Optimization of sustainable electric vehicle charging

Martijn de Jong, Markus Niebisch, Evie Nielen,
Cyriel van Velzen, Emma Verweij

Abstract

This report proposes an approach to obtain an optimal charging schedule for electric vehicles (EVs) with a focus on sustainability. The problem was posed by Groendus, an energy partner that helps businesses become more sustainable by providing charging stations, local production of green energy, and an energy management system. Smart EV charging presents a two-fold problem that can be separated into the deterministic optimization of EV charging schemes and the prediction of driver behavior. We propose a mixed-integer linear programming model to solve the optimization problem, minimizing the use of grid-bought energy while still meeting the expectations of EV drivers. Parameter estimations based on historical data and day-ahead forecasts are used to demonstrate the performance of the implemented model in various (artificial) scenarios. With our implementation, charging schemes follow the trend of solar energy and pricing. Lastly, we provide an outlook for an online model that incorporates real-time predictions. The approach and results presented in this report serve as a starting point for further studies with real data to ultimately obtain more sustainable EV charging schedules.

1 Problem formulation

This project is conducted in collaboration with Groendus, a Dutch company dedicated to advancing the green energy transition by implementing eco-friendly solutions for their corporate clientele. Among their services, Groendus provides businesses with EV charging systems, aiming to optimize energy management and sustainability. Additionally, they install solar panels to harness locally sourced renewable energy, which can power the charging systems and other operations within the business premises, such as factories.

Roughly two types of parking lots may be present at businesses where Groendus provides these services. The first type is company parking spaces that are regularly used by recurring users (mainly employees), and the second type is locations such as big warehouse parking spaces, where there is less regularity and cars remain parked for a shorter time. For all cars at the charging stations, charging schedules should be determined. This decision can depend on the amount of solar energy available, altering the energy price during the day, and the number of cars at the charging stations.

In the current situation, cars that plug into one of their charging stations are charged at the maximum power as long as this price falls below some threshold. This is implemented in a decision tree. This leads to situations where cars have been fully charged when solar energy production peaks, leading to surplus energy being sold back to the grid. This scenario is particularly prevalent in the first type of parking lot, which is predominantly used by employees. This is also the situation we have focused on, in part because we see less potential for (significant) optimization in the case of irregular and short parking sessions. Groendus, as well as their partners, would prefer a more intelligent charging system that charges at desirable times – specifically, when self-produced green energy is abundant – saving costs but, more importantly, contributing to the green energy transition.

The desired 'smart' solution would incorporate historical user data and/or predictions on remaining battery capacity and duration of stay of cars to determine individual charging schedules, with two aims: (i) to meet a minimum charge level when the car leaves (in the drivers' interests) and (ii) to minimize the feeding of green energy to the grid and buying from the grid. This second aim is in the interest of the green energy transition and, secondarily, of the owner of the charging hub, who minimizes costs.

This problem depends on a lot of parameters that are not all known. To overcome this uncertainty, we separated the problem into two parts: the first part is to solve the optimization problem under the assumption that all information is perfectly known using a Mixed-integer Linear Programming model (MILP), and the second part is to estimate the parameters that are not known. The estimated parameters are used as input for the MILP model. In doing so, we have split the problem posed by Groendus into two parts, one of which was to create a model to determine an optimal charging schedule for EVs.

The outlook of this paper is as follows. In Section 2, we introduce the MILP, where we assume all the parameters to be known. In Section 3, we determine which parameters are really known and how to estimate the remaining unknown parameters from data. Section 4 describes the implementation of the MILP and displays the model results for several scenarios. To improve the charging schedule, it could be sensible to investigate real-time updates during the day. An outline of this idea is given in Section 5. Lastly, we discuss the results in Section 6 and conclude the paper in Section 7.

2 Introduction of deterministic model

In this section, we describe the mixed-integer linear programming model that yields an optimal EV charging scheme, if the given input contains perfect information of the future.

In Table 1, we summarize the parameters in our model. We note that in our model, the total number of parking sessions during a 24-hour period N is known. Furthermore, in the table, a superscript $^+$ indicates energy purchased from the grid, whereas $^-$ indicates energy sold to the grid.

The primary objective of the model is sustainability, balanced with customer satisfaction and fair energy distribution among cars. Secondary objectives include cost minimization and revenue maximization from selling surplus energy to the grid. These objectives are achieved by minimizing a combination of total grid costs, customer demand fulfillment, and a fairness factor ensuring equitable charging across all users.

In the following, the sets over which we sum are left implicit, and we refer the reader to Table 1 for the definitions of the indices i, j and t.

Parameter	Description	Unit
n	in $\mathbb{N}$, number of charging stations	-
N	in $\mathbb{N}$, total number of parking sessions during a 24-hour period	-
Δt	duration of time slot, equal to 15 minutes = 0.25 hours	s
T	in $\mathbb{N}$, number of time slots of duration Δt	-
i	in $\{1, \ldots, n\}$, index for charging station	-
j	in $\{1, \ldots, N\}$, index for sessions	-
$i(j)$	index $i = i(j)$ of charging station corresponding to session j	
t	in $\mathcal{T} = \{1, \ldots, T\}$, index for time slots	-
E_t^+	total energy taken from the grid during time slot t	kWh
E_t^-	total energy sold to the grid during time slot t	kWh
E_t^i	total energy used by charging station i during time slot t	kWh
E_t^{prod}	total energy produced (by solar panels, e.g.) during time slot t	kWh
$E_t^{\text{oth.}}$	total energy used by business apart from car charging	kWh
p_t^+	costs of grid energy at time t	e/kWh
p_t^-	payment received for selling energy to the grid at time t	e/kWh
α	penalty encouraging usage of own production	-
$r_{\min}^j, r_{\max}^j$	min. and max. energy input for car j (per time slot)	kWh
s_j	starting time of session j	-
f_j	end time of session j	-
$\mathcal{R}_j = [s_j, f_j]$	the time car j is at the charging station, where $s_j, f_j \in \mathcal{T}$	-
$\mathcal{J}_t$	$\mathcal{J}_t = \{j \mid t \in \mathcal{R}_j\}$, at time t, the cars which are present	-
m_j	total remaining charging capacity of car j	kWh
n_j	the minimum demanded amount of energy charged into car j	kWh
$c = (c_t^i)$	$c_t^i \in \{0, 1\}$ specifies whether there is a car at (t, i)	-
y_t^i	$y_t^i \in \{0, 1\}$ specifies if charging station i is active at time t	-
g	the maximum grid connection	
	i.e., the total amount of energy available from the grid	kWh

Table 1: Summary of notation.

The first term of our objective function determined grid costs. At any time slot $t \in \mathcal{T}$, the grid energy price is given by p_t^+, and the amount of energy bought from the grid in that time slot is E_t^+, such that the total costs during a day are given by

$$\sum_t E_t^+ p_t^+.$$

For the second term of the objective function, we compute customer demand fulfillment as follows. The total charging capacity of all cars on a day is given by the sum $\sum_j m_j$, while the total energy received by the cars on a day (that is, used by all charging stations) equals $\sum_{i,t} E_t^i$. The ratio of these then gives the average state of charge (SoC), $\frac{\sum_{i,t} E_t^i}{\sum_j m_j}$.

Picking a target SoC ($s_{\text{target}} \in [0, 1]$), the deviation of the average SoC is given by $\left| s_{\text{target}} - \frac{\sum_{i,t} E_t^i}{\sum_j m_j} \right|$, giving the customer demand fulfillment term in the objective function. In our objective function, we multiply this term with a constant $M \gg 1$.

Finally, the third term of the objective function is constructed as follows. To ensure equitable charging across all users, we aim to minimize the total deviation across all users of the energy received relative to the mean energy input. The mean energy input, calculated as $\frac{1}{N}\sum_{i,t} E^i_t$, serves as a reference point. Conversely, the energy received by each session $j \in \{1, \dots, N\}$ is represented by $\sum_t E^{i(j)}t$. The total deviation from this mean energy input is then quantified as

$$\sum_j \left| \sum_t \left(E^{i(j)}t - \frac{1}{N}\sum_i E^i_t \right) \right|.$$

Combining the three terms, the objective function in our MILP model is given by[1]

$$\underset{E^t_+, E^i_t}{\text{minimize}} \underbrace{\sum_t E^+_t p^+_t}_{\text{total grid costs}} + M \underbrace{\left| s_{\text{target}} - \frac{\sum_{i,t} E^i_t}{\sum_j m_j} \right|}_{\text{customer fulfilment}} + \underbrace{\sum_j \left| \sum_t \left(E^{i(j)}_t - \frac{1}{N}\sum_i E^i_t \right) \right|}_{\text{equitable distribution of energy}} \tag{1}$$

subject to constraints

$$E^+_t + E^{\text{prod}}_t = E^-_t + \sum_i E^i_t + E^{\text{oth.}}_t, \qquad t \in \mathcal{T} \tag{2}$$

$$y^i_t r_{\text{min}} \le E^{i(j)}_t \le r^j_{\text{max}} y^i_t, \qquad t \in \mathcal{T}, i \in \{1, \dots, n\} \tag{3}$$

$$y^i_t \le c^i_t, \qquad i \in \{1, \dots, n\} \tag{4}$$

$$n_j \le \sum_{t \in \mathcal{R}_j} E^{i(j)}_t \le m_j, \qquad j \in \{1, \dots, N\} \tag{5}$$

$$E^+_t \le g, \qquad t \in \mathcal{T} \tag{6}$$

$$E^-_t \le E^{\text{prod}}_t, \qquad t \in \mathcal{T} \tag{7}$$

$$E^+_t, E^-_t, E^{\text{prod}}_t \ge 0, \qquad t \in \mathcal{T} \tag{8}$$

The first constraint, (2), enforces that the total influx of energy during time slot $t \in \mathcal{T}$, the purchased energy and the produced energy, is equal to the total outflux of energy, the energy used by the stations, the other energy used by the business and the energy sold to the grid. Constraint (3) specifies the charging range of car batteries and chargers (recall that $i = i(j)$ is the index of the charger at which car j is parked): chargers must give a minimum of 2 kW, and can give a maximum grid output of $r_{\text{max}} = 11$ kW. Moreover, typical car batteries are divided into three categories: small batteries, with a maximum power input of $r_{\text{max}} = 3.7$ kW, medium-size batteries with $r_{\text{max}} = 7.2$ kW and large batteries which can charge at the maximum charge rate of the charger itself. Hence (in terms of energy per time slot), this constraint requires the charging energy to be either zero or in this range. Now, the energy input is non-zero exactly when $y^i_t = 1$, but this can only be true when there is a car, specified by (4).

[1] The absolute signs in the function are implemented in the standard way, rewriting an absolute value as a sum of two positive numbers determined using an extra constraint.

During a session j, the total charged energy should meet a minimum demand (depending on the duration of the session) and is also capped by the maximum capacity contained in (5). Constraint (6) includes the cap given by the maximum grid connection, and (7) states that only produced energy is sold to the grid. Finally, all energy is positive, as is required in (8).

3 Parameter estimation from data

As described in the previous section, the deterministic model relies on the large set of parameters summarized in Table 1, which are assumed to be known. However, perfect information on parameters is rarely available in a real-world scenario. In the upcoming section, we discuss which parameters are known through day-ahead forecasts or real-time input from charging stations and which parameters need to be estimated. Surface level analysis of mock data provided by Groendus suggests that historical data is suitable for the estimation of feasible parameters, as we demonstrate in this section.

3.1 Known parameters

We consider two types of known parameters: those that are assumed to be known as they are completely determined by external factors, and parameters Groendus can measure or influence themselves. One of the inputs determined by external factors is the price of energy from the grid per hour, i.e., $p_t^{\pm}$. Every day at noon, the day-ahead grid prices for buying from and selling to the grid are fixed and published, and those prices will not change throughout the day. The maximal grid connection g is also known. However, due to the increasing problem of net congestion, predictions for the available energy from the grid still contain a large margin of error ($\sim$ 22%). Moreover, predictions on the available energy from own production E_t^{prod} (e.g., from solar panels) have an even larger margin of error ($\sim$ 25%) due to uncertainty in weather forecasting. Hence, we use historical data to estimate the available energy from a company's local production throughout a sunny day. To tune this prediction based on the daily weather forecast, we introduce a 'cloudy' parameter, which determines the fraction of the predicted solar power that is assumed to be delivered on a cloudy day. The second category of parameters consists of those parameters that Groendus can measure directly. Naturally, the number of charging stations n at a parking lot is known, and since Groendus works mostly with small businesses, the number of n is assumed to be around 40. Through the charging stations, the maximal charge rate $r_{\max}$ of an EV can be measured. Unfortunately, no further information can be gathered from charging stations, and hence, all other car specifications, such as the battery size and the SoC, must be estimated.

3.2 Estimation of EV specifications

For the estimation of unknown EV specifications, we use the mock data provided by Groendus. First, we note that the three EV charge rates (3.7 kW, 7.2 kW and 11 kW respectively) can be easily identified in the data, see Figure 1.

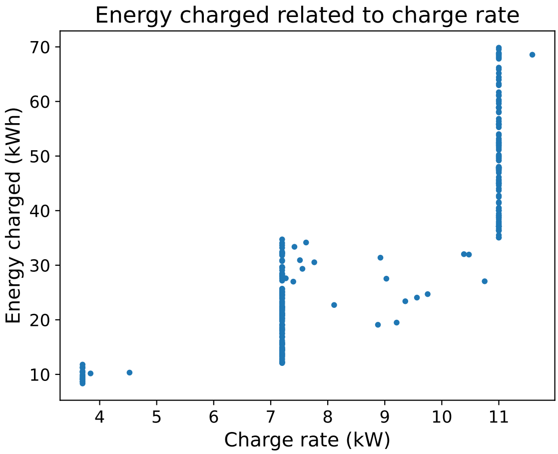

Figure 1: Scatter plot of charge rates of EVs related to the energy charged during a charging session. Three typical charge rates can be identified: 3.7 kW, 7.2 kW and 11 kW.

Additionally, Figure 1 shows a correlation between the maximal charging rate and the energy charged during a session, and gives a distinct range for the possible total demand of EVs dependent on their maximal charging rates. Following Ghotge et al. (2020), we assume feasible values for the state of charge (SoC) on arrival to be between 0 – 90%, with a mean of 50%. A truncated normal distribution was used to estimate the SoCs of EVs on arrival. Hence, the battery capacity of an EV can be roughly estimated by the upper bound of the total energy charged per charge rate, as shown in Figure 1. This gives rise to three categories of battery sizes, to wit: smaller batteries of roughly 20 kWh, mid-sized batteries of roughly 40 kWh and larger batteries of roughly 70 kWh. Due to a lack of data for EVs with a low charge rate, the energy charged by small EVs was assumed to be 50% of their total battery capacity, such that the total battery capacity is indeed 20 kWh. These battery sizes agree with typical sizes used in similar studies (Kuran et al. (2015), Limmer (2019)). Once again, truncated normal distributions were used to generate feasible battery sizes for EVs within one of these three categories. Note that the above estimations are realistic, as the EVs in the current datasets of Groendus are typically fully charged during a session, such that the maximal demand of visiting EVs can be determined. In particular, Figure 2 showcases the current approach of Groendus, in which EVs are charged at their maximal charge rate at the start of a session. Hence, the average charge during the active charging phase of a session is close to one of the three previously mentioned maximal charge rates, but the charge rate over the full session is much lower.

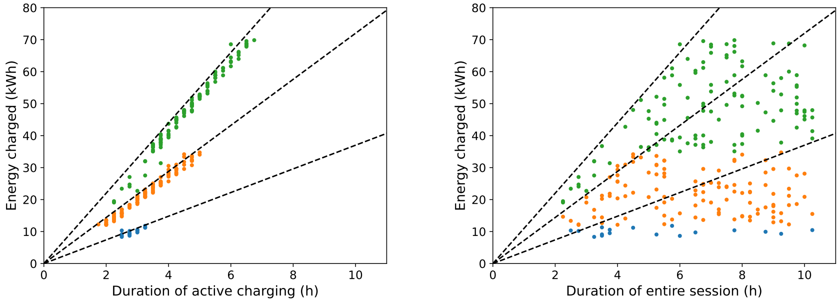

Figure 2: Scatter plots of the total energy charged in the active charging phase of a session (left) and the energy charged compared to the full length of a session (right). Maximal charging rates are shown (dashed lines), and cars with charging rates of 3.7 kW, 7.2 kW, and 11 kW are colored blue, orange and green, respectively.

The data of this particular business showed that the battery size of 9% of the visiting EVs was small, 45% mid-sized, and 46% large. Naturally, this distribution can vary for different businesses, and needs to be adjusted on a case-by-case basis. Table 2 summarizes the parameters used for synthetic data generation as input of the MILP.

Car size	**Max charge rate**	**Battery size**	**Percentage of EVs**
Small	3.7 kW	20 kWh	9%
Medium	7.2 kW	40 kWh	45%
Large	11 kW	70 kWh	46%

Table 2: Specifications of visiting EVs.

3.3 Predicting driver behavior

A final set of parameters that needs to be estimated is made up of parameters that are completely determined by driver behavior. Parameters such as arrival times, duration of stay, and a baseline demand for energy heavily depend on the behavior and expectations of individual drivers. Therefore, these parameters are difficult to estimate without input from drivers. However, these parameters can still be predicted for businesses with a relatively large group of employees that arrive and leave at fixed times. Figure 3 is generated by another (mock) dataset of around 3 months from Groendus and indeed shows the relatively 'predictable' behavior of employees. Drivers are divided into two categories: employees and visiting customers. As can be seen in Figure 3, employees typically arrive around 9:00 in the morning and leave 8 hours later, at 17:00 in the afternoon. Another slight peak of arrivals and departures around 13:00 indicates that either some employees work half days, or this is a peak hour for customers coming in for a shorter time. Truncated normal distributions for duration of stay for employees and customers are used in our model, where employees are assumed

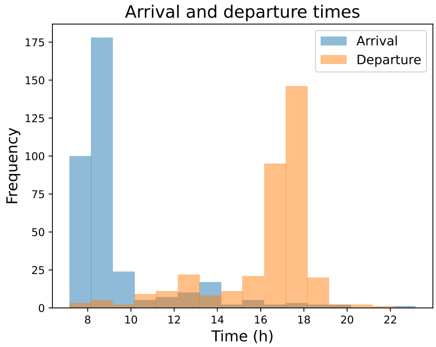

Figure 3: Arrival and departure times over a period of 3 months for a typical business working with Groendus.

to stay 8 hours on average, and customers are assumed to stay only 1 hour. In practice, however, the duration of stay needs to be adjusted on a case-by-case basis. For instance, Figure 4 shows that this particular business does not have many customers visiting the charging stations, but this could, of course, differ per location. Hence, the estimated number of customers and employees in a day can be easily tuned in the synthetic data generation, which serves as input for the MILP. In the upcoming section, we demonstrate the performance of our model for different configurations.

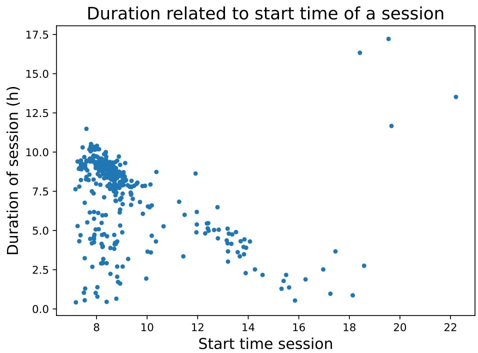

Figure 4: Arrival times of drivers related to their duration of stay.

Although crude estimations of driver behavior are sufficient for implementing the MILP, these estimations can be improved. In Section 5, an outline is given on setting up an online model that updates predictions on driver behavior in real-time. Moreover, driver expectations should also be taken into account in the form of a baseline demanded amount of energy charged over a period of time. This demand depends both on the company and the driver, and hence is a parameter that is not determined through estimation, but through conversation with all involved parties. In Section 4, several scenarios for the baseline demand of EV drivers are shown.

4 Implementation & results

4.1 Implementation details

We implemented the MILP using Python, using Google's OR-Tools (Perron and Furnon) as a wrapper to call the SCIP-solver (Bolusani et al. (2024)). The program is separated into five parts. The first part generates synthetic data used as input for the model. In particular, EV, weather, and energy price data are generated based on the data provided by Groendus, using the methods presented in Section 3. The second part contains the implementation of the model presented in Section 1. The third part focuses on the visualization of the calculation output. The fourth part contains calls to single scenarios. The fifth and last part focuses on generating test series with the option to save and load the results.

4.2 Experiments

Next, we show the empirical results of the behavior of the solver. We first explore four scenarios in detail in Section 4.3, representing two different weather and user scenarios, and visualize the trend of charging in a day, depending on available solar surplus, energy price, and present cars, as well as the resulting cost. After that, the test series results are given in Section 4.4.1, which explores the trade-off between cost and total charge ratio. Finally, we compare the fair and unfair distribution of energy in Section 4.5 to justify the additional implementation of the fairness factor into the objective function in Section 2.

4.3 Scenarios

Two weather scenarios are explored. We consider a sunny day, where the station produces a lot of solar energy and energy prices are lower. Furthermore, we consider a cloudy day, where energy prices are lower and solar production rarely satisfies the station's demand. We combine these scenarios with two different driver-type settings. In the first setting, we have 40 employees, who typically arrive in the morning and leave in the late afternoon; in the second setting, we have 40 customers who arrive randomly during the day and only stay a few hours. Additionally, we decided to set the number of electric chargers at a station equal to the number of cars, as we optimize charging schemes for plugged-in EVs rather than the scheduling of EVs at charging stations themselves.

This creates a total of four scenarios, which are explored below in Sections 4.3.1 to 4.3.4. For each scenario, we display four plots. The first two show, for a fixed target demand, how much energy we use during the day compared to the current solar surplus of the station, the maximum possible charge rate, and the energy price. The latter two show the effects of different target demands on the charging curve during the day and how the cars are individually being served.

Additionally, we compare the cost to the fulfilled demand trade-off of the four scenarios in Figures 21 and 22.

4.3.1 First scenario: 40 employees, sunny

In this scenario, we have a lot of free energy available and employees who arrive in the morning and stay until late afternoon. Figure 5 shows the charge curve during the day, set to the maximum possible charge rate at any given moment, and the solar surplus we generate during the day, that is, the total solar energy produced minus the energy used by the building. The maximum possible charge rate is strongly related to the number of cars at a charger. We note that we only charge when we have solar surplus and do not use grid energy to charge the cars.

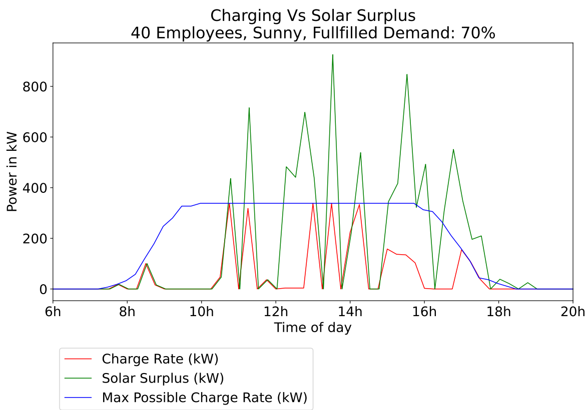

Figure 5: The blue line represents the maximal charge rate to potentially charge currently parked cars. The green line represents the solar surplus received from solar panels after fulfilling the station's needs and is, at the same time, the energy we could use either to charge the cars or to sell to the grid. The red line represents the amount of power we charge the cars with. If the red line is lower than the green line, we sell to the grid, otherwise we buy from the grid. The target demand of 70% has been fulfilled without buying additional energy from the grid.

Moreover, Figure 6 shows the cumulative charge curve during the day and sets it with the current price. Combining the information of Figures 5 and 6, we conclude that we sell energy back to the grid between 12:00 and 13:00 and between 15:00 and 17:00 when the energy price is high and charge when the price is low.

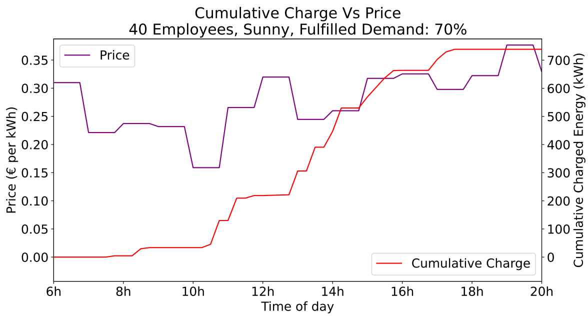

Figure 6: In purple, we see the fluctuating energy price per kWh during the day. The revenue for selling energy would be 20 cents lower. In red, we see the energy we charge during the day, which is charged mostly when the price is low or (not depicted) when the solar surplus is high.

Figures 7 and 8 show the behavior for different fulfilled demands, in 10% point steps and the respective cost. Figure 7 shows the charging curves, this time one for each fulfilled demand, and at which moments these additional demands would be fulfilled. For example, 50-80% only differ by the amount we charge around 15:00, which coincides with high energy surplus at that time (see Figure 5).

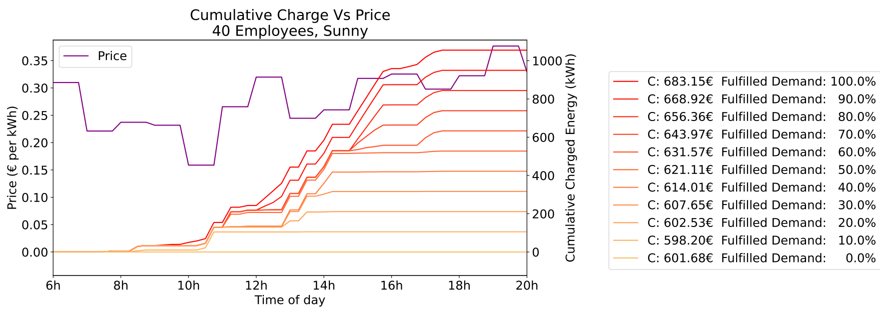

Figure 7: In purple, we see the fluctuating energy price per kWh during the day. In red, we see the energy we charge during the day for different target fulfillments. We can see several bifurcations of moments where the solver in one scenario would charge, whereas the other one would try to maximize profit.

Figure 8 shows the energy distribution towards different cars. One can see that the last 10% only serves a handful of cars which have been already charged quite fully, additionally one sees a fair distribution of energy. Furthermore, we notice that the total cost drops when the demand increases from 0 to 10% because of a negative selling price and high solar surplus around 10:30.

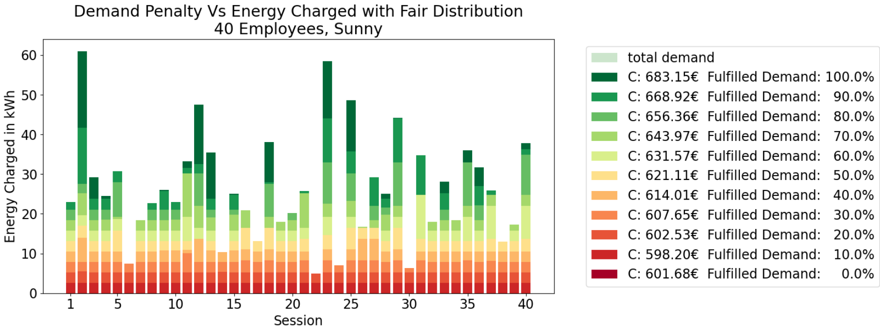

Figure 8: We can see the energy distribution towards different cars/sessions in a stacked bar plot. Every bar represents the additional energy we would charge for a target charge value compared to the next lower target charge value. For example, session 2 would be charged about 40 kWh if we set the target fulfilled demand to 90%, whereas it would be charged fully if we set the target demand to 100%.

4.3.2 Second scenario: 40 employees, cloudy

We have the same arrival times in this experiment as in Section 4.3.1, but this time, we have far less energy surplus, with only small peaks around 11:00 and at 17:00, as seen in Figure 9.

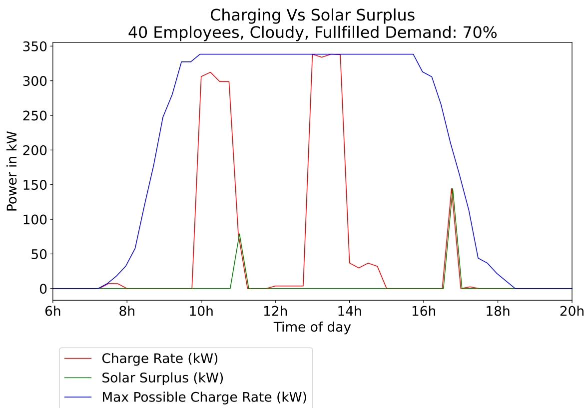

Figure 9: As there is barely solar surplus (in green) and most cars are available the whole day (in blue), the charging is mostly dependent on the current price level as seen in Figure 10 and happens at 10:00 and 13:00 and at 17:00, when surplus is available.

The lower energy surplus is reflected in Figure 10 since the cars are mostly charged at times with relatively low energy prices. This makes this scenario less complex than the previous one.

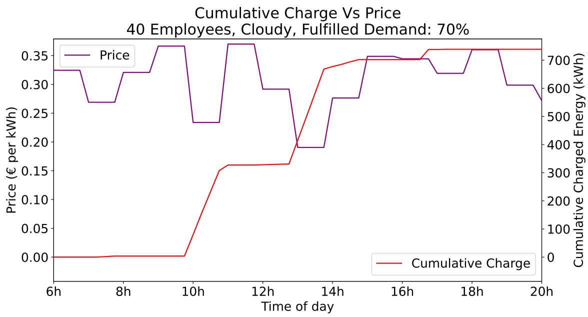

Figure 10: Charging only happens when the price is low around 10:00 and 13:00, besides a short spike around 17:00, which, as Figure 9 shows, is a moment with solar surplus.

The cumulative charging curve for different demands in Figure 11 shows the same reduced complexity. Here, the MILP will charge if the price is below a threshold. The distribution of the energy in Figure 12 is very similar to the sunny case as seen in Figure 8.

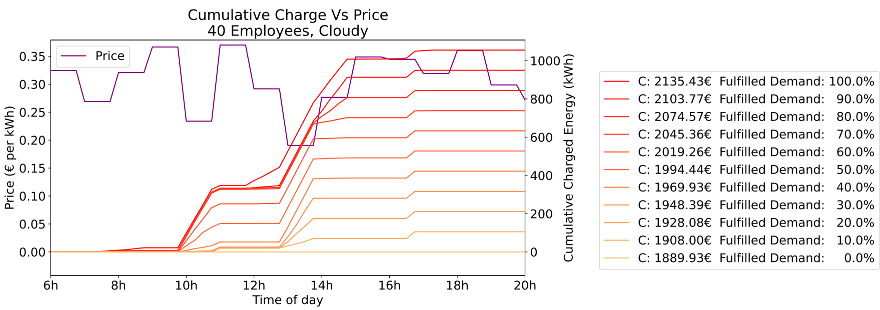

Figure 11: Similar as in Figure 10 the target charge value will determine how strong we charge around 10:00 when the price is second lowest and around 14:00 where we have the third lowest price level.

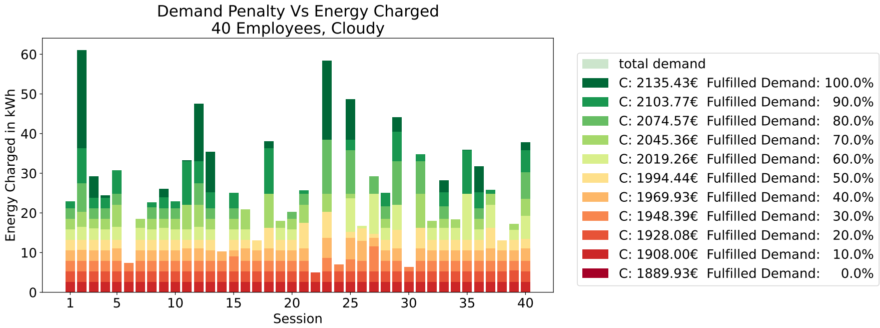

Figure 12: As in Figure 8 described, this plot shows the energy distribution towards different cars/sessions depending on the target charge value. We can see that every car can be charged fully in the 40-employee scenario, and the energy is relatively evenly distributed if the target demand is lower.

4.3.3 Third scenario: 40 customers, sunny

In this scenario, we have 40 customers who arrive over the day and a lot of available energy surplus, see Figure 13. It should be noted here that due to the short time intervals each car is available, it is only possible to charge them overall to 38%, instead of 100% of the employee scenarios.

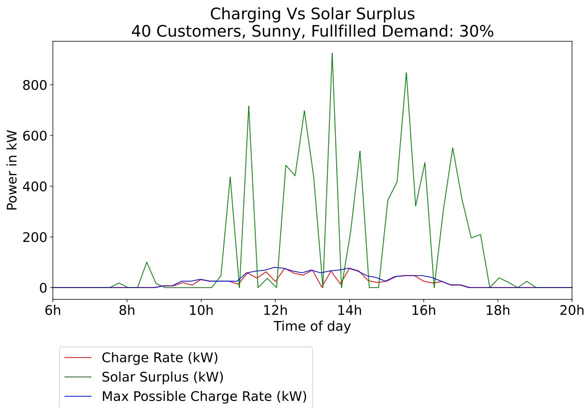

Figure 13: In this scenario, we have a massive solar surplus compared to the available cars, and we charge (in red) as much as possible throughout the day, with a slight decrease in moments where the solar surplus (in green) is lower than the potential demand (in blue).

The overall charge curve also looks closer to a sigmoid function, as the solver has to take every moment possible to charge the cars, see Figure 14. We choose a charge rate of 30% for these two plots to give the solver some potential decisions. It can be seen in Figure 13 that the charge curve adapts itself towards the solar surplus curve.

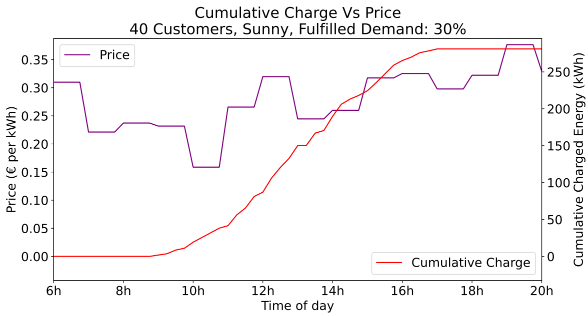

Figure 14: To fulfill the target demand, it will be charged throughout the day, which makes the curve appear like a sigmoid function, with slight dips at moments where the solar surplus is lower than our demand, which could be seen in Figure 13.

The cumulative charge vs. price plot in Figure 15 reflects that a maximal possible charge rate will constantly charge, whereas, for the smaller demands, the charging is reduced during less profitable time slots, e.g., around 13:00 and 14:45. It can be seen in the demand vs. penalty plot of Figure 16 that only three cars could be charged fully; the rest did not remain plugged-in long enough. Although we had a drop in cost when we increased from 0% to 10% in the 40 employees scenario in 4.3.1, it is not visible here, as the drop would be at 0.4%, as only two cars are present at the critical time at 11:00, where solar surplus is present together with negative price.

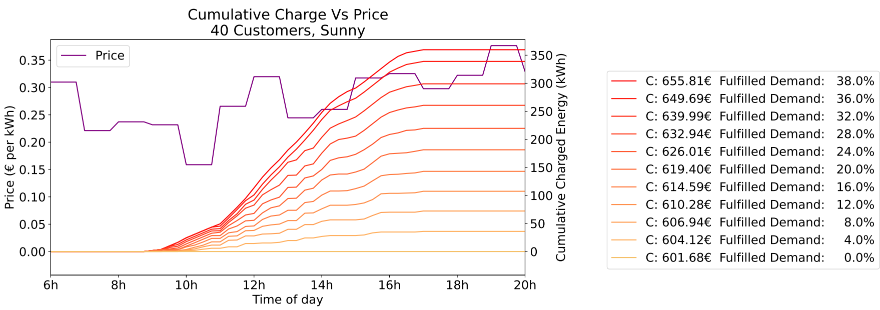

Figure 15: The different charging curves show all the same sigmoid shape, where the maximum of 38% would represent if the red curve would equal the blue curve in Figure 13. The lower charge rates will have dips at the moment where the solar surplus is lower than the maximal charge rate in the figure above.

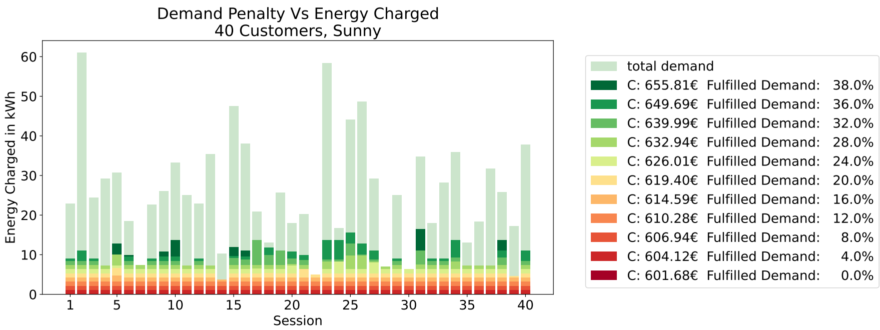

Figure 16: As in Figure 8 described, this figure shows the energy distribution to different sessions/cars at different target levels. As the cars are not available long enough, it is only possible to charge them to 38% in total, though the higher levels are distributed more unfairly due to a lack of availability.

4.3.4 Fourth scenario: 40 customers, cloudy

The cloudy scenario with 40 customers has again 40 cars that arrive during the day and stay for a short duration, but the system has only two peaks of solar surplus at 11:00 and 17:00, as seen in Figure 17. Figure 18 shows that this leads to a less complex scenario than before, where, if the solver has the choice, lower prices are preferred for faster charging, especially seen between 12:00 and 14:00.

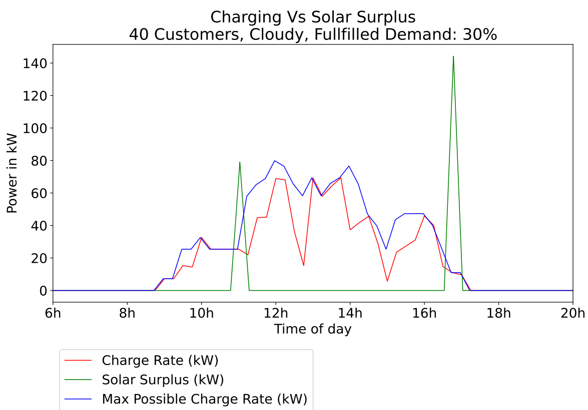

Figure 17: The charge rate in red can be seen as influenced mostly by the available cars. Additionally, the dip around 12:30 can be interpreted as a reaction to leaving cars of the morning, which have been charged before, and cars arriving, which wait to be charged at 13:00.

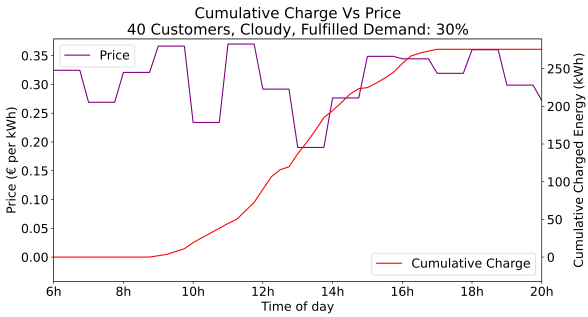

Figure 18: The cumulative charge has a sigmoid nature, although dips can be seen around 12:30, which together with the available cars of Figure 17 represent a moment where cars leave and arrive at the same time.

Figure 19 shows the charging curves during the day for different fulfillment demands. Moreover, it shows the energy price, and in the legend, we provide the total costs to guarantee the respective fulfillment demands. This figure shows smoother charging curves in comparison to Figure 15 , as the solver has to make little adaptations to fluctuating solar surplus levels.

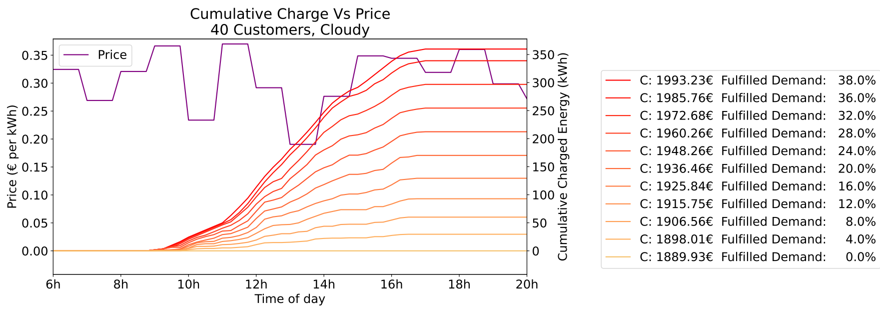

Figure 19: Given the same car distribution as in Figure 15, we perceive a similar pattern, except that the dips are more dependent on the current price and are most noticeable in the lowest price phase around 13:00.

In Figure 20, we again show the distribution of energy over the cars for several fulfillment demand values. We note that we again only can guarantee a charge level of 38% for all cars, due to the short stay of most of the cars. We also display the total demand and notice that only two cars are fully charged at the end of their session.

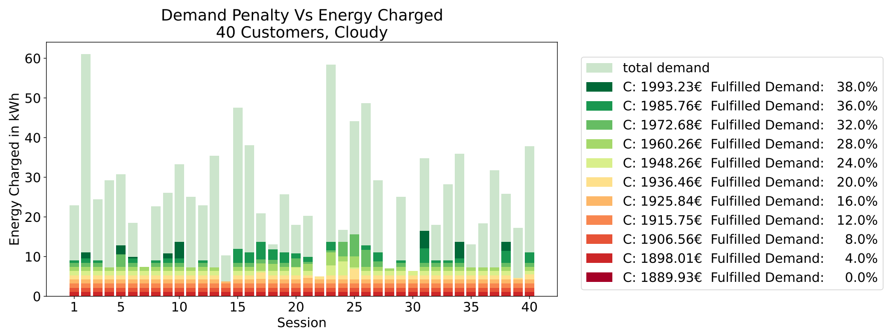

Figure 20: This result is very similar to the sunny scenario in Figure 16, where the charge levels for some sessions are the same for every target charge value. When they differ, it is the result of the incorporation of the solar surplus in the sunny scenario.

4.4 Comparison of demand fulfilment and cost

In this section, we compare the cost with the fulfilled demand, first for the four discussed scenarios and then for a test series. As for the station owner, it is most useful to know how much they would pay if, for example, all of their employees were charged only to half capacity. Figures 21 and 22 show the charge during a cloudy day grows approximately linearly: if one pays twice as much money, then they will receive roughly twice as much energy. During sunny days, it is superlinear, which can be interpreted as every kWh not charged could instead be sold, whereas, on a cloudy day, only the buying can be avoided. Despite the superlinear growth in the sunny case, the total cost of charging cars fully is much lower than in the cloudy case.

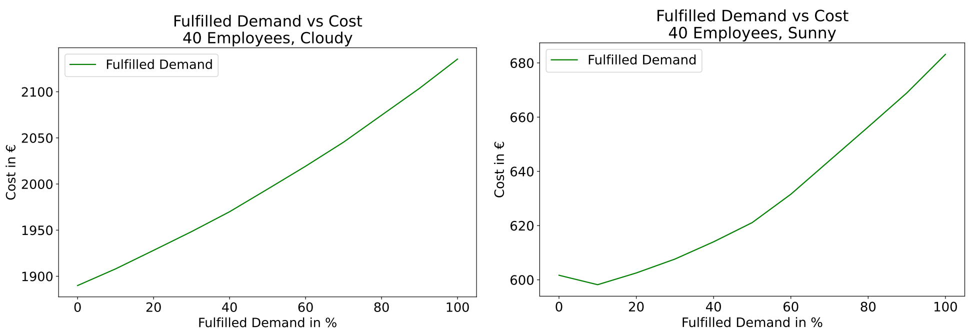

Figure 21: A comparison of the cost to demand fulfillment trade-off for the employee scenario with sunny and cloudy days. The cloudy days show a more linear behavior due to lack of sold energy, whereas the sunny days have a more quadratic shape.

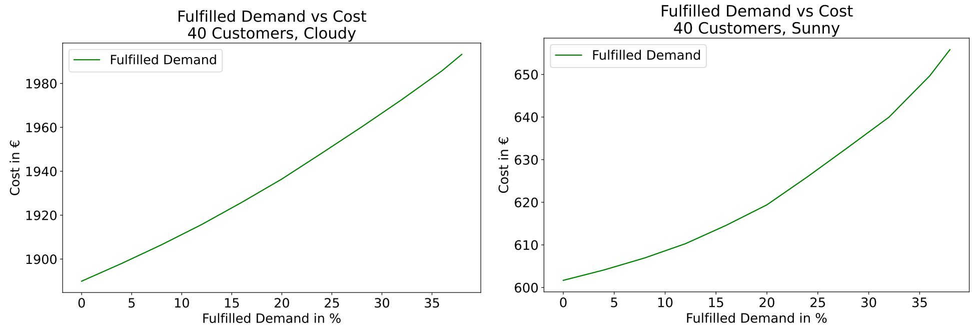

Figure 22: A comparison of the cost to demand fulfilment trade-off for the customer scenario with sunny and cloudy days. . The cloudy days show a more linear behavior due to lack of sold energy, whereas the sunny days have a more quadratic shape.

4.4.1 Test series

Next, we completed ten experiments for each of the four scenarios outlined in Section 4.3, but with different random seeds. We compare the additional cost with the fulfilled demand between sunny and cloudy days for the two 40-customer scenarios and the two 40-employee scenarios in Figure 23. In both cases, one can see a more linear relationship between cost and fulfilled demand on cloudy days and a superlinear relationship on sunny days.

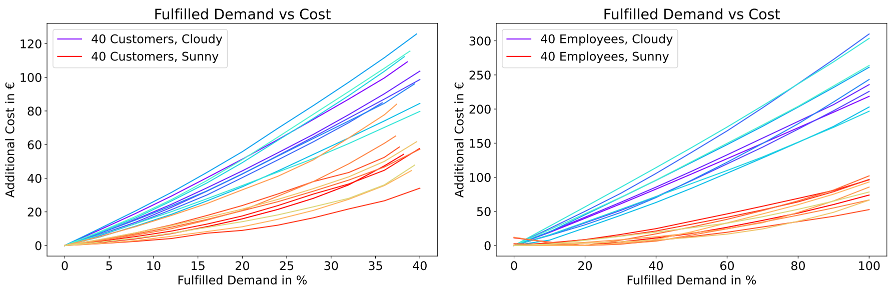

Figure 23: These two plots show ten reruns for each of the experiments mentioned above in Figures 21 and 22. The superlinear shape and the linear shape can be distinguished for cloudy days and sunny days, respectively. Additionally, there is a clear cost separation between cloudy and sunny in a 40-employee scenario, as the solver has more time for each car to optimize.

4.5 Comparison between fair and unfair distribution

Finally, we show the effect of the unfairness penalty. If we enforce fairness, as seen in Figure 24, there is a cost increase of at most €0.40, but the levels of charge for the different charging sessions are more equalized, whereas in Figure 25 the distribution of energy is much more unfair. The downside of adding a fairness factor into the objective function is a 4-5 fold longer processing time, and in rare cases, the solution is not found.

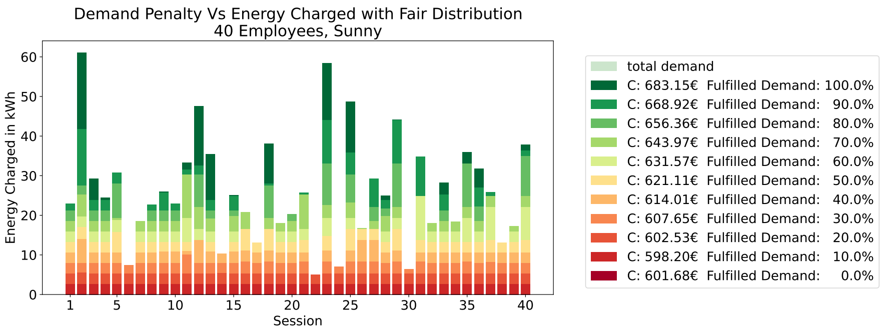

Figure 24: The energy distribution is mostly fair, with differences only appearing in the higher percentages.

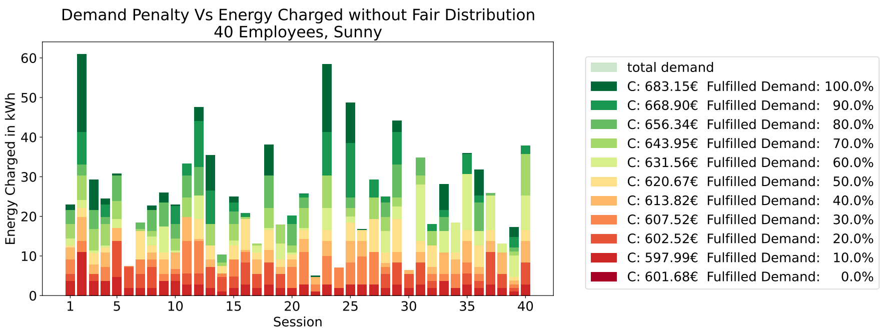

Figure 25: The energy levels here make a more chaotic impression, where some cars are charged twice or thrice as much as others for the same target value.

5 Outline on online model

Driver behavior tends to be reasonably predictable given sufficient data, as discussed before in Section **??**. From the data, we divided drivers into two broad categories: customers and employees. We identified that employee driver behavior tends to be more predictable. Usually, employees arrive around 9:00 in the morning and leave around 17:00 in the afternoon. On the other hand, customers arrive during the entire day, but we could still provide a crude estimate of customer arrival and session time. This average driver behavior provides a decent baseline for the linear program to compute a charging scheme for a given day.

Even though this average driving behavior gives decent data to the linear program, individual data could be exploited to obtain more accurate online predictions. Individual driving behavior can be best predicted if data is plentiful. Although customers may visit a company only once, employees are typically more predictable recurring drivers. Given enough data, a driver behavior profile can be constructed for every individual driver. For example, we may find that a certain employee always arrives before 9:00. If this employee still has not arrived at 10:00, it is reasonable to update our driver prediction for the day by removing this driver from our prediction as this employee will likely not show up anymore.

Assume that for some driver d, we estimated from previous data that this driver arrives with probability a_d, i.e., the arrival probability of a driver d, A_d, is Bernoulli distributed with parameter a_d. Furthermore, assume we estimated from previous data that this driver arrives at a normally distributed time $T_d \sim \mathcal{N}(\mu_d, \sigma_d)$. Then, at the start of the day, the estimated arrival time is $\mathbb{E}[T_d] = \mu_d$. However, this prediction can be updated at every time point t. If a driver has not arrived yet at time t, their conditional expected arrival time becomes

$$\begin{aligned}
\mathbb{E}\left[T_d | T_d \geq t\right] &= \mathbb{E}[\sigma_d Z + \mu_d | Z \geq t] \qquad\qquad \text{with } Z \sim \mathcal{N}(0,1) \\
&= \sigma_d \mathbb{E}[Z|Z \geq t] + \mu_d \mathbb{E}[1|Z \geq t] \\
&= \left(\sigma_d \int_t^\infty z\phi(z)\,dz + \mu_d \int_t^\infty \phi(z)\,dz\right)\mathbb{P}(Z \geq t) \\
&= \left(-\sigma_d \int_t^\infty \phi'(z)\,dz + \mu_d \int_t^\infty \Phi'(z)\,dz\right)(1 - \Phi(t)) \\
&= \frac{\sigma\phi(t)}{1-\Phi(t)} + \frac{\mu_d(1-\Phi(t))}{1-\Phi(t)} \\
&= \mu_d + \frac{\sigma_d \phi(t)}{1-\Phi(t)} \\
&> \mu_d,
\end{aligned} \tag{9}$$

where ϕ is the standard normal density function and Φ is the cumulative standard normal density function. The departure times of a driver could also be updated at every time step analogously. Furthermore, drivers could be removed from a prediction if the probability they would still arrive after time t is small enough. That is, if the current time t satisfies

$$\mathbb{P}(T_d \geq t) < \epsilon(a_d), \tag{10}$$

driver d should be removed from the prediction. This $\epsilon(a_d)$ takes values in the interval $[0, 1)$ and determines when employees should be removed. Two intuitive choices for $\epsilon(a_d)$ are either $\epsilon(a_d) \equiv \epsilon$ for some constant $\epsilon \in [0, 1)$ or $\epsilon(a_d) = 1 - a_d$. With a constant cut-off probability, employees get removed if they are unlikely to arrive after some time, if they arrive at all. With $\epsilon(a_d) = 1 - a_d$, on the other hand, $\epsilon(a_d) = 1 - a_d$, employees that are unlikely to arrive will be removed from the prediction earlier on. This may remove employees from the prediction before their expected arrival time if the arrival probability a_d is low. No matter the choice of $\epsilon(a_d)$, a cut-off time τ_d can be defined as the minimal t satisfying Equation 10. Therefore, at time t, the total number of employees in our prediction is given by

$$\begin{aligned} n(t) &= \text{Arrivals before time } t + \text{Expected arrivals after time } t \\ &= \sum_{d=1}^{D} A_d \left[\mathbb{1}_{\{T_d<t\}} + \mathbb{1}_{\{T_d\geq t\}} \cdot \mathbb{1}_{\{t<\tau_d\}}\right]. \end{aligned} \tag{11}$$

Notice that the total number of employees predicted to arrive is always smaller than the total number of employees D. The predicted number of employees that arrive during the day at time t is a priori given by the expectation of $n(t)$. Assuming that the arrival probability A_d and arrival time T_d are independent, we find that

$$\begin{aligned} \mathbb{E}[A_d\left[\mathbb{1}_{\{T_d<t\}} + \mathbb{1}_{\{T_d\geq t\}} \cdot \mathbb{1}_{\{t<\tau_d\}}\right]] &= a_d\left[\mathbb{P}(T_d < t) + \mathbb{P}(\{T_d \geq t\} \cap \{t < \tau_d\})\right] \\ &= \begin{cases} a_d & \text{if } t < \tau_d \\ a_d \cdot \mathbb{P}(T_d < t) & \text{else.} \end{cases} \end{aligned} \tag{12}$$

Therefore, we find in that $\mathbb{E}[n(t)] \leq \sum_{d=1}^{D} a_d$, with equality for $t < \min_d(\tau_d)$. At the start of the day, the expected number of predicted employees is given by $\sum_{d=1}^{D} a_d$, which is equal to the expected number of employees to show up during that day.

Intuitively, this result means that we draw a Bernoulli variable with parameter a_d for every driver d and assume arrival times at μ_d. During the day, these μ_d's are updated according to Equation 9. Furthermore, the number of predicted employees is updated every time if either an employee arrived that was not in our current prediction or the current time passes one or more cut-off points τ_d.

As a next step, this online prediction scheme could be tested as input for the linear programming model. For this, synthetic or real data could be used to model the driver behavior of employees, i.e., we estimate the arrival probability a_d, average arrival time μ_d, and variance σ_d^2 for every driver d. Consequently, a Bernoulli variable with probability a_d is drawn for every driver. Only drivers for which this Bernoulli variable is equal to 1 are included in the initial prediction. After every time step, the predicted driver behavior is updated by adding and removing drivers and updating their expected arrival and departure times. The performance of this online model could then be compared with the offline model, which only considers the initial prediction. Furthermore, it would be interesting to try several options for $\epsilon(a_d)$ to find the best performance of this online prediction model.

6 Discussion

The goal of this project is to improve the EV charging schemes currently used by Groendus. The current scheme is a greedy algorithm that always charges maximally if the current energy price is below a certain threshold. In practice, this often resulted in maximal charging rates in the morning, while the charging stations were passive afterward when the production of solar power was largest. As the cars are not charging for the entire session, there is flexibility for improvements (Figure 2). Ideally, our new model would both predict driver behavior and provide instructions when charging stations should be active. We simplified this model by decoupling the problem into an optimal charging problem and data prediction.

First, we set out to define what was considered to be a 'good' charging scheme. Groendus prefers to maximize sustainability above all else. They hypothesized that this objective would result in the cheapest charging scheme as well. Aside from this, secondary side objectives were: minimizing the overall cost of the charging station, satisfying driver demands, and distributing delivered energy fairly among cars. All these objectives were implemented in a MILP model by penalizing buying energy from the grid, especially if the grid prices are high, penalizing unsatisfied customer demands, and, thirdly, penalizing an unfair distribution of charge among the drivers.

Secondly, we estimated the state of charge for arriving cars and the driver behavior. We found three distinct groups of cars with different charge rates, which were, furthermore, highly correlated with the energy charged in a session (Figure 1). Moreover, the driver behavior of employees was found to be very predictable as historical data shows that employees generally arrive around 9:00 in the morning and stay for approximately 8 hours (Figure 2). Customer behavior was more difficult to predict due to limited historical data.

Subsequently, we used our parameter estimations to test the MILP model in four different settings. We tested our model for both sunny and cloudy days in setups with only employees or only customers. On sunny days, car charging follows the surplus of energy during the day (Figure 5). Further, the surplus of solar energy is mostly sold to the grid when the energy price is high (Figure 6). However, the energy price affected the charging schedule less if there was higher energy demand. Even so, cars were primarily scheduled to charge if a surplus of solar energy was available and secondarily if the price was low (Figure 7). On cloudy days, on the other hand, charging stations were forced to buy energy from the grid if they wanted to satisfy customer demands. If a surplus of solar energy was temporarily available, it was used to charge cars (Figure 9), but the energy prices were dominant in deciding the charge rates (Figure 11). We find similar trends when we consider a setup with many customers instead. The surplus of energy followed the solar surplus on sunny days (Figure 13), and on cloudy days, the charging scheme followed the energy price more closely (Figure 19). However, these trends were less pronounced than in the employee setting. Furthermore, even if we set our target charge rate to 100%, most cars could not charge fully due to their limited session times (Figures 16 and 20). In addition, we compared the price increase for different target demands and found that this trend was almost linear on cloudy days, but superlinear on sunny days (Figures 21, 22 and 23). Our final tests of the MILP model studied the role of our fairness penalty in the objective function. As expected, this additional term encouraged an equal distribution of charge among cars (Figure 24), whereas an unequal charge distribution was found without this penalizing term (Figure 25).

Lastly, we briefly considered an online prediction model. Given enough data, we could update the arrival times of employees as time progresses. This online prediction model is still in the early stages of development and has yet to be tested, but it could be used to improve inaccurate predictions in real time.

In contrast to the current implementation from Groendus, our MILP model primarily follows the solar surplus, and secondarily follows the energy price. Therefore, this model satisfies the wishes set by Groendus, as their main focus was to optimize the use of green energy.

Furthermore, our fairness penalty distributes charges equally among the cars if the state of charge is predicted accurately. Importantly, this fair distribution of charge barely influences the total costs and could, therefore, be implemented without additional costs. The model solution is most strongly correlated with solar surplus and energy prices in a setting with many employees because this offers the most flexibility in charging times. For customers that stay approximately an hour, we could do only so much tweaking in the charging schemes. In this setting, maximal charging is incentivized to approach the target demand as much as possible, leaving limited options for the solar surplus or energy price to affect the charging scheme. We also found that the total cost scales non-linearly with the target demand, especially on sunny days. The reason for this difference between weather settings is that on cloudy days, costs can only be reduced by not buying from the grid, whereas on sunny days, costs can additionally be reduced by selling solar surplus to the grid instead of using it to charge cars. This target demand, therefore, offers a slider to be determined by the charging station owner. On the one hand, costs can be reduced by a low target demand. On the other hand, this may lead to undesirable selling to the grid. Alternatively, this target demand may be set higher on sunny days, determined by the expected solar surplus to reduce interactions with the grid.

In the literature, different motivations for coordinated charging schemes were given, which could lead to different outcomes. In Ghotge et al. (2020), they aim to reduce local peak electricity demand, in Franco et al. (2015) they maximize energy charged and minimize power losses, and in Dukpa and Butrylo (2022) they focus on maximizing profit. This interplay between sustainability and cost reduction comes at a cost for the charging station owner. The greenest option of using all solar energy is, of course, the most sustainable option, but it may be cheaper to sell solar energy to the grid when the price is high and purchase from the grid when the price drops. However, this greedy approach is probably undesirable, given the grid congestion problem in the Netherlands. Ultimately, this is a choice that Groendus should discuss with their clients. Some companies may be satisfied with not fully charging if this drastically reduces their costs, while others may prefer to achieve maximal sustainability. The preferences of a company may also depend on whether external fueling is also paid for by the company, as charging externally is usually more expensive.

Although our MILP model gives promising results, we assume perfect knowledge in our model. In reality, only predictions can be made. We have not tested how well our model would perform if the predictions are inaccurate. The online model prediction could be implemented to test the model performance with online updates and incomplete knowledge. In fact, we believe that a new prediction should be made at every time point for the model to perform well in a real setting. If a car were to arrive earlier than expected when solar surplus is high, we would not want to wait to charge the car just because our prediction is that this car would arrive at a later time. This lack of perfect predictions may also complicate the fairness penalty. It is more difficult to fairly distribute charge if cars may leave earlier than expected. Additionally, this fairness penalty increases the computational complexity of the model significantly. The fairness penalty may, therefore, be impractical in real applications. Aside from uncertainty in predictions, the initial model predictions we provided may be inaccurate as well. Realistic scenarios were created based on synthetic data, but this synthetic data lacked customer cars. Therefore, it is difficult to make strong claims about the model's performance in a customer setting. We also have not included hybrid setups with both employees and customers, but we suspect the charging scheme follows the solar surplus and energy pricing more strongly than in the customer setting, but weaker than in the employee setting. Lastly, our research on online updating is limited. Although we have provided a method to update customer arrival times, the exclusion criterion is not fully understood. We have provided several criteria options for a cut-off time, but are uncertain about the performance of these options.

A key point for future study is to test how well this model would perform with incomplete data. This will, of course, depend on the accuracy of predictions. A sensitivity analysis could be performed to study which predictions are most important. Consequently, effort could be put into improving these predictions. Furthermore, we suggest implementing the online predictions to study if these actually improve performance, as we suspect. We advise testing at least the two suggested exclusion criteria to study the performance of both options, although better options may be available. This online model might require a lot of work to become fully functional. A simple, intermediate solution is to update the initial prediction once most employees have arrived. Most charging occurs after this time anyway, thus offering an acceptable heuristic approach. Also, the issue of a suitable target demand should still be resolved before enrolling this model in practice. A good baseline for this target demand would be to set this equal to the predicted solar surplus, but provide a minimal demand on cloudy days to satisfy customers. Lastly, we limited our study to AC chargers. The same methods can also be applied to a study with DC chargers, except that these chargers allow the charge in cars to be used as batteries for the company. Effectively, this would allow the charging rates at a charger i, E_t^i, to be negative as well, hence adding an extra constraint similar to Constraint 4. Furthermore, extra constraints should be added to reduce the risk that drivers leave with an empty car.

7 Conclusion

We set out to find a better EV charging strategy than the current scheme employed by Groendus. This is achieved by providing predictions of cars and behavior on the one hand and implementing a MILP model on the other hand. This MILP model prioritizes using green energy over minimizing costs, which is a secondary objective. This expected behavior is indeed found when testing the MILP on synthetic data, assuming perfect predictions. The charging schemes follow the trends of solar surplus and energy pricing, especially if cars stay plugged in for a long time. Although this assumption of perfect predictions is unrealistic, it is promising that our model performs well in this case. Proper, possibly online, model predictions are still required to implement this MILP model in practice and to answer the question of whether to charge or not to charge.

References

Suresh Bolusani, Mathieu Besançon, Ksenia Bestuzheva, Antonia Chmiela, João Dionísio, Tim Donkiewicz, Jasper van Doornmalen, Leon Eifler, Mohammed Ghannam, Ambros Gleixner, Christoph Graczyk, Katrin Halbig, Ivo Hedtke, Alexander Hoen, Christopher Hojny, Rolf van der Hulst, Dominik Kamp, Thorsten Koch, Kevin Kofler, Jurgen Lentz, Julian Manns, Gioni Mexi, Erik Mühmer, Marc E. Pfetsch, Franziska Schlösser, Felipe Serrano, Yuji Shinano, Mark Turner, Stefan Vigerske, Dieter Weninger, and Lixing Xu. The SCIP Optimization Suite 9.0. Technical report, Optimization Online, February 2024. URL https://optimization-online.org/2024/02/the-scip-optimization-suite-9-0/.

Andu Dukpa and Boguslaw Butrylo. Milp-based profit maximization of electric vehicle charging station based on solar and ev arrival forecasts. *Energies*, 15(15):5760, 2022.

John F Franco, Marcos J Rider, and Ruben Romero. A mixed-integer linear programming model for the electric vehicle charging coordination problem in unbalanced electrical distribution systems. *IEEE Transactions on Smart Grid*, 6(5):2200–2210, 2015.

Rishabh Ghotge, Yitzhak Snow, Samira Farahani, Zofia Lukszo, and Ad van Wijk. Optimized scheduling of ev charging in solar parking lots for local peak reduction under ev demand uncertainty. *Energies*, 13(5):1275, 2020.

Mehmetkr Şükrü Kuran, Aline Carneiro Viana, Luigi Iannone, Daniel Kofman, Gregory Mermoud, and Jean P Vasseur. A smart parking lot management system for scheduling the recharging of electric vehicles. *IEEE Transactions on Smart Grid*, 6(6):2942–2953, 2015.

Steffen Limmer. Evaluation of optimization-based ev charging scheduling with load limit in a realistic scenario. *Energies*, 12(24):4730, 2019.

Laurent Perron and Vincent Furnon. Or-tools. URL https://developers.google.com/optimization/.