged, the battery state of charge is actually 10% or 80%, respectively, for safety and durability reasons.

The general *resource-constrained shortest path problem* [26] (with energy as the resource) is still computationally hard, especially when resource replenishment (i.e., charging) is considered. However, with these assumptions, it is possible to transform the road graph into a *feasibility graph*, where feasible paths between charging stations are pre-computed to improve run-time efficiency. We denote this directed and weighted graph $\mathcal{G}^{\text{feasible}}\left(\mathcal{V}^{\text{feasible}}, \mathcal{E}^{\text{feasible}}, \boldsymbol{\tau}^{\text{feasible}}\right)$.

We simply let $\mathcal{V}^{\text{feasible}} = \mathcal{V}^{\text{charge}}$ be the set of charging stations. Then, for any given path $\boldsymbol{p}$ through $\mathcal{G}^{\text{road}}$, let $\tau_{\boldsymbol{p}}^{\text{road}} = \sum_{e \in \boldsymbol{p}} \tau_e^{\text{road}}$ be the travel time of the path and $\varepsilon_{\boldsymbol{p}}^{\text{road}} = \sum_{e \in \boldsymbol{p}} \varepsilon_e^{\text{road}}$ be the total energy consumption of the path. We create a new set of edges $\mathcal{E}^{\text{path}} = \left\{(u, u') \in \mathcal{V}^{\text{feasible}} \times \mathcal{V}^{\text{feasible}}\right\}$, where each edge $(u, u') \in \mathcal{E}^{\text{path}}$ corresponds to the shortest path $\boldsymbol{p}^*_{(u,u')} = \arg\min_{\boldsymbol{p} \in \mathcal{P}^{\text{road}}_{(u,u')}} \tau_{\boldsymbol{p}}^{\text{road}}$ between the charging stations, such that we have $\tau^{\text{path}}_{(u,u')} = \tau^{\text{road}}_{\boldsymbol{p}^*_{(u,u')}}$ and $\varepsilon^{\text{path}}_{(u,u')} = \varepsilon^{\text{road}}_{\boldsymbol{p}^*_{(u,u')}}$. Finally, given a

maximum battery capacity $\varepsilon^{\max}$ and minimum battery capacity $\varepsilon^{\min}$, we define the set of feasible edges as the set of shortest paths between charging stations where the battery capacity exceeds the energy consumption, i.e., $\mathcal{E}^{\text{feasible}} = \left\{ e \in \mathcal{E}^{\text{path}} \mid \varepsilon_e^{\text{path}} \leq \varepsilon^{\max} - \varepsilon^{\min} \right\}$.

Moreover, we define a vector of travel times for the edges of the feasibility graph, $\boldsymbol{\tau}^{\text{feasible}}$. For each edge $(u, u') \in \mathcal{E}^{\text{feasible}}$, we let $\tau_{(u,u')}^{\text{feasible}} = \tau_{(u,u')}^{\text{path}} + \tau_{u'}^{\text{queue}} + \tau_{(u,u')}^{\text{charge}}$, where the charging time $\tau_{(u,u')}^{\text{charge}} = \varepsilon_{(u,u')}^{\text{path}} / \varrho_{u'}^{\text{charge}}$ depends on both the energy consumed on the edge (u, u') and the provided charging power at u', while we assume that the queue time $\tau_{u'}^{\text{queue}}$ only depends on u'.

We note that, even though we assume that each charging session fully charges the battery, the construction of the feasibility graph (and the entire online learning framework presented in this work) can be extended in a straightforward way to allow for partial charging. If the state of charge is discretized into a finite number of *levels*, vertices can be added for these to each charging station, where edges between them represent partial charging choices. Then, feasibility graph *layers* may be computed for each of the state of charge levels. See, e.g., [41], for another approach using discretized charging levels for partial charging, applied to a setting with fixed and known parameters.

3.3 Probabilistic queue and charging times

As stated earlier, we consider both the queue time and the charging power of each charging station to be stochastic and unknown, only to be revealed after the station has been visited. In contrast, we assume that we are given the travel time and energy consumption of each road segment in the road network graph (and, in practice, that they are fixed). We further assume that the queue time and charging power are independently distributed, both with respect to each other, as well as between different charging stations. In reality, they exhibit a complex interdependence, where low charging power might cause queues to appear, and the simultaneous charging of many vehicles may cause the available power to decrease.

3.3.1 Queue time model

The queuing behavior at a particular charging location may be complex, depending on the characteristics of the location. A charging location has few or many charging stations, where each may have multiple connectors. The stations may also differ in the maximum charging power provided, as well as the price of charging. These, and other factors, impact the preferences of drivers towards different stations, especially if many of the stations at the same location are occupied simultaneously.

Rather than modelling the queues in detail, we take inspiration from a simple model of queuing theory, the *M/M/1 queue* [30], and assume that the queue time τ_u^{queue} of each charging station $u \in \mathcal{V}^{\text{feasible}}$ is exponentially distributed according to an unknown rate parameter λ_u^{queue}. The likelihood function of the queue time model can then be defined as

$$P\left(\tau_u^{\text{queue}} | \lambda_u^{\text{queue}}\right) = \text{Exp}(\lambda_u^{\text{queue}}). \tag{2}$$

We also take a Bayesian view, and assume that the rate parameter is drawn from a known prior distribution. In principle, any suitable (positive support) distribution can be used as prior, but for this likelihood and parameter, the gamma distribution is a conjugate prior (meaning that the posterior distribution given observations is also a gamma distribution, and thereby the posterior parameters can be efficiently computed). The prior is then given by

$$P\left(\lambda_u^{\text{queue}} | \alpha_{u,0}^{\text{queue}}, \beta_{u,0}^{\text{queue}}\right) = \text{Gamma}\left(\alpha_{u,0}^{\text{queue}}, \beta_{u,0}^{\text{queue}}\right). \tag{3}$$

Given a sequence of observed queue times $y_1, \ldots, y_t$, the parameters $\alpha_{u,t}^{\text{queue}}$ and $\beta_{u,t}^{\text{queue}}$ of the gamma posterior distribution are given by $\alpha_{u,t}^{\text{queue}} = \alpha_{u,0}^{\text{queue}} + t$ and $\beta_{u,t}^{\text{queue}} = \beta_{u,0}^{\text{queue}} + \sum_{i=1}^{t} y_i$. Similarly, incremental updates of the parameters can be performed using $\alpha_{u,t}^{\text{queue}} = \alpha_{u,t-1}^{\text{queue}} + 1$ and $\beta_{u,t}^{\text{queue}} = \beta_{u,t-1}^{\text{queue}} + y_t$.

3.3.2 Charging power model

Ideally, which is also often the case, any given charging station $u \in \mathcal{V}^{\text{feasible}}$ should be able to provide the specified maximum charging power $\varrho_u^{\max}$. Occasionally, however, some charging stations provide

less power. Reasons for this may include, e.g., intermittent high load in the surrounding electric grid, limitations of the charging station, etc. Moreover, in this work, we assume that the vehicle is able to fully utilize the charging power provided by the charging station.

While a Gaussian model could be sufficient for the anomalous cases described here, it would have to be truncated or rectified to represent the sharp peak in density at $\varrho_u^{\max}$ for a charging station functioning as intended. Then, conjugacy properties may not be used for efficient posterior parameter updates. An alternative, which we describe here, is to use a gamma distribution to model the charging power. In practice, we still rectify the charging power distribution below $\varrho_u^{\min}$ in Section 4 to prevent negative or zero charging power, but this should have a relatively minor impact on the results since the density of the charging power distribution is often concentrated close to $\varrho_u^{\max}$ with the prior distributions that we consider. We define the likelihood function as

$$P\left(\varrho_u^{\max} - \varrho_u^{\text{charge}} | \alpha_u^{\text{charge}}, \beta_u^{\text{charge}}\right) = \text{Gamma}\left(\alpha_u^{\text{charge}}, \beta_u^{\text{charge}}\right). \tag{4}$$

A conjugate prior distribution for both parameters of the gamma likelihood was derived by [13] and further analyzed by [33], which [13] refers to as the *Gamcon-II prior*. The joint prior distribution over α_u^{charge} and β_u^{charge} has a set of parameters $\pi_{u,0}^{\text{charge}} > 0$, $\gamma_{u,0}^{\text{charge}} > 0$ and $\xi_{u,0}^{\text{charge}} > 0$, where $\sqrt[\xi_{u,0}^{\text{charge}}]{\pi_{u,0}^{\text{charge}}} < 1$. Decomposed, the conjugate prior over β_u^{charge} conditional on α_u^{charge} is also a gamma distribution, defined as

$$P\left(\beta_u^{\text{charge}} | \alpha_u^{\text{charge}}, \gamma_{u,0}^{\text{charge}}, \xi_{u,0}^{\text{charge}}\right) = \text{Gamma}\left(\xi_{u,0}^{\text{charge}} \cdot \alpha_u^{\text{charge}}, \gamma_{u,0}^{\text{charge}}\right). \tag{5}$$

Whereas the prior distribution over β_u^{charge} has a convenient form for both sampling and moment computation, only the unnormalized probability density function for the marginal conjugate prior distribution over α_u^{charge} is available. It is defined as

$$\begin{aligned} P\left(\alpha_u^{\text{charge}} | \pi_{u,0}^{\text{charge}}, \gamma_{u,0}^{\text{charge}}, \xi_{u,0}^{\text{charge}}\right) \propto \exp\Big(&\alpha_u^{\text{charge}} \ln \pi_{u,0}^{\text{charge}} - \xi_{u,0}^{\text{charge}} \alpha_u^{\text{charge}} \ln \gamma_{u,0}^{\text{charge}} \\ &- \xi_{u,0}^{\text{charge}} \ln \Gamma\left(\alpha_u^{\text{charge}}\right) + \ln \Gamma\left(\xi_{u,0}^{\text{charge}} \alpha_u^{\text{charge}}\right)\Big), \end{aligned} \tag{6}$$

where $\Gamma(\cdot)$ is the well-known gamma function. The joint unnormalized prior distribution over α_u^{charge} and β_u^{charge} is then the product of Eq. 5 and Eq. 6. With an observed charging power z_t, the incremental updates for the parameters of the joint posterior are given by $\pi_{u,t}^{\text{charge}} = \pi_{u,t-1}^{\text{charge}} \cdot (\varrho_u^{\max} - z_t)$, $\gamma_{u,t}^{\text{charge}} = \gamma_{u,t-1}^{\text{charge}} + (\varrho_u^{\max} - z_t)$ and $\xi_{u,t}^{\text{charge}} = \xi_{u,t-1}^{\text{charge}} + 1$. Despite lacking a normalization constant, Eq. 6 can be used to efficiently find the mode of the posterior, since it is log-concave on the (positive real) domain. An unnormalized density function can also be used in adaptive rejection sampling methods to efficiently generate exact samples from the posterior distribution.

4 CMAB formulation

We formulate the problem of selecting paths and charging stations through the road network as a sequential decision-making problem under uncertainty. Specifically, we see it as a *combinatorial semi-bandit* (CMAB) problem [10, 17], a variant of the classical *multi-armed bandit* (MAB) problem. For a finite horizon T, and each iteration $t \in [T]$, the agent has to select and execute an action. In the CMAB setting, this action consists of a subset of objects from a *ground set*. Often, there are constraints on which subsets are allowed to be selected by the agent. The environment gives feedback for each of the objects selected (called *semi-bandit feedback*), the set of which determines the *reward* received by the agent for taking the action.

In our setting, the ground set corresponds to the set of edges in the feasibility graph, i.e., $\mathcal{E}^{\text{feasible}}$. For a source vertex $u^{\text{src}} \in \mathcal{V}^{\text{feasible}}$ and a target vertex $u^{\text{trg}} \in \mathcal{V}^{\text{feasible}}$, fixed for a particular problem instance, the set of allowed actions corresponds to the set of paths $\mathcal{P}^{\text{feasible}}_{(u^{\text{src}}, u^{\text{trg}})}$ from the source to the target in the feasibility graph. In each iteration $t \in [T]$, the agent selects and travels a path $\boldsymbol{p}_t \in \mathcal{P}^{\text{feasible}}_{(u^{\text{src}}, u^{\text{trg}})}$, and receives the path travel time $\tau^{\text{path}}_{(u,u')}$, queue time $\tau^{\text{queue}}_{u'}$ and charging time $\tau^{\text{charge}}_{(u,u')}$ as feedback for each edge $(u, u') \in \boldsymbol{p}_t$. Since the shortest path problem is a minimization problem, we say that an action

has a *loss* instead of a reward, where the loss of the travelled path is $L_t(\boldsymbol{p}_t) = \sum_{(u,u')\in\boldsymbol{p}_t} \tau^{\text{feasible}}_{(u,u')}$, where $\tau^{\text{feasible}}_{(u,u')} = \tau^{\text{path}}_{(u,u')} + \tau^{\text{queue}}_{u'} + \tau^{\text{charge}}_{(u,u')}$. Let $\boldsymbol{\theta}$ be an arbitrary vector of model parameters for the entire feasibility graph, where for each vertex $u \in \mathcal{V}^{\text{feasible}}$ we let $\boldsymbol{\theta}_u = \left(\lambda^{\text{queue}}_u, \alpha^{\text{charge}}_u, \beta^{\text{charge}}_u\right)$. Then, we define the *expected loss function* of a path $\boldsymbol{p}$ as

$$f_{\boldsymbol{\theta}}(\boldsymbol{p}) = \sum_{(u,u')\in\boldsymbol{p}} \left(\tau^{\text{path}}_{(u,u')} + \frac{1}{\lambda^{\text{queue}}_{u'}} + g_{\alpha^{\text{charge}}_{u'},\beta^{\text{charge}}_{u'}}(u,u')\right), \tag{7}$$

where $g_{\alpha^{\text{charge}}_{u'},\beta^{\text{charge}}_{u'}}(u,u') = \mathbb{E}\left[\varepsilon^{\text{path}}_{(u,u')}/\max\left(\varrho^{\min}_{u'}, \varrho^{\text{charge}}_{u'}\right) \mid \alpha^{\text{charge}}_{u'}, \beta^{\text{charge}}_{u'}\right]$ is the expected charging time, given the parameters $\alpha^{\text{charge}}_{u'}$ and $\beta^{\text{charge}}_{u'}$. Here, the charging power in the denominator is rectified below the minimum charging power $\varrho^{\min}_{u'}$. Throughout this work, $g_{\alpha^{\text{charge}}_{u'},\beta^{\text{charge}}_{u'}}(u,u')$ is approximated as $\hat{g}_{\alpha^{\text{charge}}_{u'},\beta^{\text{charge}}_{u'}}(u,u')$ using Monte Carlo estimation, by averaging over charging time values computed with samples from the charging power model defined in Eq. 4.

MAB algorithms are usually evaluated using the notion of *regret* until a horizon T, which is defined as the sum over all iterations $t \in [T]$ of the difference in expected loss of the best action $\boldsymbol{p}^*$ (defined as in Eq. 1, but for the feasibility graph) and the action $\boldsymbol{p}_t$ selected by the algorithm, such that

$$\text{Regret}(T) = \sum_{t\in[T]} \left(f_{\boldsymbol{\theta}^*}(\boldsymbol{p}_t) - f_{\boldsymbol{\theta}^*}(\boldsymbol{p}^*)\right), \tag{8}$$

where $\boldsymbol{\theta}^*$ is the true underlying parameter vector (in which the parameters of the queue time and charging power distributions are assumed to be drawn from their respective prior distributions). The objective is to find a policy which minimizes the expected regret, where a sub-linear growth with respect to T is generally desired.

5 CMAB methods

We adapt three CMAB algorithms for our problem setting: Epsilon-greedy, Thompson Sampling and BayesUCB. For all three algorithms, a shortest path algorithm is used to find the shortest path through the feasibility graph. This is usually called an *oracle* in CMAB literature. While Dijkstra's algorithm is a commonly used oracle for CMABs with shortest path problems (see e.g., [17, 32, 44, 2]), we may use the more efficient A* algorithm since the feasibility graph admits a suitable heuristic function. Since the vector of path travel times $\boldsymbol{\tau}^{\text{path}}$ is fixed and known, the direct (beeline) distance between each pair of charging stations can be divided by the maximum allowed speed in the road network (e.g., 120 km/h) to get a value which is guaranteed to underestimate the travel time between those stations. We do not explicitly consider the queue time or charging time in the heuristic function, but clearly, both are non-negative and implicitly underestimated by zero.

All three algorithms follow the same general structure, as outlined in Algorithm 1, which closely corresponds to the CMAB description in Section 4, while also including explicit posterior parameter updates and other details. The primary bottleneck in the computational efficiency of Algorithm 1 is the shortest path computation on the feasibility graph, with a run-time complexity for each iteration $t \in [T]$ of $\mathcal{O}\left(|\mathcal{E}^{\text{feasible}}| + |\mathcal{V}^{\text{feasible}}| \log |\mathcal{V}^{\text{feasible}}|\right)$ if Dijkstra's algorithm is used.

5.1 Epsilon-greedy

In each iteration $t \in [T]$, the Epsilon-greedy MAB algorithm selects actions either greedily, according to current parameter estimates of the loss distributions, or uniformly at random. It selects uniform exploration with a small probability ϵ_t (decreasing with t) and greedy otherwise.

In line 3 of Algorithm 1, we retrieve $\hat{\tau}^{\text{queue}}_u$, $\hat{\alpha}^{\text{charge}}_u$ and $\hat{\beta}^{\text{charge}}_u$ through MAP estimation, i.e., by finding the mode of each posterior distribution. This can be done analytically for the gamma prior / posterior in Eq. 3, such that $\hat{\tau}^{\text{queue}}_u \leftarrow \beta^{\text{queue}}_{u,t-1}/(\alpha^{\text{queue}}_{u,t-1} - 1)$, where we assume that $\alpha^{\text{queue}}_{u,t-1} > 1$. For $\hat{\alpha}^{\text{charge}}_u$ and $\hat{\beta}^{\text{charge}}_u$, the Gamcon-II prior / posterior over α^{charge}_u in Eq. 6 has no analytical formula for the mode, but it can be found numerically. With the mode $\hat{\alpha}^{\text{charge}}_u$, we can calculate $\hat{\beta}^{\text{charge}}_u \leftarrow (\xi^{\text{charge}}_{u,t-1} \cdot \hat{\alpha}^{\text{charge}}_u - 1)/\gamma^{\text{charge}}_{u,t-1}$.

Algorithm 1 CMAB charging station selection

Input: $\alpha^{\text{queue}}_{u,0}, \beta^{\text{queue}}_{u,0}, \pi^{\text{charge}}_{u,0}, \gamma^{\text{charge}}_{u,0}, \xi^{\text{charge}}_{u,0}$

1: **for** $t = 1, \ldots, T$ **do**
2: **for** $u \in \mathcal{V}^{\text{feasible}}$ **do**
3: Compute $\hat{\tau}^{\text{queue}}_u, \hat{\alpha}^{\text{charge}}_u$ and $\hat{\beta}^{\text{charge}}_u$ using specified CMAB method and current posterior parameters $\alpha^{\text{queue}}_{u,t-1}, \beta^{\text{queue}}_{u,t-1}, \pi^{\text{charge}}_{u,t-1}, \gamma^{\text{charge}}_{u,t-1}, \xi^{\text{charge}}_{u,t-1}$
4: **end for**
5: **for** $(u, u') \in \mathcal{E}^{\text{feasible}}$ **do**
6: $\hat{\tau}^{\text{feasible}}_{(u,u')} \leftarrow \tau^{\text{path}}_{(u,u')} + \hat{\tau}^{\text{queue}}_{u'} + \hat{g}_{\hat{\alpha}^{\text{charge}}_u, \hat{\beta}^{\text{charge}}_u}(u, u')$
7: **end for**
8: $\boldsymbol{p}_t \leftarrow \arg\min_{\boldsymbol{p} \in \mathcal{P}^{\text{feasible}}_{(u^{\text{src}}, u^{\text{trg}})}} \sum_{(u,u') \in \boldsymbol{p}} \hat{\tau}^{\text{feasible}}_{(u,u')}$
9: **for** each travelled edge $(u, u') \in \boldsymbol{p}_t$ **do**
10: Observe feedback $\tau^{\text{queue}}_{u'}$ and $\tau^{\text{charge}}_{(u,u')}$
11: $\alpha^{\text{queue}}_{u',t} \leftarrow \alpha^{\text{queue}}_{u',t-1} + 1$
12: $\beta^{\text{queue}}_{u',t} \leftarrow \beta^{\text{queue}}_{u',t-1} + \tau^{\text{queue}}_{u'}$
13: $\varrho^{\text{charge}}_{u'} \leftarrow \frac{\varepsilon^{\text{path}}_{(u,u')}}{\tau^{\text{charge}}_{(u,u')}}$
14: $\pi^{\text{charge}}_{u',t} \leftarrow \pi^{\text{charge}}_{u',t-1} \cdot \left(\varrho^{\max}_{u'} - \varrho^{\text{charge}}_{u'}\right)$
15: $\gamma^{\text{charge}}_{u',t} \leftarrow \gamma^{\text{charge}}_{u',t-1} + \left(\varrho^{\max}_{u'} - \varrho^{\text{charge}}_{u'}\right)$
16: $\xi^{\text{charge}}_{u',t} \leftarrow \xi^{\text{charge}}_{u',t-1} + 1.$
17: **end for**
18: **end for**

In greedy iterations, the calculated estimates are used directly in line 8 to find a path to travel. In exploration iterations, however, line 8 is changed to provide random exploration of the feasibility graph. In [12] (supplementary material), a CMAB version of Epsilon-greedy was introduced, which we adapt here. First, a vertex $u^{\text{rand}} \in \mathcal{V}^{\text{feasible}}$ is selected uniformly at random. Then, we find the paths $\boldsymbol{p}^{(1)}_t \leftarrow \arg\min_{\boldsymbol{p} \in \mathcal{P}^{\text{feasible}}_{(u^{\text{src}}, u^{\text{rand}})}} \sum_{(u,u') \in \boldsymbol{p}} \hat{\tau}^{\text{feasible}}_{(u,u')}$ and $\boldsymbol{p}^{(2)}_t \leftarrow \arg\min_{\boldsymbol{p} \in \mathcal{P}^{\text{feasible}}_{(u^{\text{rand}}, u^{\text{trg}})}} \sum_{(u,u') \in \boldsymbol{p}} \hat{\tau}^{\text{feasible}}_{(u,u')}$, which we concatenate to get $\boldsymbol{p}_t$.

5.2 Thompson Sampling

Thompson Sampling [42] is one of the oldest MAB algorithms, which has recently been adapted to CMAB problems [43], including shortest path problems with stochastic edge weights for various applications [43, 2, 4]. Like Epsilon-greedy, it performs randomized exploration, but it does so in every iteration and in a more guided way. It utilizes the knowledge encoded in the prior and posterior distributions, by sampling paths according to the probability that they are optimal (given the prior beliefs and the observations from the environment).

In Algorithm 1, only line 3 needs to be adapted to this method. Here, the expected queue time $\hat{\tau}^{\text{queue}}_u$ and charging power parameters $\hat{\alpha}^{\text{charge}}_u$ and $\hat{\beta}^{\text{charge}}_u$ are calculated using parameters sampled from the current posterior distributions. For the queue time prior distribution in Eq. 3, sampling the rate parameter $\hat{\lambda}^{\text{queue}}_u$ from the gamma distribution is straightforward, which gives an expected queue time of $\hat{\tau}^{\text{queue}}_u \leftarrow 1/\hat{\lambda}^{\text{queue}}_u$. Similar to the mode calculations in Section 5.1, sampling $\hat{\alpha}^{\text{charge}}_u$ from the Gamcon-II prior and posterior distributions is not as convenient. However, we can utilize *adaptive rejection sampling* (ARS) [18] to generate exact posterior samples, since it only requires a (log-concave, but not necessarily normalized) probability density function, like Eq. 6. In this work, we specifically use an extension called *transformed density rejection* (TDR) [24]. Once a sample of $\hat{\alpha}^{\text{charge}}_u$ is obtained, the conditional gamma prior distribution in Eq. 5 can be used to sample $\hat{\beta}^{\text{charge}}_u$.

5.3 BayesUCB

MAB algorithms based on *upper confidence bounds* (UCB) [6] use high probability overestimates of actions' expected rewards to explore the environment. By doing this, UCB methods follow the principle of *optimism in the face of uncertainty* to select promising actions. UCB methods have been shown to have good performance in many different problem settings, and have been adapted to CMAB settings [12], including shortest path problems.

We adapt a Bayesian version of UCB called *BayesUCB* [28, 27] to this setting, so that we can utilize the prior distributions for exploration. Again, like for Thompson Sampling, we only have to modify line 3 of Algorithm 1 to implement this method. BayesUCB uses (lower, in this case) quantiles of the posterior distributions over expected action losses as optimistic estimates. Given

a probability distribution χ and a probability ν, the quantile function $Q(\nu, \chi)$ is defined such that $\Pr_{x\sim\chi}\{x \leq Q(\nu,\chi)\} = \nu$. For the queue time $\hat{\tau}_u^{\text{queue}}$, a high rate parameter $\hat{\lambda}_u^{\text{queue}}$ results in a low expected travel time $\hat{\tau}_u^{\text{queue}} \leftarrow 1/\hat{\lambda}_u^{\text{queue}}$. Hence, an upper quantile of Eq. 3 should be sought, i.e., $\hat{\lambda}_u^{\text{queue}} \leftarrow Q\left(1-1/t, P\left(\lambda_u^{\text{queue}} | \alpha_{u,t-1}^{\text{queue}}, \beta_{u,t-1}^{\text{queue}}\right)\right)$, where we use the probability value $(1-1/t)$ suggested by [28]. Since the Gamcon-II prior and posterior distributions do not admit a convenient way of computing quantile values, we settle on using the mode of Eq. 6 to obtain $\hat{\alpha}_u^{\text{charge}}$. However, we utilize the mode $\hat{\alpha}_u^{\text{charge}}$ to compute an upper confidence bound for β_u^{charge} using Eq. 5, such that $\hat{\beta}_u^{\text{charge}} \leftarrow Q\left(1-1/t, P\left(\beta_u^{\text{charge}} | \hat{\alpha}_u^{\text{charge}}, \gamma_{u,t-1}^{\text{charge}}, \xi_{u,t-1}^{\text{charge}}\right)\right)$. Then, as before, we can calculate an optimistic (low) estimate of the mean charging time.

6 Experiments

To evaluate the feasibility graph construction procedure described in Section 3 and the CMAB methods outlined in Section 5, we perform realistic experiments in country-sized road networks. We define three different problem instances, characterized by their origins and destinations, across the northern European countries of Sweden, Norway and Finland. We utilize open datasets for the road network map data [35] and the charging station data [34] of each country. Our CMAB simulation framework is based on the code of [38], though significantly modified.

6.1 Energy consumption and travel time

For the vehicle energy consumption, we use a simplified vehicle longitudinal dynamics model based on [21], where we only consider the maximum speed of each road segment in the road network, i.e., disregard accelerations, decelerations and altitude changes. The vehicle parameters that we use are (arbitrarily) for a medium duty truck, similar to the one used in [2]. For an edge $e \in \mathcal{E}^{\text{road}}$, the model is defined as

$$\varepsilon_e^{\text{road}} = \frac{mgC_r d_e^{\text{road}} + 0.5 C_d A \rho d_e^{\text{road}} (v_e^{\text{road}})^2}{3600\eta}, \tag{9}$$

where m is the vehicle mass (13700 kg), g is the gravitational acceleration (9.81 m/s^2), C_r is the rolling resistance coefficient (0.0064, assumed to be the same for the entire road network), d_e^{road} is the length of the road segment (m), C_d is the air drag coefficient of the vehicle (0.7), A is the frontal surface area of the truck (8 m^2), ρ is the air density (1.2 kg/m^3, assumed to be the same everywhere), v_e^{road} is the maximum speed of the road segment (m/s), and η is the battery-to-wheel energy conversion efficiency (assumed to perfect, i.e., 1). The battery capacity ($2.5 \cdot 10^8 \text{Ws} \approx 69.4\text{kWh}$) of the vehicle is assigned to be very low, so that it is required to charge often. The travel time of the edge is assumed to be $\tau_e^{\text{road}} = d_e^{\text{road}} / v_e^{\text{road}}$.

6.2 Experimental setup

First, each of the country road network graphs is transformed into a feasibility graph according to the procedure described in Section 3. For simplicity, we remove all charging stations with lower specified power than 10 kW, since slower charging stations should be less relevant for long-distance travel. Furthermore, we assume that each charging location has a single charging station (by removing all except the one with the highest specified charging power, as well as any duplicates). The sizes of the original road networks $\mathcal{G}^{\text{road}}\left(\mathcal{V}^{\text{road}}, \mathcal{E}^{\text{road}}\right)$ and the feasibility graphs $\mathcal{G}^{\text{feasible}}\left(\mathcal{V}^{\text{feasible}}, \mathcal{E}^{\text{feasible}}\right)$, as well as further details about the results are outlined in Appendix A.

For each vertex $u \in \mathcal{V}^{\text{feasible}}$ and the queue time prior distribution defined in Eq. 3, we assign the prior parameters as $\alpha_{u,0}^{\text{queue}} = 2$ and $\beta_{u,0}^{\text{queue}} = 2400$, and for the charging power prior distribution in Eq. 5 and Eq. 6, we set the parameters so that $\pi_{u,0}^{\text{charge}} = \exp(13.5)$, $\gamma_{u,0}^{\text{charge}} = 300$ and $\xi_{u,0}^{\text{charge}} = 3$. Furthermore, we scale the samples and expected values of the gamma distribution in Eq. 4 by 300 to achieve a sufficiently high charging power variance. For the rectification of the charging power below $\varrho_u^{\min}$, we assume $\varrho_u^{\min} = \varrho_u^{\max}/2$. We choose $N = 1000$, and can then estimate the expected charging time using Monte Carlo sampling, such that $\hat{g}_{\alpha_{u'}^{\text{charge}}, \beta_{u'}^{\text{charge}}}(u,u') = \frac{1}{N}\sum_{k=1}^{N} \frac{\varepsilon_{(u,u')}^{\text{path}}}{\max(\varrho_u^{\min}, \varrho_u^{\max} - 300 z_k)}$, where $z_k \sim \text{Gamma}\left(\alpha_{u'}^{\text{charge}}, \beta_{u'}^{\text{charge}}\right)$.

Method	Sweden	Norway	Finland
GR	$1.8 \cdot 10^6(\pm 1.3 \cdot 10^6)$	$6.2 \cdot 10^5(\pm 8.6 \cdot 10^5)$	$1.8 \cdot 10^6(\pm 1.7 \cdot 10^6)$
E-GR	$4.4 \cdot 10^6(\pm 5.4 \cdot 10^5)$	$3.2 \cdot 10^6(\pm 1.1 \cdot 10^6)$	$2.3 \cdot 10^6(\pm 9.2 \cdot 10^5)$
TS	$4.0 \cdot 10^5(\pm 2.7 \cdot 10^5)$	$2.2 \cdot 10^5(\pm 2.2 \cdot 10^5)$	$4.4 \cdot 10^5(\pm 1.9 \cdot 10^5)$
B-UCB	$7.6 \cdot 10^5(\pm 2.4 \cdot 10^5)$	$2.7 \cdot 10^5(\pm 1.3 \cdot 10^5)$	$5.2 \cdot 10^5(\pm 3.3 \cdot 10^5)$

Table 1: Final average ($\pm$ standard deviation) of regret at iteration $T = 1000$ for all problem instances

For Epsilon-Greedy, we let $\epsilon_t = 1/\sqrt{t}$. Moreover, for Thompson Sampling, we experience that TDR occasionally fails to produce samples when the posterior distribution over α_u^{charge} gets too concentrated, typically after a few hundred observations. When this happens, we switch to the mode of the distribution for the specific charging station u, while continuing posterior sampling for all other charging stations. Besides Epsilon-Greedy (E-GR), Thompson Sampling (TS) and BayesUCB (B-UCB), we also include a pure greedy method (GR) in the experiments (i.e., Epsilon-Greedy with $\epsilon_t = 0$), as a baseline.

6.3 Results

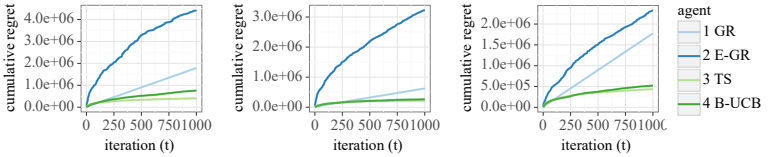

(a) Sweden (b) Norway (c) Finland

Figure 1: Plots of cumulative regret as a function of the iteration t, for each of the problem instances, and the CMAB methods Greedy (GR), Epsilon-Greedy (E-GR), Thompson Sampling (TS) and BayesUCB (B-UCB)

We run all experiments with a horizon $T = 1000$. We include the following problem instances: Sweden (Gothenburg to Stockholm), Norway (Oslo to Trondheim) and Finland (Helsinki to Vaasa). For all pairs of problem instances and CMAB methods, we run the same experiment 10 times, with different random seeds. The regret results are summarized in Table 1, as well as through regret plots in Figures 1a, 1b and 1c.

In general, the Epsilon-Greedy method incurs the highest final regret of all the methods, for all problem instances. This can be explained by the uniformly random selection of charging stations, which means that the method takes very long detours, sometimes to the other side of the country. Following close behind is the Greedy method, which quickly converges to sub-optimal paths. This becomes even more apparent in the regret plots, which show the regret (averaged over 10 runs) as a function of the iteration t. For each of the problem instances, the Greedy method exhibits a linear increase in regret, while the other methods (even Epsilon-greedy) continuously find better paths. For all problem instances, Thompson Sampling performs slightly better than BayesUCB, which may be due to the more optimistic exploration of BayesUCB resulting in a wider spread of explored paths (and consequently, occasionally more regret incurred).

7 Conclusion

In this work, we developed a combinatorial semi-bandit framework for navigation and charging station selection in road networks where the queue time and charging power of each charging station are stochastic with unknown distributions, the parameters of which are generated from known prior distributions. We utilized conjugate prior distributions for the exponential and gamma models to estimate the loss distributions and induce exploration. We then demonstrated the performance of our framework on several country-sized road and charging networks.

Acknowledgments and Disclosure of Funding

This work is funded by the Strategic Vehicle Research and Innovation Programme (FFI) of Sweden, through the project EENE (reference number: 2018-01937). Map data copyrighted OpenStreetMap contributors and available from `https://www.openstreetmap.org`.

References

[1] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In Shie Mannor, Nathan Srebro, and Robert C. Williamson, editors, *Proceedings of the 25th Annual Conference on Learning Theory*, volume 23 of *Proceedings of Machine Learning Research*, pages 39.1–39.26, Edinburgh, Scotland, 25–27 Jun 2012. PMLR.

[2] Niklas Åkerblom, Yuxin Chen, and Morteza Haghir Chehreghani. An online learning framework for energy-efficient navigation of electric vehicles. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 2051–2057, Jul 2020. doi: 10.24963/ijcai.2020/284.

[3] Niklas Åkerblom, Yuxin Chen, and Morteza Haghir Chehreghani. Online learning of energy consumption for navigation of electric vehicles. *Artificial Intelligence*, 317, 2023. doi: 10.1016/j.artint.2023.103879.

[4] Niklas Åkerblom, Fazeleh Sadat Hoseini, and Morteza Haghir Chehreghani. Online learning of network bottlenecks via minimax paths. *Machine Learning*, 112(1):131–150, 2023. doi: 10.1007/s10994-022-06270-0.

[5] Andreas Artmeier, Julian Haselmayr, Martin Leucker, and Martin Sachenbacher. The shortest path problem revisited: Optimal routing for electric vehicles. In Rüdiger Dillmann, Jürgen Beyerer, Uwe D. Hanebeck, and Tanja Schultz, editors, *KI 2010: Advances in Artificial Intelligence*, pages 309–316, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. ISBN 978-3-642-16111-7. doi: 10.1007/978-3-642-16111-7_35.

[6] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *J. Mach. Learn. Res.*, 3:397–422, Mar 2002. ISSN 1532-4435.

[7] Moritz Baum, Julian Dibbelt, Andreas Gemsa, Dorothea Wagner, and Tobias Zündorf. Shortest feasible paths with charging stops for battery electric vehicles. In *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems*, SIGSPATIAL '15, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450339674. doi: 10.1145/2820783.2820826. URL `https://doi.org/10.1145/2820783.2820826`.

[8] Richard Bellman. On a routing problem. *Quarterly of applied mathematics*, 16:87–90, 1958. doi: 10.1090/qam/102435.

[9] Sébastien Bubeck and Che-Yu Liu. Prior-free and prior-dependent regret bounds for thompson sampling. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'13, pages 638–646, Red Hook, NY, USA, 2013. Curran Associates Inc.

[10] Nicolò Cesa-Bianchi and Gábor Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012. ISSN 0022-0000. doi: 10.1016/j.jcss.2012.01.001. JCSS Special Issue: Cloud Computing 2011.

[11] Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In *Proceedings of the 24th International Conference on Neural Information Processing Systems*, NIPS'11, pages 2249–2257, Red Hook, NY, USA, 2011. Curran Associates Inc. ISBN 9781618395993.

[12] Wei Chen, Yajun Wang, and Yang Yuan. Combinatorial multi-armed bandit: General framework and applications. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 151–159, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.

[13] Eivind Damsleth. Conjugate classes for gamma distributions. *Scandinavian Journal of Statistics*, 2(2):80–84, 1975. ISSN 03036898, 14679469. URL http://www.jstor.org/stable/4615580.

[14] Rina Dechter and Judea Pearl. Generalized best-first search strategies and the optimality of a*. *J. ACM*, 32(3):505–536, Jul 1985. ISSN 0004-5411. doi: 10.1145/3828.3830.

[15] Edsger W Dijkstra. A note on two problems in connexion with graphs. *Numerische mathematik*, 1:269–271, 1959. doi: 10.1007/BF01386390.

[16] Lester R Ford Jr. Network flow theory. Technical Report P-932, Rand Corporation, Santa Monica, CA, 1956.

[17] Yi Gai, Bhaskar Krishnamachari, and Rahul Jain. Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations. *IEEE/ACM Transactions on Networking*, 20(5):1466–1478, 2012. doi: 10.1109/TNET.2011.2181864.

[18] W. R. Gilks and P. Wild. Adaptive rejection sampling for gibbs sampling. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 41(2):337–348, 1992. doi: https://doi.org/10.2307/2347565.

[19] Thore Graepel, Joaquin Quiñonero Candela, Thomas Borchert, and Ralf Herbrich. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft's bing search engine. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, pages 13–20, Madison, WI, USA, 2010. Omnipress. ISBN 9781605589077.

[20] Marianne Guillet, Gerhard Hiermann, Alexander Kröller, and Maximilian Schiffer. Electric vehicle charging station search in stochastic environments. *Transportation Science*, 56(2):483–500, 2022. doi: 10.1287/trsc.2021.1102.

[21] Lino Guzzella, Antonio Sciarretta, et al. *Vehicle propulsion systems*. Springer, 2007. doi: 10.1007/978-3-540-74692-8.

[22] Aric A Hagberg, Daniel A Schult, and Pieter J Swart. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conference (SciPy2008)*, pages 11–15, 2008.

[23] Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968. doi: 10.1109/TSSC.1968.300136.

[24] Wolfgang Hörmann. A rejection technique for sampling from t-concave distributions. *ACM Trans. Math. Softw.*, 21(2):182–193, jun 1995. ISSN 0098-3500. doi: 10.1145/203082.203089. URL https://doi.org/10.1145/203082.203089.

[25] Huiwen Jia, Cong Shi, and Siqian Shen. Multi-armed bandit with sub-exponential rewards. *Operations Research Letters*, 49(5):728–733, 2021. ISSN 0167-6377.

[26] H.C Joksch. The shortest route problem with constraints. *Journal of Mathematical Analysis and Applications*, 14(2):191–197, 1966. ISSN 0022-247X. doi: https://doi.org/10.1016/0022-247X(66)90020-5. URL https://www.sciencedirect.com/science/article/pii/0022247X66900205.

[27] Emilie Kaufmann. On bayesian index policies for sequential resource allocation. *The Annals of Statistics*, 46(2):842–865, 2018. ISSN 00905364, 21688966.

[28] Emilie Kaufmann, Olivier Cappe, and Aurelien Garivier. On bayesian upper confidence bounds for bandit problems. In Neil D. Lawrence and Mark Girolami, editors, *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, volume 22 of *Proceedings of Machine Learning Research*, pages 592–600, La Palma, Canary Islands, 21–23 Apr 2012. PMLR.

[29] Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In Nader H. Bshouty, Gilles Stoltz, Nicolas Vayatis, and Thomas Zeugmann, editors, *Algorithmic Learning Theory*, pages 199–213, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. ISBN 978-3-642-34106-9. doi: 10.1007/978-3-642-34106-9_18.

[30] David G. Kendall. Stochastic processes occurring in the theory of queues and their analysis by the method of the imbedded markov chain. *The Annals of Mathematical Statistics*, 24(3): 338–354, 1953. ISSN 00034851. URL `http://www.jstor.org/stable/2236285`.

[31] Ki-Beom Lee, Mohamed A. Ahmed, Dong-Ki Kang, and Young-Chon Kim. Deep reinforcement learning based optimal route and charging station selection. *Energies*, 13(23), 2020. ISSN 1996-1073. doi: 10.3390/en13236255. URL `https://www.mdpi.com/1996-1073/13/23/6255`.

[32] Keqin Liu and Qing Zhao. Adaptive shortest-path routing under unknown and stochastically varying link states. In *2012 10th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, pages 232–237, 2012.

[33] Robert B. Miller. Bayesian analysis of the two-parameter gamma distribution. *Technometrics*, 22(1):65–69, 1980. doi: 10.1080/00401706.1980.10486102.

[34] Open Charge Map. Open Charge Map — the global public registry of electric vehicle charging locations. `https://openchargemap.org`, 2022.

[35] OpenStreetMap contributors. Planet dump retrieved from https://planet.osm.org . `https://www.openstreetmap.org`, 2022.

[36] Tao Qian, Chengcheng Shao, Xiuli Wang, and Mohammad Shahidehpour. Deep reinforcement learning for ev charging navigation by coordinating smart grid and intelligent transportation system. *IEEE Transactions on Smart Grid*, 11(2):1714–1723, 2020. doi: 10.1109/TSG.2019.2942593.

[37] Daniel Russo and Benjamin Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243, 2014.

[38] Daniel J. Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1):1–96, 2018. ISSN 1935-8237. doi: 10.1561/2200000070.

[39] Martin Sachenbacher, Martin Leucker, Andreas Artmeier, and Julian Haselmayr. Efficient energy-optimal routing for electric vehicles. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, AAAI'11, pages 1402–1407. AAAI Press, 2011.

[40] Alfonso Shimbel. Structure in communication nets. In *Proceedings of the symposium on information networks*, pages 199–203. Polytechnic Institute of Brooklyn, 1954.

[41] Timothy M. Sweda, Irina S. Dolinskaya, and Diego Klabjan. Adaptive routing and recharging policies for electric vehicles. *Transportation Science*, 51(4):1326–1348, 2017. doi: 10.1287/trsc.2016.0724.

[42] W.R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3–4):285–294, 1933. doi: 10.2307/2332286.

[43] Siwei Wang and Wei Chen. Thompson sampling for combinatorial semi-bandits. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5114–5122. PMLR, 10–15 Jul 2018.

[44] Zhenhua Zou, Alexandre Proutiere, and Mikael Johansson. Online shortest path routing: The value of information. In *2014 American Control Conference*, pages 2142–2147, 2014. doi: 10.1109/ACC.2014.6859133.

A Appendix

In this Appendix, we report further details on the experiments than in the main paper. We run all experiments with a horizon $T = 1000$. We include the following problem instances: Sweden (Gothenburg to Stockholm), Norway (Oslo to Trondheim), and Finland (Helsinki to Vaasa). Table 2 shows the number of vertices and edges in the initial road graph and constructed feasibility graph of each problem instance. Figures 2a, 2c and 2e visualize all edges of the road network graphs of Sweden, Norway and Finland, respectively, as well as examples of the explored paths and charging stations visited when Thompson Sampling is applied to the problem. Figures 2b, 2d and 2f show the corresponding feasibility graphs for each of the networks. In the feasibility graphs, we can see that some parts of the road networks are unreachable from the rest of the road network, given the specified battery capacity of 69.4 kWh. We want to emphasize that the methods developed in this work may be applied to problem instances with changing source and target vertices (for either single or multiple vehicles), which would require exploration of much greater parts of the road networks than what is shown in Figure 2.

For all pairs of problem instances and CMAB methods, we run the same experiment 10 times, with different random seeds. To further illustrate the regret results summarized in the main paper, we also include regret plots with standard error regions in Figures 3a, 3b and 3c. In Table 3, we report the average and standard deviation of the per-iteration run-time (in seconds) of each method and problem instance. The CMAB methods are implemented with Python using the SciPy library, for implementations of statistical functions, including the Transformed Density Rejection method [24], and the NetworkX library [22], for the implementation of the A* algorithm. All run-time measurements were performed on a single core of a laptop with an Intel(R) Core(TM) i7-10850H CPU (2.70 GHz) and 32.00 GB RAM.

The notation used throughout the paper is summarized in Table 4.

Instance	$\lvert\mathcal{V}^{\text{road}}\rvert$	$\lvert\mathcal{E}^{\text{road}}\rvert$	$\lvert\mathcal{V}^{\text{feasible}}\rvert$	$\lvert\mathcal{E}^{\text{feasible}}\rvert$
Sweden	$6.8 \cdot 10^5$	$1.5 \cdot 10^6$	$1.7 \cdot 10^3$	$1.6 \cdot 10^5$
Norway	$3.6 \cdot 10^5$	$7.6 \cdot 10^5$	$1.1 \cdot 10^3$	$8.5 \cdot 10^4$
Finland	$4.3 \cdot 10^5$	$9.5 \cdot 10^5$	$5.5 \cdot 10^2$	$7.0 \cdot 10^4$

Table 2: Sizes of the road and feasibility graphs for all problem instances

Method	Sweden	Norway	Finland
GR	$2.04(\pm 0.18)$	$2.61(\pm 0.27)$	$3.09(\pm 0.37)$
E-GR	$2.80(\pm 0.51)$	$2.83(\pm 0.79)$	$3.38(\pm 0.77)$
TS	$2.79(\pm 0.78)$	$2.92(\pm 0.62)$	$2.93(\pm 0.36)$
B-UCB	$3.47(\pm 0.32)$	$2.97(\pm 0.37)$	$3.46(\pm 0.42)$

Table 3: Average (over 1000 iterations) per-iteration run-time (s) ($\pm$ standard deviation) of each CMAB method, for different problem instances

Notation	Description
$\mathcal{G}^{\text{road}}$	Road graph
$\mathcal{V}^{\text{road}}$	Road graph vertices
$\mathcal{E}^{\text{road}}$	Road graph edges
$\boldsymbol{\tau}^{\text{road}}$	Vector of road graph edge travel times
τ_e^{road}	Road graph travel time of edge e
$\tau_{\boldsymbol{p}}^{\text{road}}$	Road graph travel time of path $\boldsymbol{p}$
$\varepsilon_e^{\text{road}}$	Road graph energy consumption of edge e
$\varepsilon_{\boldsymbol{p}}^{\text{road}}$	Road graph energy consumption of path $\boldsymbol{p}$

Table 4: Summary of the notation used throughout the paper

Continued on the next page

Notation	Description
$\mathcal{V}^{\text{charge}}$	Charging station vertices
$\varrho_u^{\text{charge}}$	Charging station power (actual) of vertex u
$\varrho_u^{\max}$	Charging station maximum power of vertex u
$\varrho_u^{\min}$	Charging station minimum power of vertex u
u	A vertex
u'	A vertex (alternative)
u^{src}	Source vertex
u^{trg}	Target vertex
e	An edge
$\boldsymbol{p}$	A path
$\boldsymbol{p}^*$	Shortest path (implicitly between specified source and target vertices)
$\boldsymbol{p}^*_{(u,u')}$	Shortest path between vertices u and u'
$\mathcal{P}^{\text{road}}_{(u^{\text{src}},u^{\text{trg}})}$	Set of paths in road graph between source and target vertices
$\mathcal{G}^{\text{feasible}}$	Feasibility graph
$\mathcal{V}^{\text{feasible}}$	Feasibility graph vertices
$\mathcal{E}^{\text{feasible}}$	Feasibility graph edges
$\mathcal{E}^{\text{path}}$	Edges corresponding to shortest paths between charging stations in road graph
$\tau^{\text{path}}_{(u,u')}$	Path travel time between vertices u and u'
$\tau^{\text{charge}}_{(u,u')}$	Charging time required for travel between vertices u and u'
τ^{queue}_u	Queue time required for charging station at vertex u
τ^{feasible}_e	Feasibility graph travel time of edge e
$\varepsilon^{\text{path}}_{(u,u')}$	Path energy consumption between vertices u and u'
$\varepsilon^{\max}$	Maximum battery capacity of vehicle
$\varepsilon^{\min}$	Minimum battery capacity of vehicle
λ_u^{queue}	Queue time distribution rate parameter of vertex u
$\alpha^{\text{queue}}_{u,0}$	Shape prior parameter for queue time rate parameter, of vertex u
$\beta^{\text{queue}}_{u,0}$	Rate prior parameter for queue time rate parameter, of vertex u
$\alpha^{\text{queue}}_{u,t}$	Shape posterior parameter for queue time rate parameter, at time t and vertex u
$\beta^{\text{queue}}_{u,t}$	Rate posterior parameter for queue time rate parameter, at time t and vertex u
α_u^{charge}	Charging power distribution shape parameter of vertex u
β_u^{charge}	Charging power distribution rate parameter of vertex u
$\pi^{\text{charge}}_{u,0}$	First prior parameter for charging power parameters, of vertex u
$\gamma^{\text{charge}}_{u,0}$	Second prior parameter for charging power parameters, of vertex u
$\xi^{\text{charge}}_{u,0}$	Third prior parameter for charging power parameters, of vertex u
$\pi^{\text{charge}}_{u,t}$	First posterior parameter for charging power parameters, at time t and vertex u
$\gamma^{\text{charge}}_{u,t}$	Second posterior parameter for charging power parameters, at time t and vertex u
$\xi^{\text{charge}}_{u,t}$	Third posterior parameter for charging power parameters, at time t and vertex u
t	Time step
T	Time horizon
$\boldsymbol{p}_t$	Path (action) selected by CMAB algorithm at time t
$L_t(\boldsymbol{p})$	Loss of path $\boldsymbol{p}$ at time t
$\boldsymbol{\theta}$	Vector of all parameters of feasibility graph
$\boldsymbol{\theta}^*$	True underlying (unknown) vector of all parameters of feasibility graph
$\boldsymbol{\theta}_u$	Vector of parameters $\left(\lambda_u^{\text{queue}}, \alpha_u^{\text{charge}}, \beta_u^{\text{charge}}\right)$ of vertex u
$f_{\boldsymbol{\theta}}(\boldsymbol{p})$	Expected loss function, w.r.t. parameter vector $\boldsymbol{\theta}$, applied to path $\boldsymbol{p}$
$g_{\alpha,\beta}(u,u')$	Expected charging time of edge (u,u'), given parameters α and β
$\hat{g}_{\alpha,\beta}(u,u')$	Estimated charging time of edge (u,u'), given parameters α and β
$\text{Regret}(T)$	Regret of CMAB algorithm at horizon T
ϵ_t	Exploration probability of the Epsilon-greedy CMAB algorithm
$\hat{\tau}_u^{\text{queue}}$	Estimate of mean queue time, for vertex u

Table 4: Summary of the notation used throughout the paper

Continued on the next page

Notation	Description
$\hat{\lambda}_u^{\text{queue}}$	Estimate of queue time rate parameter, for vertex u
$\hat{\alpha}_u^{\text{charge}}$	Estimate of charging power shape parameter, for vertex u
$\hat{\beta}_u^{\text{charge}}$	Estimate of charging power rate parameter, for vertex u
$Q(\nu, \chi)$	Quantile function for distribution χ and probability value ν
m	Vehicle mass
g	Gravitational acceleration
C_r	Rolling resistance coefficient
d_e^{road}	Length of the road segment, for edge e
C_d	Air drag coefficient
A	Frontal surface area of vehicle
v_e^{road}	Maximum speed of the road segment, for edge e
ρ	Air density
η	Battery-to-wheel energy conversion efficiency

Table 4: Summary of the notation used throughout the paper.

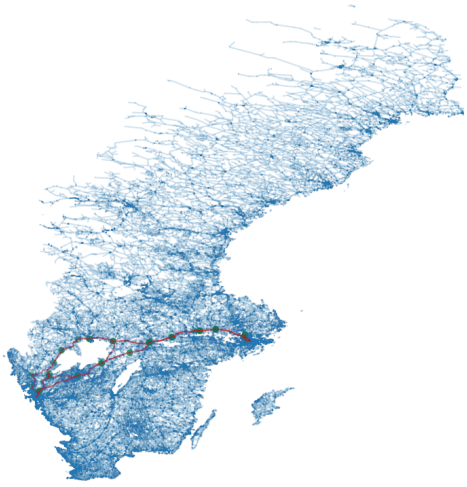

(a) Road graph: Sweden

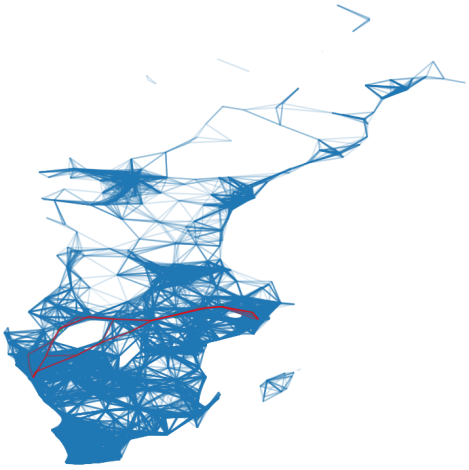

(b) Feasibility graph: Sweden

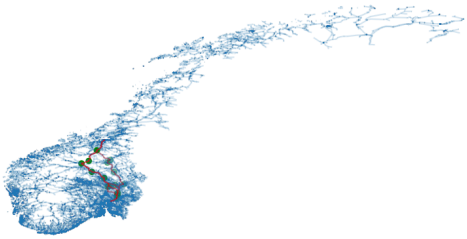

(c) Road graph: Norway

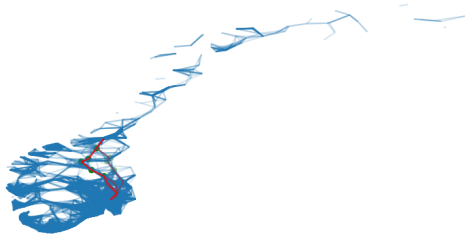

(d) Feasibility graph: Norway

(e) Road graph: Finland

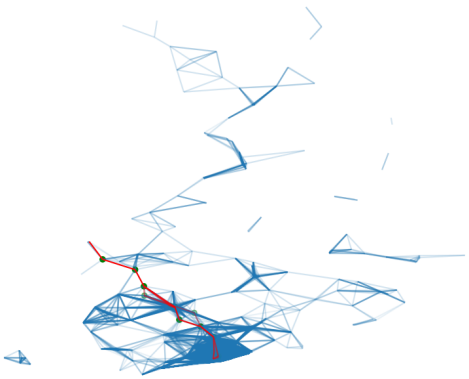

(f) Feasibility graph: Finland

Figure 2: Road and feasibility graphs for each of the problem instances in blue, with Thompson Sampling exploration of paths in red and charging stations in green, where opaqueness indicates degree of exploration for both

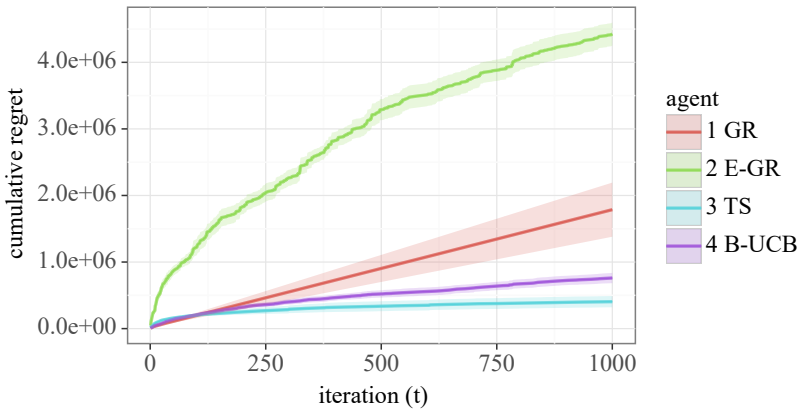

(a) Cumulative regret: Sweden

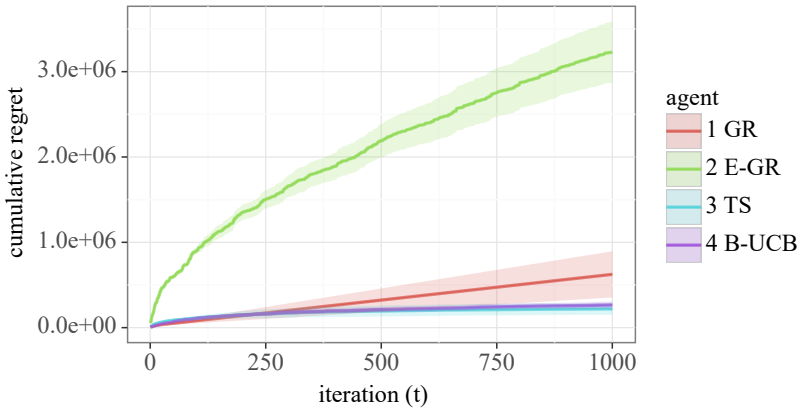

(b) Cumulative regret: Norway

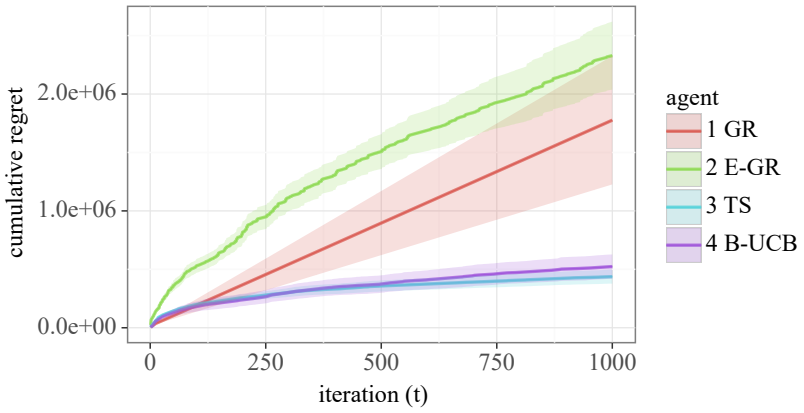

(c) Cumulative regret: Finland

Figure 3: Plots of cumulative regret averaged over 10 runs (with standard error regions), as a function of the iteration t, for each of the problem instances, and the CMAB methods Greedy (GR), Epsilon-Greedy (E-GR), Thompson Sampling (TS) and BayesUCB (B-UCB)