Honda Research Institute EU
Honda Research Institute Europe GmbH
https://www.honda-ri.de/

Design of Fair and Interpretable Electric Vehicle Charging Policies through Genetic Programming

Steffen Limmer, Angus Kenny, Tapabrata Ray, Felix Lanfermann, Hemant Kumar Singh, Andrea Castellani

2025

Preprint:

Copyright Elsevier 2025. This manuscript version is made available under the CC-BY-NC-ND 4.0 license http://creativecommons.org/licenses/by-nc-nd/4.0/

Design of Fair and Interpretable Electric Vehicle Charging Policies through Genetic Programming

Steffen Limmer[a,*], Angus Kenny[b], Tapabrata Ray[b], Felix Lanfermann[a], Hemant Kumar Singh[b], Andrea Castellani[a]

[a]*Honda Research Institute Europe GmbH, 63073 Offenbach am Main, Germany*
[b]*The University of New South Wales, Canberra ACT 2600, Australia*

Abstract

Controlled charging of a group of electric vehicles (EVs) subject to power limits, such as those imposed by transformer capacity constraints, may result in inefficient and unfair energy distribution among EVs. The present work aims to learn fair and efficient charging policies on historical data using multi-objective genetic programming (GP). A formula, or model, is evolved through GP that takes features of connected EVs as input variables and computes scores to guide the distribution of available energy. Two variants of this approach are proposed and evaluated in simulation experiments, considering a residential charging scenario. In this setting, the dissatisfaction of EV users resulting from a certain charging control policy is quantified as the additional time they required for charging externally, compared to uncontrolled charging. The efficiency and fairness of a charging control policy are measured as the mean and maximum additional external charging time over all users. In the simulation experiments, the proposed approach is compared to several baseline methods, including different manually designed charging policies, such as equal distribution and first-come-first-served, as well as a recent approach for combining multiple fixed charging policies. The experimental results show that the proposed approach increases efficiency by at least 13 % and fairness by at least 12 %. An analysis of the automatically designed policies in terms of interpretability concludes that the best-performing policies are highly complex, containing more than 13 variables and more than 16 operators, on average. However, it is shown that it is possible to significantly reduce this complexity without substantial loss in the quality of the charging control, through appropriate control of the maximum GP tree size.

Keywords: electric vehicles, fairness, genetic programming, multi-objective optimization, explainable artificial intelligence

1. Introduction

In 2024, more than 20 % of all new cars sold globally were electric vehicles (EVs). Compared to 2023, the number of EV sales increased by around 25 % [1]. The increasing number of EVs imposes issues in the operation of the power grid. One such issue is that transformers at building or street level are often not able to handle the load arising from the uncoordinated charging of a high number of EVs. It can be expected that this will not change in the near future, since an appropriate extension of the grid is not only expensive but also time-consuming to implement [2]. Besides dynamic pricing approaches, controlled charging—also termed smart charging—is considered a promising solution for the successful grid integration of EVs [3].

This work considers the problem of charging control for a group of EVs that share a power connection with limited capacity. The connection cannot simultaneously meet the demand of all users, which may lead to reduced user satisfaction. The aim is to design a suitable charging control strategy that minimizes this dissatisfaction and improves fairness among users. A popular approach is the computation of charging plans based on optimization, for example, with mixed integer linear programming [4]. However, this typically requires prior knowledge of users' departure

*Corresponding author
Email address: steffen.limmer@honda-ri.de (Steffen Limmer)

times and energy demands, and it is reasonable to assume that users are often unable, or unwilling, to provide this information accurately. Rule-based approaches represent a valid option for EV charging control in the absence of this information. They also have the advantage of being less compute-intensive than optimization-based approaches.

Typically, efficient charging control that results in high overall user satisfaction is desired. However, high efficiency does not necessarily guarantee fairness, since a high total satisfaction may still be unevenly distributed among users. The present work proposes a novel approach for the automated design of efficient and fair policies (i.e., rules) for rule-based EV charging control. By investigating whether this approach improves residential EV charging experiences compared to existing control approaches, this research aims to contribute to the development of smarter, more sustainable charging solutions that can mitigate strain on the grid while supporting the widespread adoption of EVs. It is proposed to learn charging policies on historical data with the help of genetic programming (GP) [5], a variant of evolutionary algorithms that evolves programs or expressions. To balance the trade-off between efficiency and fairness, multi-objective GP is employed. In recent years, GP has received increasing attention in the context of explainable artificial intelligence, since it produces white-box models [6]. The transparency of the EV charging control strategy can be considered to contribute to the users' trust in the system. Two variants for applying the evolved GP models to charging control are proposed. In the first variant, a GP model is used to compute a priority-based ordering of the connected EVs; in the second, it is used to determine the shares of available energy allocated. The two variants are evaluated in simulation experiments and compared to different manually designed charging policies, as well as a previously proposed approach that combines multiple existing policies to create new ones [7]. Furthermore, the evolved GP models and their complexity are analyzed in more detail. The main contributions of this work are:

- Proposal and evaluation of two variants of an approach for the automated design of efficient and fair EV charging policies based on multi-objective GP
- Analysis of the resulting GP models, especially in terms of their interpretability
- Investigation of the tradeoff between model complexity and charging control quality

The rest of the paper is organized as follows: Section 2 discusses related work on fair EV charging control. Section 3 describes the considered problem setting. Section 4 introduces the proposed approach. Section 5 describes the setup of the experiments and discusses their results. Finally, Section 6 provides a summary and conclusion.

2. Related Work

Several approaches for EV charging control can be found in the literature [8, 9]. The most popular are rule-based and optimization-based approaches. Rule-based charging control allocates charging power to EVs according to a typically manually designed rule or policy. Rule-based controllers are usually applied for real-time control, allocating charging power only for the current time period, without planning into the future. In contrast, optimization-based approaches optimize charging power allocation with respect to a certain objective, subject to various constraints. For the actual optimization, methods such as linear programming (LP), mixed integer linear programming (MILP), and mixed integer quadratic programming (MIQP) are often employed. Optimization-based approaches typically plan ahead, and thus they commonly require knowledge or a prediction of the system's future state. Specifically, they usually require knowledge of the EVs' departure times and energy requirements. More recently, reinforcement learning-based approaches have emerged as an alternative to traditional rule-based and optimization-based methods. An overview of such approaches can be found in [10]. Several studies consider fairness in EV charging control under limited power, which is discussed in more detail below.

It is very common to express user satisfaction based on the amounts of charged and requested energy and to try to distribute the satisfaction as equally as possible over the users. For example, Ottesen [11] proposes an MIQP approach that takes fairness into account by formulating the objective function in a way that the relative deviations from the desired amounts of energy are distributed as equally as possible over the EVs. An analogous approach was proposed by Phan et al. [12], who additionally developed a distributed approach to speed up the solution of the optimization problem. Aswantara et al. [13] formulate a MILP model to minimize charging costs while ensuring that the unmet ratio of the target state of charge (SOC) is distributed equally among the EVs. Zeballos et al. [14] compare different charging policies, such as earliest-deadline-first or least-laxity-first, for charging control under limited capacity. They

find that while the total amount of unfinished work is the same across all considered policies, some distribute this unfinished work more unfairly than others. In their study, a distribution is considered fair if the ratio between the missing amount of energy and the requested amount of energy is equal for all EVs. To promote a fairer allocation of available energy, they propose a new policy called "least laxity ratio". Hussain and Musilek [15] investigate EV charging during power outages, where energy is limited to that stored in stationary batteries or generated by photovoltaic (PV) systems. They introduce the concept of "essential energy demand": for EVs performing critical tasks, such as ambulances, this is the energy required to complete the upcoming task(s); for other EVs, it is the energy needed to reach a charging station unaffected by the outage. The authors propose a rule-based charging control approach designed to maximize the satisfaction of essential energy demand and to distribute it equally across EVs.

Similar to the previously described works, Danner and de Meer [16] seek an equal distribution of satisfaction. However, they do not measure satisfaction solely by delivered and requested energy, but instead define several quality-of-service metrics for EV charging. They propose a charging policy called dynamically weighted fair queuing to achieve an equal distribution of quality of service among EVs.

Several works adopt the so-called "min-max fairness principle", in which the goal is not an equal distribution of satisfaction, but rather to maximize the satisfaction of the least satisfied individual. One example is the work of Frendo et al. [17], who consider workplace charging and assume that all users want to charge to 100 % SOC. They present a MILP formulation for day-ahead charging planning together with a heuristic for real-time adaptation, which prioritizes charging all EVs to a minimum SOC rather than to full SOC. In [18], the heuristic approach from [17] is extended by incorporating a prediction of the EVs' departure times. In [19], another extension is proposed, which considers nonlinear charge profiles. Shariatzadeh et al. [20] also consider workplace charging of a group of EVs with time-of-use energy prices. They propose a multi-objective MILP approach to compute charging schedules, where the first objective is to minimize the total energy cost and the second objective is to minimize the maximum deviation from the desired SOC over all users. The desired SOC values are predicted using a data-driven approach. Khele et al. [21] present a MILP model to plan the charging of multiple EVs with the option of energy transfer between them. They introduce constraints in the formulation to prevent excessive discharge of some of the EVs. Zhou et al. [22] assume that EVs may delay their departure until their respective energy requirements have been satisfied. They propose a charging policy that accounts for this delay, and demonstrate through simulations that this policy outperforms others with respect to efficiency and fairness, where fairness is defined in terms of the maximum delay over all EVs. Tsaousoglou et al. [23] assume an analogous scenario and propose a MILP-based approach to minimize the maximum delay. In addition, they consider that the charging stations are spread over the distribution network and consider network constraints. Limmer and Rodemann [7] present an approach that combines multiple existing charging policies by distributing different fractions of the available energy with the different policies. The values of the fractions are optimized in a multi-objective optimization framework with two objectives: maximizing overall user satisfaction, and maximizing the minimum satisfaction over all users. The (dis-)satisfaction of users is quantified by the time they have to spend for charging at locations other than the one under consideration.

Besides equality and min-max fairness, other notions of fairness have been discussed in the EV charging literature. Gerding et al. [24], for example, define a distribution of energy as fair if it is envy-free. Under this definition, EV i is considered envious of EV j, if i did not receive the requested (and technically feasible) amount of energy, while j received more energy than i during the time i was connected to the charging station. The authors propose two rolling horizon MILP approaches, which aim to achieve a good trade-off between efficiency and fairness, and compare them to rule-based control with different charging policies. Panayiotou et al. [25] address fair EV charging in parking structures with a limited number of available charging spots. They define user utilities based on the satisfaction of their respective charging requirements and seek to allocate utilities fairly among users. As a fairness principle, α-fairness is employed, allowing control of the trade-off between social welfare and min-max fairness via a parameter α. The authors present and evaluate a MILP formulation and an ant colony optimization (ACO) approach for computing α-fair charging schedules. In [26], Panayiotou and Ellinas introduce the "fair-enough" scheme as a less compute-intensive alternative to α-fairness for balancing EV charging efficiency and fairness. Ardakanian et al. [27] address the allocation of power to multiple charging stations in a distribution network, avoiding congestion. They propose a centralized and a distributed approach based on convex optimization for proportionally fair power allocation in each timestep. Under this scheme, an allocation is proportionally fair if it is not possible to increase the power allocated to a station by a certain percentage ϵ without decreasing the powers allocated to other stations so that the percentage decreases sum up to at least ϵ. Carvalho et al. [28] describe a similar approach to proportionally fair power allocation

considering network constraints. Al Zishan et al. [29] also consider proportional fair power allocation, but weight EVs with more urgent charging demands higher than EVs with less urgent charging demands. Furthermore, they suggest penalizing EVs that have declared false charging deadlines in the past. Dreucci et al. [30] adopt a linear programming-based approach for the fair downregulation of charging powers at multiple charging stations to prevent congestion of distribution grid elements. They propose to prioritize stations based on the urgency of the charging jobs of the connected EVs and evaluate the fairness of the distribution of unsatisfied energy demands in terms of proportional fairness.

Table 1: Summary of related work on fair EV charging.

Year	**Fairness Interpretation**	**Satisfaction Metric**	**Charging Type**	**Charging Requirements Known**	**Approach**
[11] (2017)	Equality	Missing energy	General	Yes	MIQP
[12] (2012)	Equality	Missing energy	General	Yes	MIQP
[13] (2013)	Equality	Missing SOC	General	Yes	MILP
[14] (2019)	Equality	Ratio of missing to requested energy	General	Yes	Real-time policy
[15] (2022)	Equality	Charged (essential) energy	Public	Yes	Real-time policy
[16] (2021)	Equality	Quality of Service	General	Yes	Real-time policy
[17] (2019)	Min-max	SOC	Workplace	Yes	MILP and Planning heuristic
[18] (2021)	Min-max	SOC	Workplace	Departure times predicted	Planning heuristic
[19] (2020)	Min-max	SOC	Workplace	Yes	Planning heuristic
[20] (2024)	Min-max	Missing SOC	Workplace	Departure times known and target SOCs predicted	Multi-objective MILP
[21] (2023)	Min-max	Amount of discharging	General	Yes	MILP
[22] (2014)	Min-max	Delay in departure	Residential	Yes	Real-time policy
[23] (2023)	Min-max	Delay in departure	General	Yes	MILP
[7] (2024)	Min-max	External charging time	Residential	No	Combination of real-time policies
[24] (2019)	Envy-freeness	Charged energy	General	Yes	MILP
[25] (2021)	α-fairness	Utility	Public	Yes	MILP and ACO
[26] (2024)	α-fairness and fair-enough scheme	Utility	Public	Yes	MILP
[27] (2013)	Proportional fairness	Power allocation	General	No	Convex optimization
[28] (2015)	Proportional fairness	Power allocation	General	No	Convex optimization
[29] (2020)	Proportional fairness	Power allocation	General	Yes	Convex optimization

[30] (2025)	Proportional fairness	Ratio of missing to requested energy	General	Yes	LP
Present work	Min-max	External charging time	Residential	No	Evolved real time-policy

Table 1 provides an overview of the discussed works on fair EV charging, as well as the present work. It lists for each work the considered fairness interpretation, how (dis-)satisfaction of users is expressed or measured, the considered type of EV charging (public, workplace, or residential) if explicitly stated in the work ("general" otherwise), if the EVs' charging requirements (energy demands and departure times) have to be known in advance, and the proposed approach for the charging control. As can be seen, the majority of the approaches rely on the knowledge or at least a good prediction of the users' charging requirements. In the present work, this information is assumed to be unavailable for charging control, which renders nearly all of the listed approaches inapplicable. The works [27] and [28] do not assume knowledge of user requirements. However, they address fair allocation of power to charging stations distributed in the power grid, and thus the corresponding approaches are not suited for the residential charging scenario considered in the present work. Of the reviewed methods, only the approach from [7] is suitable for the considered scenario. Similar to the method proposed here, it automatically generates a charging policy. However, it is limited to the combination of existing policies, which cannot be expected to behave significantly differently from the single policies. The approach proposed here does not have this limitation and can generate completely new policies. As mentioned previously, there are EV charging control approaches based on reinforcement learning, which share some similarities in this respect. However, such approaches typically lack interpretability. To the best of the authors' knowledge, the design of EV charging policies through genetic programming has not been proposed in the literature before. This study adopts the problem setting from [7], which is described in more detail in the following section, and compares the proposed approach in simulation experiments to several baselines, including the method from [7].

3. Problem Setting

Nomenclature

Δt Length of each timestep
$\hat{T}^{\text{add}}(CP)$ Maximum additional daily external charging time over all users with control policy CP
$\overline{T}^{\text{add}}(CP)$ Mean additional daily external charging time over all users with control policy CP
$B^{\text{arr}}_{n,d}$ Arrival battery level of user n on day d
$B^{\text{dep}}_{n,d}$ Departure battery level of user n on day d
$B_{n,d,t}$ Battery level of user n after timestep t of day d
C_n Battery capacity of user n
D Number of considered days
$E_{n,d}$ Energy consumption of user n on day d
H Number of timesteps in the time horizon per day
N Number of EV users
P^{max}_n Maximum charging power of user n
P^{lim} Limit of total EV charging power
$P_{n,d,t}$ Charging power of user n in timestep t of day d
S^{min} Minimum SOC ensured by users through external charging
$t^{\text{arr}}_{n,d}$ Arrival timestep of user n on day d
$t^{\text{dep}}_{n,d}$ Departure timestep of user n on day d
T_n Average daily external charging time of user n
$T_n(CP)$ Average daily external charging time of user n resulting from charging control with policy CP
$T_n(UCTL)$ Average daily external charging time of user n resulting from uncontrolled charging

$T_n^{\text{add}}(CP)$ Additional daily external charging time of user n with control policy CP compared to uncontrolled charging

A set of N EV users are considered, each with an individual charging spot in a shared residential garage. The total charging power is limited by a value P^{lim} due to transformer constraints, which is significantly lower than the sum of the maximum charging powers of the EVs. To ensure that this limit is not violated, a rule-based charging control is used, which periodically (e.g., every 15 minutes) distributes the available power among connected EVs that are not yet fully charged, according to a given charging policy. Information on departure times and energy requirements is not available to the controller. The charging control results in dissatisfaction of users whose charging demands are satisfied less compared to uncontrolled charging. It is assumed that the users can charge externally at public charging stations or their workplaces, but prefer to charge at home. This is consistent with studies of EV users' behavior, which conclude that charging at home is most popular, followed by workplace and public charging [31]. Therefore, in this work, dissatisfaction of users is measured as the additional time they have to spend per day on external charging compared to uncontrolled charging. The general problem is to find a charging policy that minimizes this additional external charging time and thereby maximizes user satisfaction.

Historical data of charging behaviors and demands of the users over D days in the past is assumed to be available, allowing simulation-based evaluation of charging policies. This makes it possible to estimate users' additional daily external charging times during the historical time period under any given policy. The objective is to learn a charging policy that performs well on the historical data, in the hope that it will also perform well in the future.

In this study, the following more specific assumptions are made: for each of the D historical days, a time horizon consisting of H timesteps of a length of Δt hours each is considered[1]. Every day d, each user n initiates a charging session within the associated time horizon, characterized by arrival and departure timesteps $t_{n,d}^{\text{arr}}$ and $t_{n,d}^{\text{dep}}$, as well as an arrival battery level $B_{n,d}^{\text{arr}}$ in kWh. In each timestep t (with $t_{n,d}^{\text{arr}} \leq t \leq t_{n,d}^{\text{dep}}$) of a user n's charging session, the user's EV is charged with a certain charging power $P_{n,d,t} \leq P_n^{\max}$, where $P_n^{\max}$ denotes the maximum charging power of the user's EV in kW. After the timestep, the battery level $B_{n,d,t}$ of the user's EV is

$$B_{n,d,t} = \begin{cases} \min\{C_n, B_{n,d}^{\text{arr}} + \Delta t \cdot P_{n,d,t}\} & \text{if } t = t_{n,d}^{\text{arr}} \\ \min\{C_n, B_{n,d,t-1} + \Delta t \cdot P_{n,d,t}\} & \text{else,} \end{cases} \tag{1}$$

where C_n is the battery capacity of the EV in kWh. The battery level $B_{n,d,t_{n,d}^{\text{dep}}}$ after timestep $t_{n,d}^{\text{dep}}$ is the departure battery level $B_{n,d}^{\text{dep}}$ of user n on day d. After charging their EV on day d, user n drives a certain distance, consuming an energy amount $E_{n,d}$, before returning for charging on the next day $d+1$. It is assumed that the users want to keep their SOCs above a certain minimum SOC $S^{\min} \in [0, 1]$ and that they recharge their EVs externally if needed to avoid dropping below this level. Thus, it is assumed that the arrival battery level of a user n on day d (except the first day) is

$$B_{n,d}^{\text{arr}} = \max\{B_{n,d-1}^{\text{dep}} - E_{n,d-1}, S^{\min} \cdot C_n\}. \tag{2}$$

Assuming that users can recharge with their respective maximum charging power externally, a user n has to spend on average

$$T_n = \frac{60}{D} \sum_{d=1}^{D} \frac{\max\{0, S^{\min} \cdot C_n + E_{n,d} - B_{n,d}^{\text{dep}}\}}{P_n^{\max}} \tag{3}$$

minutes per day charging externally (note that conversion is made from hours to minutes in (3), multiplying by 60). As stated previously, a rule-based controller allocates the charging power for each connected EV at every timestep, ensuring a limit P^{lim} for the total charging power. Let $T_n(CP)$ denote the *total* daily external charging time for user n under charging policy CP. Analogously, let $T_n(UCTL)$ denote the daily external charging time under uncontrolled charging conditions—i.e., without power limits and with each EV charging at its maximum power. Consequently, the user's dissatisfaction under policy CP is defined as the *additional* daily external charging time over the uncontrolled baseline:

$$T_n^{\text{add}}(CP) = T_n(CP) - T_n(UCTL). \tag{4}$$

[1]Please note that a time horizon is not necessarily aligned to a calendar day and might overlap with two calender days. Thus, day d actually refers to the d-th time horizon.

The efficiency of a charging policy CP, in terms of the overall user satisfaction, is measured as the mean additional charging time per day $\overline{T}^{\text{add}}(CP)$ over all users:

$$\overline{T}^{\text{add}}(CP) = \frac{\sum_{n=1}^{N} T_n^{\text{add}}(CP)}{N}, \tag{5}$$

where a lower value of $\overline{T}^{\text{add}}(CP)$ indicates a higher efficiency. However, even with a high efficiency, some users may still be treated very disadvantageously. Thus, in addition to efficiency, fairness is considered, which is measured according to the commonly used min-max fairness principle, as the maximum additional charging time per day $\hat{T}^{\text{add}}(CP)$ over all users:

$$\hat{T}^{\text{add}}(CP) = \max\{T_n^{\text{add}}(CP) | n = 1, \ldots, N\}, \tag{6}$$

where, again, a lower value indicates a higher fairness. Ultimately, the problem consists in learning a charging policy which minimizes the two objective functions $f1$ and $f2$ on the historical data:

$$\min_{CP} f1 = \overline{T}^{\text{add}}(CP) \tag{7}$$

$$\min_{CP} f2 = \hat{T}^{\text{add}}(CP). \tag{8}$$

Please note that the approach proposed in this work is not limited to the previously described problem setting. For instance, it can also be applied when EVs do not have exactly one charging session per day. Moreover, it is independent of the specific metrics used to assess the quality of the charging service. For example, the (dis-)satisfaction of users could be expressed in a different way than their external charging times, or fairness could be evaluated using criteria other than min-max fairness.

4. Approach

In the previous section, the problem was formulated as an optimization problem with two objectives. However, the search space has not yet been clearly defined, as "charging policy" remains a rather abstract term. A charging policy can be regarded as a model or an algorithm that, at each timestep, determines the amount of energy to allocate to each connected EV based on information about the charging system's current state, particularly the states of the connected EVs themselves. This work proposes using genetic programming (GP) to learn a model aimed at achieving high performance with respect to the two objectives. GP, as first described by Koza [5], is a special type of evolutionary algorithm that evolves programs, expressions or formulas. It thereby produces inherently white-box models, which are generally more interpretable and explainable than the black-box representations used by methods such as neural networks [6]. Studies have shown that algorithmic transparency positively impacts trust and perceived fairness [32]. A GP model is a combination of variables, constants and operators. The sets of variables and operators that can be combined are predefined, while constants are either predefined or evolved by the GP approach. Commonly, models are represented internally as trees. For example, the expression $5 \cdot X1 + 2 \cdot X2$, where $X1$ and $X2$ are variables, can be represented as the tree shown in Figure 1. To search for a model that performs well with respect to a given

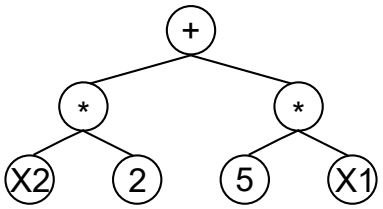


Figure 1: Illustration of a genetic programming tree that represents the expression $5 \cdot X1 + 2 \cdot X2$.

objective function, GP adheres to the typical workflow of evolutionary algorithms, as shown in Figure 2. It operates on a set, also called *population*, of a fixed number of candidate solutions, also referred to as *individuals*. Starting from

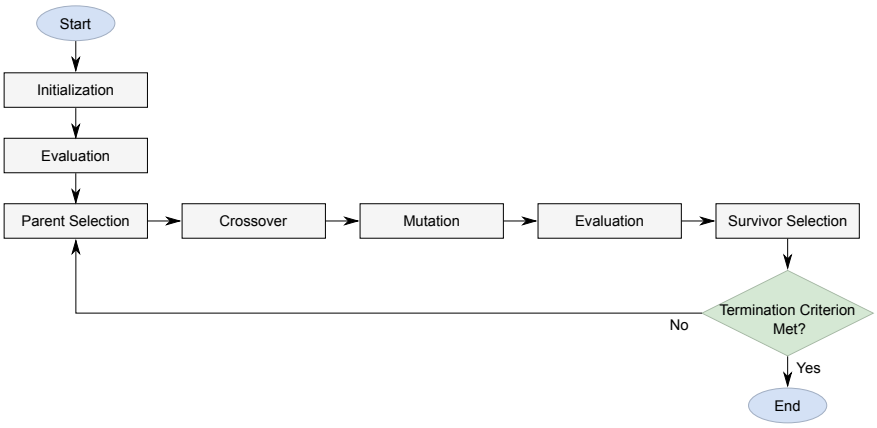


Figure 2: The workflow of genetic programming and evolutionary algorithms in general.

an initial population, the algorithm iteratively varies the population. In each iteration, also denoted as *generation*, parent individuals are selected from the population and offspring individuals are generated by applying crossover and mutation operators to the selected parent individuals. The crossover operator combines two or more parents, while the mutation operator applies random changes to an individual. After evaluating the offspring individuals with the objective function, survivor individuals are selected from the parent population and the newly generated offspring, which form the population of the next generation.

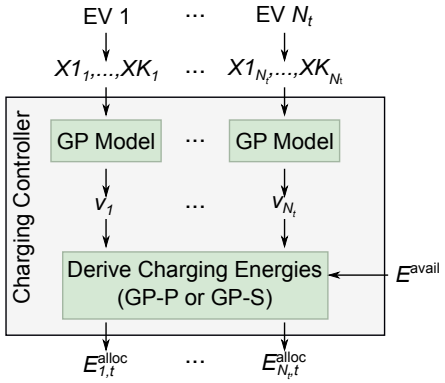


Figure 3: Illustration of the charging control in a timestep t.

In this work, a GP model is evolved for charging control, as illustrated in Figure 3. At the beginning of each timestep t, where the number N_t of connected EVs is greater than zero, the GP model is used to compute for each connected EV n a score v_n based on K numerical features $X1, \ldots, XK$ of the EV (e.g., SOC or maximum charging power). From the resulting scores $v_1, \ldots, v_{N_t}$, the energy allocations $E^{\text{alloc}}_{1,t}, \ldots, E^{\text{alloc}}_{N_t,t}$ for the connected EVs in the current timestep t are derived. Theoretically, the energy allocated to an EV could be set directly to its corresponding score. However, this would neither ensure that the amount of available energy $E^{\text{avail}} = P^{\text{lim}} \cdot \Delta t$ is not exceeded, nor that as much as possible of the available energy is allocated to the EVs. This study evaluates two different approaches for deriving allocations from the scores computed with the GP model.

The first approach, denoted GP-P, interprets the computed scores as priorities. The EV n with the highest score v_n, is allocated as much energy as possible, defined as the minimum of the available energy $E^{\text{avail}} = P^{\text{lim}} \cdot \Delta t$ and the maximum amount that the EV can technically feasibly charge in the current timestep:

$$E^{\text{alloc}}_{n,t} = \min\{E^{\text{avail}}, \min\{C_n - B_{n,t}, P^{\text{max}}_n \cdot \Delta t\}\}, \tag{9}$$

where C_n is the battery capacity, $B_{n,t}$ is the battery level at the beginning of timestep t and P^{max}_n is the maximum

charging power of EV n. Subsequently, as much of the remaining available energy as possible is allocated to the EV with the second highest score, then to the third, and so on, until either all connected EVs have received energy or all of the available energy is fully allocated. This procedure is detailed in Algorithm 1 in the appendix.

The second approach, denoted GP-S, allocates shares of the available energy to the connected EVs according to the scores computed by the GP model. Let C denote the set of connected EVs that are not fully charged. For each EV $n \in C$ the energy allocation is defined as

$$E_{n,t}^{\text{alloc}} = \min\{\frac{|v_n|}{\sum_{k \in C} |v_k|} \cdot E^{\text{avail}}, \min\{C_n - B_{n,t}, P_n^{\max} \cdot \Delta t\}\}. \tag{10}$$

Thus, the normalized absolute values of the scores computed for the connected EVs determine the shares of the available energy allocated to the EVs. If any energy remains unallocated, it is distributed in the same way to EVs that can still accept more. See Algorithm 2 in Appendix A for more details. The GP-P approach can be regarded as an "on/off" strategy. In each timestep, high-priority EVs are charged as much as possible, while low-priority EVs are not charged at all. By contrast, GP-S distributes the available energy more evenly among the connected EVs. Please note that the values of the EVs' features typically change from timestep to timestep. For example, the SOC of an EV increases if it was charged in the previous timestep. Hence, the computed scores and the resulting amounts of allocated energy typically change from timestep to timestep, even in the case that the connected EVs stay the same.

As explained, GP is used to search for models that provide good objective values when applied in the GP-P or GP-S charging control approach on historical data. However, the two considered objective functions are partially conflicting. Therefore, multi-objective GP is employed, which computes a set of *non-dominated* solutions. Given a multi-objective minimization problem with search space Ω and M objective functions $f_1, \ldots, f_M$, a solution $x_1 \in \Omega$ is said to dominate a solution $x_2 \in \Omega$ if x_1 is better than x_2 in one objective and at least as good as x_2 in the other objectives:

$$f_m(x_1) \leq f_m(x_2) \, \forall m \in \{1, \ldots, M\} \quad \text{and} \quad \exists m \in \{1, \ldots, M\} : f_m(x_1) < f_m(x_2). \tag{11}$$

A solution $x_1 \in \Omega$ is called non-dominated if there is no other solution $x_2 \in \Omega$ that dominates it. The set of all non-dominated solutions is called the *Pareto set*, and the corresponding set of objective values is called *Pareto front*. The goal of multi-objective optimization is to approximate this Pareto set with a representative set of mutually non-dominating solutions. Genetic programming, and evolutionary algorithms in general, are considered well-suited for multi-objective optimization since they inherently operate on a set of candidate solutions. At the end of the optimization process, the non-dominated solutions in the final population are returned as an approximation of the Pareto set. To guide the optimization towards a high-quality approximation, specialized operators for the selection of parent and survivor individuals have been proposed in the literature. As explained in more detail later, the selection strategy of the widely used Non-dominated Sorting Genetic Algorithm (NSGA-II) [33] is applied to multi-objective GP in the experiments described in the present work.

Figure 4 summarizes the proposed framework and approach. Using historical data on user charging behavior,

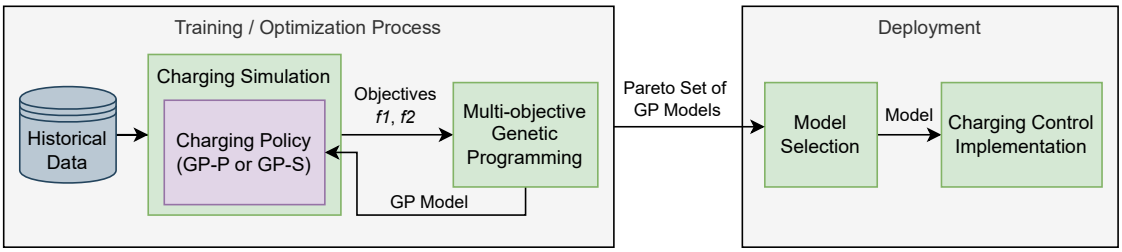


Figure 4: Overview of considered framework and proposed approach.

a simulation framework is set up, enabling EV charging over the historical period to be simulated under different controlled-charging policies. With this simulation, charging policies are evaluated with respect to two objectives, $f1$ and $f2$, which represent the mean and maximum daily additional external charging times of the users. Multi-objective GP is employed to search for models that perform well with respect to these objectives. The GP models are incorporated into the charging control policy either to prioritize EVs (GP-P) or to directly compute shares of the

available energy allocated to the EVs (GP-S). This results in a set of mutually non-dominating policies, from which one can be selected by a decision maker for implementation in practice.

5. Experiments

In simulation experiments, the two variants, GP-P and GP-S, of the proposed approach are evaluated and compared to different baselines. The setup of the simulations used in the experiments is described in the following subsection. Details on the implementation of the multi-objective GP are provided in Section 5.2, and the baseline approaches considered in the experiments are described in Section 5.3. The experimental results are presented and discussed in Section 5.4.

5.1. Simulation Setup

As described previously, it is assumed that, in practice, simulations are set up using historical data. Due to a lack of real-world historical data, we assume artificially generated user behaviors and characteristics in the simulations, which are set up as follows: the charging of $N = 100$ users over D days is simulated under the assumption that all days are working days—i.e., no weekends or public holidays are considered in the simulations. This assumption simplifies the simulation and is not expected to qualitatively affect the overall trends in the results. For each day, $H = 96$ timesteps of length Δt of 0.25 h (i.e., 15 min) are simulated, representing the period from 12 pm to 12 pm on the following day. The power limit P^{lim} is set to 20 kW. This value was chosen based on prior experiments indicating that a power limit of this magnitude often does not satisfy all users' charging requirements. For simplicity, the limit is assumed to be constant. In practice, however, the power limit may vary because a base load must be taken into account in addition to the EV charging load.

To simulate user behavior—namely arrival/departure times and energy consumption—an average profile is sampled for each user and is used to generate specific values for each simulated day. The setting of arrival and departure times is based on an analysis of real-world residential charging data by Sørensen et al. [34], which shows that on working days, the number of EV arrivals peaks in the time between 3 pm and 8 pm and the number of departures peaks between 6 am and 8 am. From this, the following procedure for setting the arrival and departure timesteps is derived: For each user n, a mean arrival time μ_n^{arr} is sampled from a truncated Gaussian distribution with a lower bound of 5 (i.e., 1 pm), a mean of 25 (i.e., 6 pm), an upper bound of 45 (i.e., 11 pm), and a standard deviation of 8 (i.e., 2 hours). Daily arrival timesteps of a user n are then sampled from a truncated Gaussian distribution with lower bound 1 (i.e., 12 pm), a mean of μ_n^{arr}, an upper bound of 96 (i.e., 11:45 am the following day), and a standard deviation of 2 (i.e., 30 minutes) and are rounded to the nearest integers. Analogously, for each user n, a mean departure time μ_n^{dep} is sampled from a truncated Gaussian distribution with a lower bound of 65 (i.e., 4 am, following day), a mean of 77 (i.e., 7 am, following day), an upper bound of 89 (i.e., 10 am, following day) and a standard deviation of 4 (i.e., 1 hour). Daily departure timesteps are then sampled from a truncated Gaussian distribution with the corresponding arrival timesteps as the lower bounds, a mean of μ_n^{dep}, an upper bound of 96 and a standard deviation of 2 (i.e., 30 minutes).

The daily energy consumptions $E_{n,d}$ of the users are set based on daily driving ranges $R_{n,d}$:

$$E_{n,d} = c \cdot R_{n,d}, \tag{12}$$

where c denotes the energy consumption per km, which is assumed to be constant and identical for all EVs. Based on an analysis of real-world data, Weiss et al. [35] report energy consumption values between 0.10 and 0.33 kWh/km for electric passenger cars. Accordingly, a simplified assumption of a constant consumption of $c = 0.2$ kWh/km is used in the simulations. To compute daily driving ranges, a mean driving range μ_n^R per user n is first sampled from a truncated Gaussian distribution with a lower bound of 15 (km), a mean of 30 (km), an upper bound of 100 (km) and a standard deviation of 10 (km). Daily driving ranges of user n are then sampled uniformly distributed from the interval $[0.9 \cdot \mu_n^R, 1.1 \cdot \mu_n^R]$.

EV models are randomly assigned to the users, considering the five EV models with the highest market share in Germany in 2020 according to [36]. The battery capacities C and maximum charging powers P^{max} of these models are obtained from EV database [37]. It is assumed that the users have access only to AC charging stations for home and external charging. The probabilities P^{model} of the EV models to be assigned to a user are set proportional to the sales

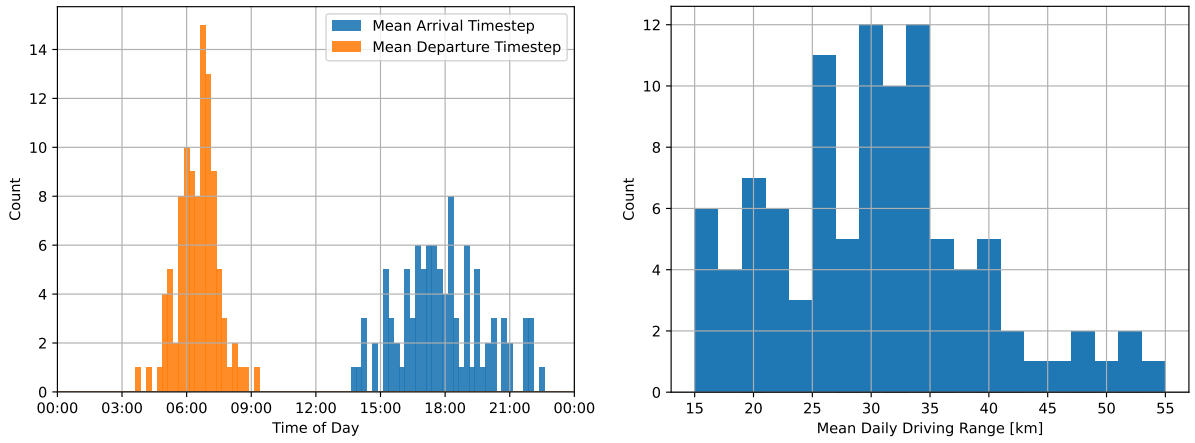
numbers reported by [36]. The considered EV models, together with their technical specifications and assignment probabilities, are listed in Table 2.

Table 2: EV models considered in the simulations.

Model	$P^{\max}$ **(AC)**	C	P^{model}
Renault Zoe	22 kW	41 kWh	0.33
Volkswagen e-Golf	7.2 kW	40 kWh	0.19
Tesla Model 3	11 kW	49 kWh	0.17
Volkswagen ID.2	11 kW	58 kWh	0.16
Hyundai Kona	11 kW	64 kWh	0.15

Users are assumed to ensure a minimum arrival SOC $S^{\min}$ of 0.1 through external charging if necessary. The initial SOCs at the beginning of the simulation are set to 0.5. In the simulation, EV charging under a given charging policy is simulated, with energy distributed to connected EVs at each timestep according to the policy. This, in combination with the daily energy requirements of the users, allows to compute for each day the users' arrival and departure SOCs and their daily additional external charging times as outlined in Section 3.

Unless otherwise stated, the experiments are carried out on a simulation of 1100 days initialized with a fixed random number generator seed. In the following, this simulation is referred to as the "default" simulation. The first 100 simulated days are used for training, representing the historical period on which charging policies are learned. The remaining 1000 days are used to test whether the resulting policies generalize to the future. The training period is set notably shorter than the test period since, in practice, data must be collected over the training period, making it desirable to keep it short. Figure 5 shows histograms of the users' mean arrival and departure times and their mean daily driving ranges in the default simulation.



(a) (b)

Figure 5: Histograms of the 100 users' mean arrival and departure times (a) and mean daily driving ranges (b) in the default simulation.

5.2. Multi-objective Genetic Programming

At each timestep, the GP model is used to assign each connected EV either a priority (in the GP-P approach) or an energy share (in the GP-S approach). In preliminary experiments, the following six features of an EV n were identified as promising input variables for such a GP model:

- $X1$: The number of timesteps since arrival of the EV computed as $t - t_n^{\text{arr}} + 1$, where t_n^{arr} is the arrival timestep of the EV n and t is the current timestep.

- $X2$: The battery level of the EV in kWh.
- $X3$: The SOC of the EV, a scalar value in the interval [0,1].
- $X4$: The average amount of charged energy per timestep since arrival computed as $(B_n - B_n^{\text{arr}} + 0.0001)/(t - t_n^{\text{arr}} + 1)$, where B_n^{arr} is the arrival battery level and B_n is the current battery level of the EV. A small constant is added to the numerator to avoid zero values, which could otherwise lead to divisions by zero in the GP model.
- $X5$: The maximum charging power $P_n^{\max}$ of the EV in kW.
- $X6$: The arrival SOC of the EV, a scalar value in the interval [0,1].

It is recalled that energy requirements and departure times of EVs are assumed not to be available as inputs for the charging control and therefore cannot be used as input variables for the GP model.

For the implementation of the multi-objective GP, the Python libraries DEAP [38] and pymoo [39] were used. DEAP was employed for the internal representation of GP models in the form of trees, for the sampling of the initial population, for crossover and mutation, and for the conversion of trees to expressions during the evaluation. The general workflow and the selection process were carried out by the NSGA-II algorithm implemented in pymoo. The GP trees were composed of constants, variables as listed above, and arithmetic operators from the set $\{+, -, \times, /\}$. Unless otherwise stated, the maximum depth and maximum size of GP trees were set to 10 and 50, respectively. The population was initialized with randomly generated trees of a maximum depth of five. The population size was set to 100, and in each optimization run, 200 generations were executed. In each generation, 100 offspring were produced. Parent individuals were selected by tournament selection, in which five individuals were randomly sampled from the population, and the best among them was chosen.

For the comparison of individuals, NSGA-II's non-dominated sorting strategy was applied. This strategy partitions the population into several levels: the first level contains the non-dominated individuals of the total population, the second level contains the non-dominated individuals from the remaining ones, and so forth. Individuals at a lower level are considered superior to those at a higher level. Two individuals at the same level are compared by their *crowding distance* [33], a measure of their distance to nearest neighbors in the objective space. Individuals with higher crowding distance are preferred over those with lower values.

With a probability of 0.95, two offspring were generated from two selected parents through crossover; otherwise, the parents were directly copied. In the crossover, randomly chosen subtrees are exchanged between the parents. With a probability of 0.2, mutation was applied to an offspring, randomly altering a node in the tree. If mutation or crossover produced a tree that exceeded the maximum depth or size, a random parent was used instead. After evaluation of the offspring, the 100 best individuals from the combined set of parents and offspring were selected to form the population of the next generation, again using NSGA-II's non-dominated sorting strategy for comparison.

Table 3 summarizes the setting of the most important parameters.

Table 3: Parameter Setting used for the Multi-objective GP.

Parameter	**Value**
Population Size	100
Number of Generations	200
Crossover Rate	0.95
Mutation Rate	0.20
Allowed Operators	$\{+, -, \times, /\}$

5.3. Baseline Approaches

As baselines for the evaluation, the following seven manually designed EV charging policies, as well as the combined control approach described in [7], are considered:

Equal Distribution (ED): At each timestep, the available energy is distributed as equally as possible among the connected EVs. More precisely, each EV n in the set C of connected EVs is allocated an amount of energy

$$E_n^{\text{alloc}} = \min\{\frac{E^{\text{avail}}}{|C|}, \min\{C_n - B_n, P_n^{\max} \cdot \Delta t\}\}, \tag{13}$$

where E^{avail} is the available energy, C_n is the battery capacity, B_n is the battery level, and $P_n^{\max}$ is the maximum charging power of EV n. If any energy remains, it is reallocated analogously to EVs that can still accept additional energy. This reallocation is repeated until either the available energy is exhausted or all EVs have reached their maximum charging capacity.

First-come-first-served (FCFS): EVs with earlier arrival times are prioritized over those with later arrival times as follows: As much as possible of the available energy is allocated to the EV that arrived first. From the remaining energy, as much as possible is allocated to the EV that arrived second, and so on.

Less Energy First (LEF): EVs with lower battery levels are prioritized over those with higher values, with allocation carried out analogously to the FCFS policy.

Lower SOC First (LSF): EVs with lower SOCs are prioritized over those with higher SOCs, with allocation carried out analogously to the FCFS policy.

Less Charged First (LCF): EVs with lower amounts of energy charged per timestep since arrival are prioritized over those with higher values, with allocation carried out analogously to the FCFS policy.

Lower Power First (LPF): EVs with lower maximum charging power are prioritized over those with higher maximum charging power, with allocation carried out analogously to the FCFS policy.

Lower Arrival SOC First (LAF): EVs with lower arrival SOC are prioritized over those with higher arrival SOC, with allocation carried out analogously to the FCFS policy.

Combined Control (CC): The combination of the seven previously described policies is applied as proposed in [7]: First a fraction w_1 of the available energy is distributed to the connected EVs using the ED policy, then another fraction w_2 of the available energy is distributed using the FCFS policy, then a fraction w_3 is distributed using the LEF policy, and so on, with the fractions constrained such that $\sum_{k=1}^{7} w_k = 1$, and $w_k \in [0, 1]$ for $k = 1, \ldots, 7$. The fractions are optimized through multi-objective optimization with respect to the two objective functions $f1$ and $f2$, in a manner analogous to the multi-objective optimization of models with genetic programming. As a result, a complete Pareto front approximation of charging policies (encoded as settings of fraction values) is obtained. For the optimizations, the real-valued NSGA-II implementation from the C++ library pagmo [40] is used. Similar to the multi-objective GP, the population size is set to 100, the number of generations to 200, the crossover probability to 0.95, and the mutation probability to 0.20. Distribution indexes $\eta_c = 10$ and $\eta_m = 20$ are used for the real-valued crossover and mutation, respectively.

It should be noted that the proposed GP-based approach can be regarded as a generalization of manually designed charging policies such as those described above. For instance, the GP-S variant with a GP model that simply outputs a constant value would be equivalent to the ED charging policy, while the other four policies can be reproduced by the GP-P variant.

5.4. Experimental Results

In the following, the results of the simulation experiments are discussed. More precisely, the results on the 100 days training data and the 1000 days test data are presented in Sections 5.4.1 and 5.4.2, respectively. In Section 5.4.3, results on data with unseen user characteristics are shown. Section 5.4.4 provides an analysis of the models evolved by the genetic programming approach. In Section 5.4.5, it is investigated how a reduction of the complexity of the models affects their performance. Finally, Section 5.4.6 discusses the impact of the number of training days on the performance on the test data.

5.4.1. Results on Training Data

The different approaches were first evaluated on the 100 days training data from the default simulation. Table 4 reports the mean and maximum daily additional external charging times obtained with the seven manually designed charging policies. Of the manually designed policies, the LPF and LAF policies achieved the best results. Giving priority to EVs with low maximum charging power, as done by the LPF policy, is reasonable, since they require

Table 4: Results of the manually designed charging policies on the 100 days training data. Best results are highlighted in bold.

Policy	$f1$ ($\overline{T}^{\text{add}}$) **[min]**	$f2$ ($\hat{T}^{\text{add}}$) **[min]**
ED	10.240	49.818
FCFS	10.481	58.954
LCF	11.464	54.539
LEF	9.815	44.641
LSF	9.986	53.120
LPF	**6.144**	24.567
LAF	6.440	**23.071**

more time to recharge a certain amount of missing energy externally than EVs with higher maximum charging power. EVs that have to recharge externally on one day arrive back with a low SOC and are prioritized by the LAF policy on the following day. This yields the fairest charging control out of the manually designed controllers. Since both the proposed genetic programming-based approach and the CC baseline approach employ stochastic optimization, 21 runs with different seeds for the optimization algorithm were conducted. One run of the GP optimization took approximately three minutes on a host equipped with an Intel Core i5-13400@4.6 GHz 16-core CPU. For each run, the *hypervolume* (HV) of the resulting Pareto front approximation was computed. The hypervolume [41] is a popular metric used in multi-objective optimization to assess the quality of a Pareto front approximation. It quantifies the portion of the objective space dominated by the Pareto front approximation with respect to a reference point, with higher values indicating better performance.

The reference point for hypervolume computations was chosen as (11.464, 58.954), corresponding to the worst objective values obtained by any of the manually designed policies. Figure 6 visualizes the results of the manually designed policies together with the approximated Pareto fronts obtained with the optimization-based approaches CC, GP-S, and GP-P on the training data. Figure 6a presents the fronts combined over the 21 runs, i.e., the mutually non-

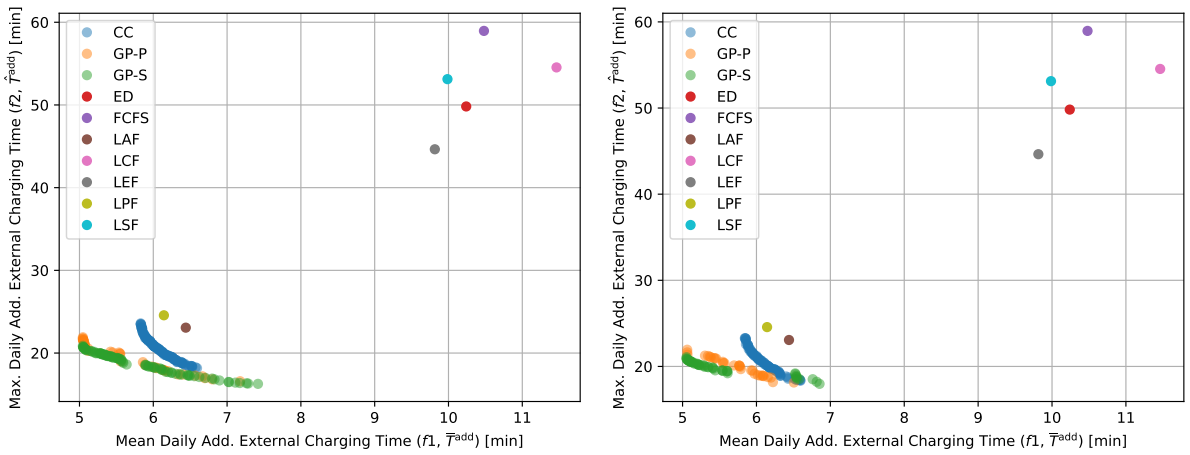


Figure 6: Results obtained under the five manually designed charging policies and the optimization-based approaches CC, GP-P and GP-S on the 100 days training data. For the optimization-based approaches, the Pareto front approximation combined over all 21 runs is shown in (a) and the Pareto front approximation resulting from the runs that yielded the median hypervolume is shown in (b).

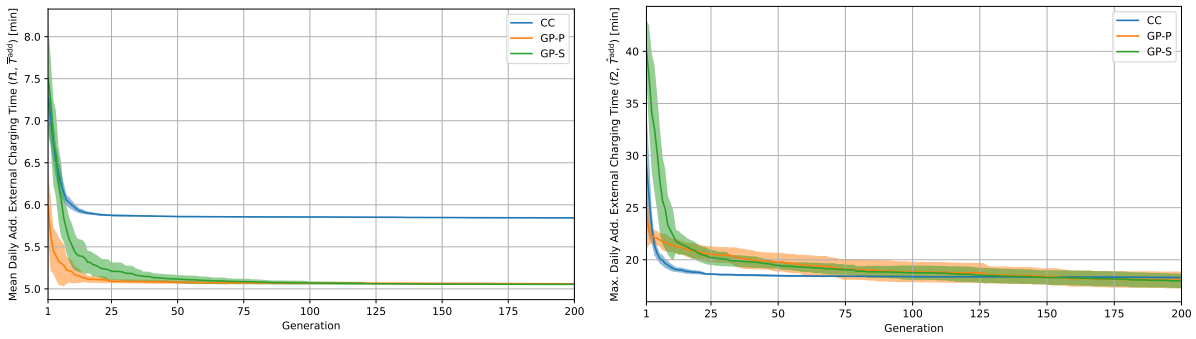
dominated solutions in the union of the Pareto front approximations computed in the 21 runs. Figure 6b depicts the fronts that yielded the median hypervolume over the 21 runs. In practice, a single solution—i.e., a charging control policy—must ultimately be selected for implementation from the approximated Pareto front. Such a solution may be selected manually by a decision maker, automatically (for example, by selecting a knee point of the front [42]) or semi-automatically by considering decision maker preferences [43]. In this study, the decision-making process itself

is not examined further. Instead, attention is directed to the extreme points of the computed Pareto fronts, i.e., the solutions that are optimal with respect to one of the two objectives, for a more detailed analysis. Table 5 shows the mean and standard deviation of the hypervolumes of the fronts obtained with the optimization-based approaches over the 21 runs per approach. Furthermore, Table 5 provides information on the extreme points of the computed Pareto

Table 5: Results of 21 runs using the combined control approach CC and the genetic programming-based approaches GP-P and GP-S on the 100 days training data. For each approach and each objective functions $f1$ and $f2$, the table shows the best objective values observed in any run, the mean μ, and the standard deviation σ of these best values across all runs. Furthermore, the mean and standard deviation of the hypervolumes of the Pareto fronts across runs are reported. Best results are highlighted in bold.

	$f1$ ($\overline{T}^{\text{add}}$) [min]			$f2$ ($\hat{T}^{\text{add}}$) [min]			Hypervolume	
Approach	**Best**	μ	σ	**Best**	μ	σ	μ	σ
CC	(5.829, 23.591)	5.845	0.007	(6.592, 18.176)	18.298	0.070	227.172	0.449
GP-P	(**5.046**, 21.909)	5.061	0.009	(7.176, 16.593)	18.066	0.806	259.040	4.104
GP-S	(**5.046**, 20.877)	**5.053**	0.005	(7.419, **16.276**)	**17.978**	0.701	**260.709**	3.654

front approximations. For each objective, $f1$ and $f2$, the best value over the 21 runs is shown in the column "Best". For example, the best $f1$ value found by the CC approach was 5.829, with a corresponding $f2$ value of 23.591. The Table 5 also provides the average μ and standard deviation σ over the best objective values across single runs. From Figure 6 and Tables 4 and 5, it is evident that while the CC baseline approach finds charging policies that are more efficient and fair than the manually designed ones, it is outperformed by the GP-based approaches. The GP-based methods reduce the mean additional external charging time by about 13 % compared to the CC approach, and about 18 % compared to the manually designed policies. The maximum additional charging time is roughly 12 % lower than with CC and 29 % lower than with any of the manually designed policies. When selecting a single solution from the front, a trade-off between objectives is unavoidable. However, the "Best" columns in Table 5 show that solutions found by the GP-based approaches, which are best in one objective, still perform reasonably well in the other objective. On average, GP-S performs slightly better than GP-P, but the differences are marginal, indicating that both approaches are practically equivalent on the training data. The small standard deviations in hypervolume and objective values (Table 5), along with the close similarity of the combined and median fronts (Figure 6a and Figure 6b), indicate the robustness of all three optimization-based approaches. This is further supported by Figure 7, which illustrates the progress of the best objective value over 200 generations. The curves show the mean and standard deviation across

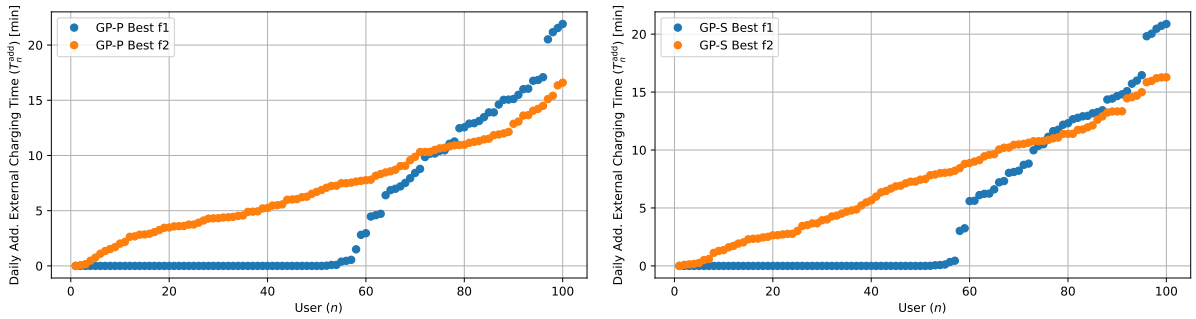


(a) (b)

Figure 7: Mean and standard deviation of the best values found for objective functions $f1$ (a) $f2$ (b) in the 21 runs in each generation for the CC, GP-P and GP-S approaches.

the 21 runs. Subfigures 7a and 7b detail the results for the first and second objective, respectively. For the first objective, the standard deviations are very small, indicating that all approaches converged within 200 generations. Although Figure 7b suggests that both GP-based approaches might continue improving the second objective beyond 200 generations, preliminary experiments showed no significant gains with longer runs. Interestingly, at the beginning

of the optimization process, GP-S performs notably worse than GP-P, but ultimately yields similar or even slightly better final results. Figure 8 shows the impact of the charging policies corresponding to the extreme points of the combined fronts computed with the two GP variants on the single users. It illustrates each user's daily additional external charging time, sorted in non-decreasing order, resulting from the policies that achieved the lowest mean and lowest maximum additional charging times (see the "Best" column in Table 5). Figure 8a shows them for the policies



(a) (b)

Figure 8: Daily additional external charging times of the individual users on the training data resulting from the extreme points of the Pareto front approximation combined over the 21 trials with GP-P (a) and GP-S (b) sorted in non-decreasing order.

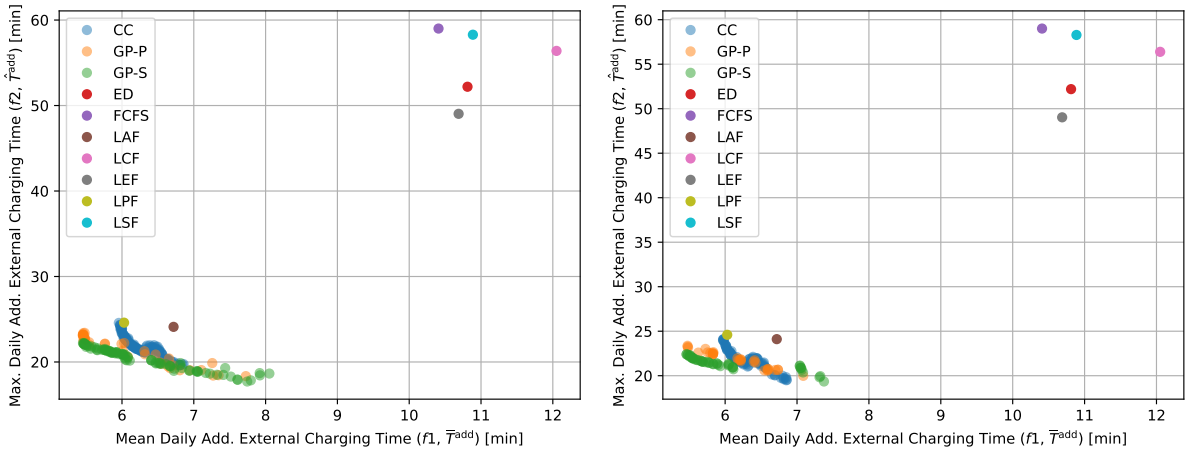
computed with GP-P, and Figure 8b for those computed with GP-S. Although the two variants perform similarly overall, their respective extreme-point solutions differ markedly. Under the policy that minimizes the mean daily additional external charging time (i.e., best on the first objective), more than half of the users experience no additional external charging. Conversely, under the policy that minimizes the maximum daily additional external charging time (the second objective), the peak charging burden across users is reduced, yet almost all users still require some external charging. Thus, while the objective values of the extreme points shown in the "Best" column of Table 5 may not differ greatly, the user-level impact remains non-negligible.

5.4.2. Results on Test Data

The results discussed so far were all on the training data, consisting of the first 100 days of the default simulation. To assess generalization to unseen conditions, both the proposed approach and the baselines were also evaluated on the test data comprising the remaining 1000 days of the default simulation. Figure 9 presents the results analogous to Figure 6, and Table 6 lists the objective values achieved on the test data by the manually designed policies and by the policies corresponding to the extreme points of the Pareto fronts, combined over the 21 runs. For example, CC_f1 denotes the policy from the CC approach that achieved the best (lowest) value on the first objective during training, while GP-P_f2 denotes the policy from the GP-P approach that achieved the best value on the second objective during training. Since these policies are selected based on their training-set performance, they do not necessarily correspond to the best results on the test data. Analogously, the combined and median fronts shown in Figure 9 are determined on the training data, and thus the contained solutions are not necessarily mutually non-dominating on the test data. As can be observed, the GP-based policies yield higher objective values on the test set than on the training set, but this also holds for most baseline approaches. Consistent with the results on the training data, the GP-S and GP-P variants exhibit no significant performance difference. Both variants yield relative and absolute improvements over the baseline approaches on the test data that are comparable to those on the training data. Thus, it can be concluded that the proposed GP approach generalizes well to the test data. Indeed, the results on the training data (Figure 6) are very similar to those on the test data (Figure 9). Although the test set comprises different days than the training set, the user characteristics (average arrival and departure times, mean daily driving ranges, and EV models) remain the same. The results suggest that day-to-day variations in user behavior average out overall.

5.4.3. Results on Completely Unseen Data

The previous section showed that the policies learned on historical data of a charging site generalize well to the future of the same site with the same users. To test whether the policies obtained by the GP-based approach generalize



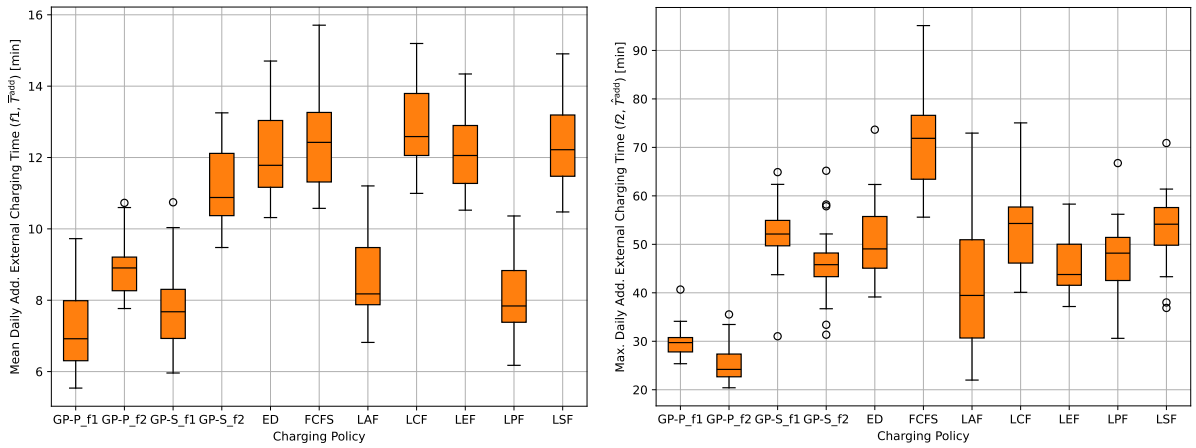
(a) (b)

Figure 9: Results on the 1000 days test data analogous to Fig. 6. (a) shows the results with the Pareto front approximations combined over all 21 runs and (b) shows the Pareto front approximations which yielded the median hypervolume in the 21 runs on the training data.

Table 6: Results of the manually designed charging policies and the extreme points of the combined fronts computed with the optimization-based approaches on the 1000 days test data. Best results are highlighted in bold.

Policy	$f1$ ($\overline{T}^{\text{add}}$) **[min]**	$f2$ ($\hat{T}^{\text{add}}$) **[min]**
ED	10.812	52.198
FCFS	10.408	59.005
LCF	12.051	56.392
LEF	10.688	49.031
LSF	10.887	58.285
LPF	6.029	24.599
LAF	6.718	24.109
CC_f1	5.956	24.571
CC_f2	6.858	19.752
GP-P_f1	5.460	23.363
GP-P_f2	7.725	**18.345**
GP-S_f1	**5.457**	22.217
GP-S_f2	8.053	18.647

even to other sites with user characteristics not seen during training, they were evaluated in 20 simulations other than the default one. In these simulations, the users have different EV models than in the default simulation, and their mean arrival and departure times, as well as their mean daily driving ranges, are no longer as depicted in Figure 5. Each of the 20 new simulations covers 1000 days. Figure 10 presents the distributions of objective values achieved by the manually designed policies and by the extreme-point policies from the combined Pareto fronts of the GP-based approaches, across the 20 simulations. Figure 10a shows these distributions for the first objective, and Figure 10b for the second. Again, note that the extreme-point policies were trained and selected on the default simulation's training data. As shown, the policies computed with the GP-P variant continue to perform well under unseen user characteristics and clearly outperform the manually designed charging policies. However, the policies computed with the GP-S variant perform notably worse and show no clear advantage over the manually designed policies. It can therefore be concluded that the GP-P variant exhibits stronger generalization behavior than the GP-S variant.



(a) (b)

Figure 10: Distribution of mean (a) and maximum (b) additional external charging times over 20 simulations of 1000 days with different seeds resulting from manually designed charging policies and the extreme points of the combined Pareto fronts computed with GP-P and GP-S approaches.

5.4.4. Analysis of Models

So far, results have been analyzed only in the objective space. To gain further insight, the models evolved by the GP variants were examined. The Pareto sets from all 21 runs were merged, the contained models or expressions were simplified with the Python library SymPy [44], and duplicates were removed. This resulted in 610 unique simplified models for the GP-P variant and 1181 for the GP-S variant. Table 7 summarizes the statistics of these models. For each of the six input variables (features), the table shows how many models include it, along with the average number of unique variables per model, the average total number of variables, and the average number of operators per model (the latter computed with SymPy). The variable $X5$ (maximum charging power) appears in every model. For the GP-P

Table 7: Statistics of unique models in the Pareto set approximations computed in the 21 runs with GP-P and GP-S.

	GP-P	**GP-S**
Number of Models	610	1181
Number of Models Containing $X1$	449 (73.71 %)	1032 (87.38 %)
Number of Models Containing $X2$	553 (90.66 %)	1180 (99.92 %)
Number of Models Containing $X3$	526 (86.23 %)	1115 (94.41 %)
Number of Models Containing $X4$	475 (77.87 %)	814 (68.92 %)
Number of Models Containing $X5$	610 (100.00 %)	1181 (100.00 %)
Number of Models Containing $X6$	606 (99.34 %)	1154 (97.71 %)
Average Number of Unique Variables per Model	5.28	5.48
Average Total Number of Variables per Model	13.36	14.57
Average Number of Operators per Model	16.17	18.72

variant, the second most frequent variable is $X6$ (arrival SOC), whereas for GP-S it is $X2$ (battery level). Typically, all six features occur in a model, and the operator counts are relatively high. On average, GP-S models are more complex than those evolved by the GP-P variant. This may explain why GP-S performs worse early in the optimization process: GP-S may require more complex models to achieve good performance, which demands more generations to evolve. The (simplified) expressions corresponding to the extreme points of the combined Pareto fronts computed by GP-P

and GP-S are:

$$\text{GP-P_f1} = -0.3238 \cdot X2 \cdot X3 + X3 \cdot (-X2 \cdot X6 - \frac{X3}{X5}) \cdot (X2 + X3 + X4 - X6) - X5 \tag{14}$$

$$\text{GP-P_f2} = \frac{X2}{(-X2 + X3) \cdot (2 \cdot X3 - X6)} - 0.4434 \cdot X3 + X6 \cdot (X3 + X5) \cdot (-X2 \cdot X4 - X2 + \frac{X3^2}{X1}) - X6 + \frac{3.5426}{X6} \tag{15}$$

$$\text{GP-S_f1} = \frac{1}{X1 \cdot X2^2 \cdot X3 \cdot X5^2 \cdot X6^3 \cdot 0.2193 \cdot (-X2 + X5) \cdot (2 \cdot X2 - X5) \cdot (X5^2 - 0.8232) \cdot (-X2 + X3 + \frac{X3}{X6} + X5)} \tag{16}$$

$$\text{GP-S_f2} = \frac{3.1899}{X3 \cdot X6 \cdot (-X2 + X5) \cdot (X6 - 0.0984) \cdot (X2 \cdot X6 + X4)}. \tag{17}$$

5.4.5. Reduction of Model Size Limit

As outlined earlier, GP is considered well-suited for explainable artificial intelligence since it produces white-box models. However, the models (14)–(17) can hardly be considered explainable or interpretable, as they are far too complex. One way to improve interpretability is to restrict GP to simpler models by imposing a lower maximum tree size during optimization. To evaluate whether such restrictions indeed result in simpler models without sacrificing performance, another 21 optimization runs were conducted on the default simulation's training data using the GP-P and GP-S variants with model size limits of 10 and 5 instead of 50 as in the experiments before. These reduced-model-size variants are denoted GP-P-MS10, GP-P-MS5, GP-S-MS10, and GP-S-MS5, respectively. Figure 11 shows the resulting combined Pareto fronts, and the corresponding mean and median hypervolume values and details of the extreme points of the fronts are provided in Table 8 (analogous to Table 5). For comparison, the table also

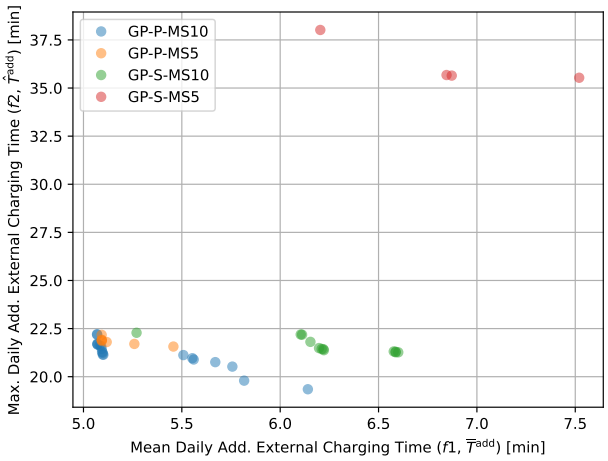


Figure 11: Pareto front approximations combined over 21 runs of the GP-P and GP-S approach with a maximum model size of 10 and 5.

includes results for CC and for GP-P and GP-S using the default model size limit of 50. As shown, the GP-P variant already outperforms GP-S with a model size limit of 10, and the performance gap widens further with a limit of 5. This is consistent with the earlier observation that GP-S requires more complex models than GP-P to perform well. In particular, for the second objective—the maximum daily additional external charging time—GP-S ceases to be competitive once the model size limit is reduced. Although the performance of the GP-P variant also deteriorates with a lower model size limit, the reduction relative to the default limit of 50 may be acceptable in exchange for higher interpretability. Even with a limit of 5, GP-P continues to be better or at least competitive with the CC approach. The observation that fairness (f_2) deteriorates more significantly under stricter model size constraints suggests that achieving high fairness may inherently require more complex models than achieving high efficiency. A deeper analysis of the models evolved with a limit of 50 (cf. Table 7) shows that their complexity (number of operators and variables)

Table 8: Results of 21 runs with CC and with the GP-P and GP-S approaches with different settings of the maximum model size on the 100 days training data analogous Table 5.

	$f1$ ($\overline{T}^{\text{add}}$) [min]			$f2$ ($\hat{T}^{\text{add}}$) [min]			**Hypervolume**	
Approach	**Best**	μ	σ	**Best**	μ	σ	μ	σ
CC	(5.829, 23.591)	5.845	0.007	(6.592, 18.176)	18.298	0.070	227.172	0.449
GP-P	(**5.046**, 21.909)	5.061	0.009	(7.176, 16.593)	18.066	0.806	259.040	4.104
GP-P-MS10	(5.069, 22.208)	5.077	0.005	(6.141, 19.348)	20.246	0.597	245.973	3.157
GP-P-MS5	(5.092, 22.174)	5.093	0.002	(5.457, 21.566)	21.787	0.126	236.747	0.771
GP-S	(**5.046**, 20.877)	**5.053**	0.005	(7.419, **16.276**)	**17.978**	0.701	**260.709**	3.654
GP-S-MS10	(5.270, 22.283)	5.410	0.159	(6.600, 21.268)	21.595	0.791	224.773	9.848
GP-S-MS5	(6.204, 38.017)	6.210	0.005	(7.520, 35.534)	35.547	0.055	121.246	0.495

is indeed positively correlated with objective $f1$ and negatively correlated with objective $f2$, meaning that the models with high fairness tend to be more complex than the models with high efficiency.

Table 9 summarizes statistics of the unique and simplified models in the Pareto front approximations computed with the GP variants with lower model size limits. Compared to the limit of 50 (Table 7), the number of unique models

Table 9: Statistics of unique models in the Pareto set approximations computed in the 21 runs with GP-P and GP-S with model size limits of 10 and 5.

	GP-P-MS10	**GP-P-MS5**	**GP-S-MS10**	**GP-S-MS5**
Number of Models	212	32	85	53
Number of Models Containing $X1$	22 (10.38 %)	1 (3.13 %)	68 (80.00 %)	3 (5.66 %)
Number of Models Containing $X2$	169 (79.72 %)	27 (84.38 %)	29 (34.12 %)	2 (3.77 %)
Number of Models Containing $X3$	65 (30.66 %)	6 (18.75 %)	62 (72.94 %)	30 (56.60 %)
Number of Models Containing $X4$	53 (25.00 %)	2 (6.25 %)	79 (92.94 %)	3 (5.66 %)
Number of Models Containing $X5$	212 (100.00 %)	32 (100.00 %)	85 (100.00 %)	53 (100.00 %)
Number of Models Containing $X6$	182 (85.85 %)	2 (6.25 %)	16 (18.82 %)	1 (1.89 %)
Average Number of Unique Variables per Model	3.32	2.19	3.99	1.74
Average Total Number of Variables per Model	3.86	2.34	4.20	1.75
Average Number of Operators per Model	4.27	3.00	4.29	2.51

decreases, and both the number of operators and variables in the models are also reduced. As before, the variable $X5$ (maximum charging power) appears in all models, indicating that it is by far the most important feature. This is not surprising, given the fact that the Lower Power First (LPF) controller is one of the best-performing manually designed controllers. Interestingly, with a limit of 5, the variable $X6$ (arrival SOC) appears only in a few models, although the Lower Arrival SOC First (LAF) controller is also one of the best manually designed controllers.

The models corresponding to the extreme points of the combined Pareto front computed with GP-S-MS5 are:

$$\text{GP-P-MS5_f1} = -X2 + X4 - X5 \tag{18}$$

$$\text{GP-P-MS5_f2} = X3 \cdot (0.3463 - X5). \tag{19}$$

These models can reasonably be considered interpretable. They are not much more complex than the simple, manually designed policies, and thus, the operator and end users should be able to understand how they work. At the same time, they provide notably higher efficiency and fairness than the manually designed charging policies. Recall that in the GP-P variant, the models are used to compute charging priorities for connected EVs. Figure 12 visualizes the priorities for different battery levels, different amounts of charged energy per timestep since arrival, and different maximum charging powers, according to model GP-P-MS5_f1. As one might expect, lower battery levels and lower maximum charging powers result in higher priorities. Surprisingly, priority also increases with the amount of energy charged per timestep since arrival. Figure 13 visualizes the priorities as functions of the maximum charging power and the SOC, according to model GP-P-MS5_f2. Again, as expected, lower SOC and maximum charging power results in

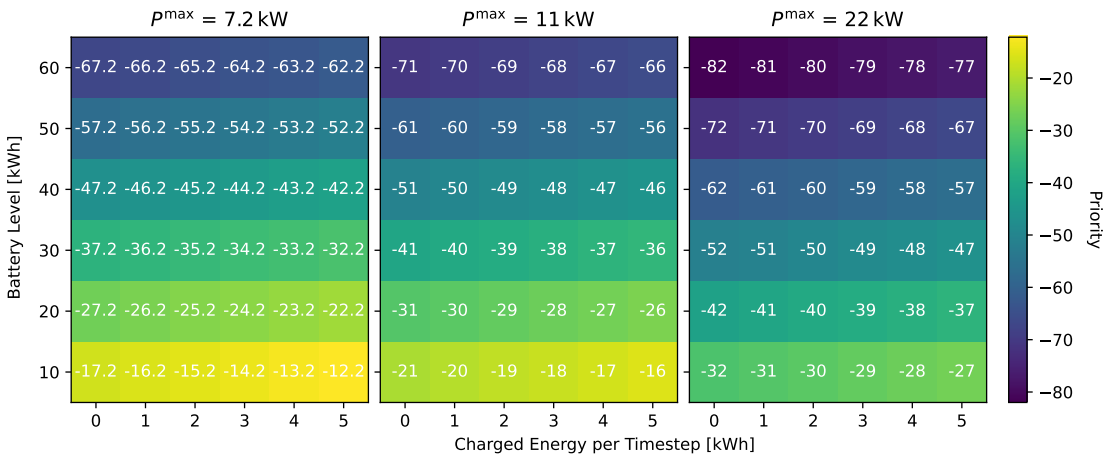


Figure 12: Charging priorities according model GP-P-MS5_f1 (Eq. (18)) for different maximum charging powers, battery levels and charged amount of energy since timestep of arrival.

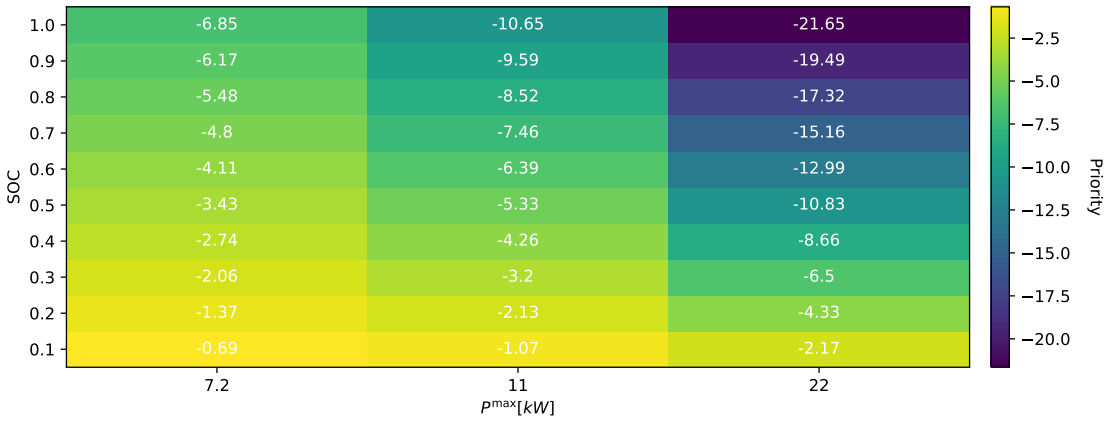


Figure 13: Charging priorities according model GP-P-MS5_f2 (Eq. (19)) for different maximum charging powers and SOCs.

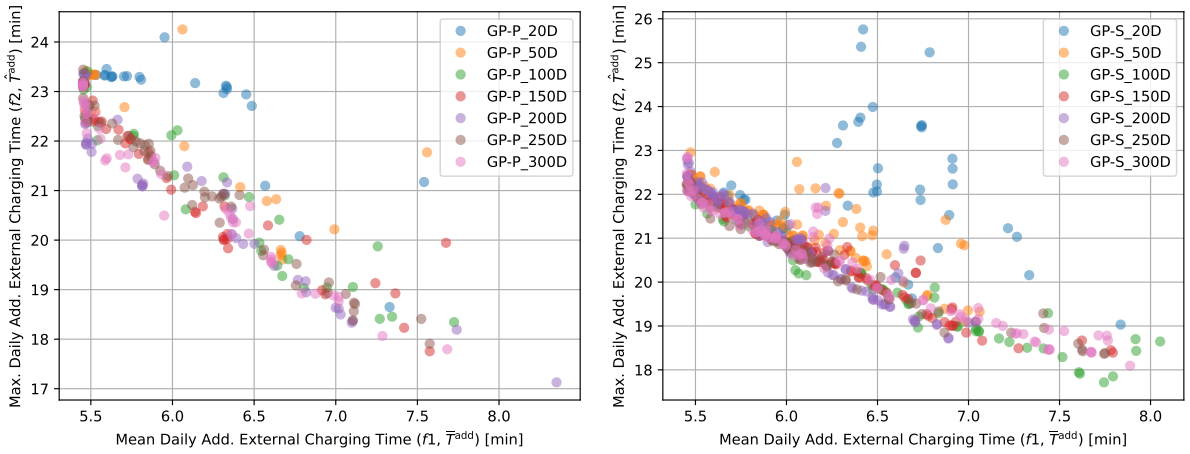
higher priorities.

5.4.6. Impact of the Number of Training Days

In the experiments so far, 100 days of training data were used for the optimizations. A lower number would imply that less historical data needs to be collected, whereas a higher number may lead to better generalization on the test data. To investigate the impact of the size of the training data on the results, the optimizations with the two GP variants were repeated with numbers of training days ranging from 20 to 300. As before, 21 runs were conducted for each GP variant and each training-day setting, and the default model size limit of 50 was applied. The resulting hypervolume values of the combined fronts (across the 21 runs) on the test data are reported in Table 10. The reference point for hypervolume calculation was defined by the worst objective values obtained by any of the manually designed policies on the test data. The corresponding Pareto front approximations are depicted in Figure 14. Compared to 100 training days, the results of the GP-P variant improve slightly when more than 100 days are provided. This trend is not observed for the GP-S variant. With 50 training days, both variants perform worse than with 100 days, but still maintain competitive performance. With only 20 training days, both GP variants no longer yield a benefit over the baseline approaches (cf. Figure 9).

Table 10: Hypervolumes of combined fronts on test data resulting from the two GP variants using different numbers of training days in the optimization.

Number of Training Days	GP-P	GP-S
20	238.050	249.358
50	248.365	254.087
100	254.313	258.357
150	257.312	258.190
200	256.855	256.838
250	258.176	257.436
300	258.303	256.898



(a) (b)

Figure 14: Combined fronts on the 1000 days test data resulting from using different numbers of training days in the optimization with (a) the GP-P variant and (b) the GP-S variant.

6. Summary and Conclusion

Two variants of a multi-objective genetic programming (GP) approach for learning electric vehicle (EV) charging policies on historical data, GP-P and GP-S, were proposed. In simulation experiments, these variants were evaluated on three settings: the training data used for learning, test data with similar characteristics but different daily user behavior, and a third set with entirely new user characteristics. Both variants achieved comparable performance on the training and test data and outperformed the considered baseline approaches in terms of efficiency and fairness. On the unseen data, however, GP-P maintained strong performance, whereas GP-S generalized less effectively.

An analysis of the resulting GP models showed that, with the default maximum model size of 50, the evolved models were too complex to be considered interpretable. However, additional experiments revealed that GP-P could be constrained to much simpler models—with only a slight drop in performance—while GP-S continued to require larger trees to perform well. Overall, GP-P demonstrated superior performance over GP-S: it achieves high efficiency and fairness, generalizes well to unseen data, and permits a tolerable reduction in performance in exchange for a significantly reduced model complexity, if interpretability is a concern.

A limitation of the proposed approach lies in its reliance on detailed historical data on user charging and driving behavior, which can be challenging to obtain in practice. However, as demonstrated by the simulation experiments, even with very limited historical data, charging policies with reasonable performance can still be learned.

The proposed approach is not limited to the specific problem setting considered in this work and can also be applied to scenarios beyond residential charging. Moreover, it is independent of the metric used to quantify user (dis-)satisfaction. In this study, the additional external charging time was used as metric. However, the satisfaction could be expressed differently, for example, based on the ratio of requested to delivered energy. It would also be

possible to consider other objectives than fairness and efficiency. The main prerequisite for adapting the approach to another scenario is that it has to be possible to evaluate charging policies in simulation on historical data. Naturally, it cannot be guaranteed that the approach will perform equally well in other scenarios. Hence, as future work, the proposed approach will be evaluated on other problem settings. One promising case is public charging, where charging requirements are known in advance. In such a setting, users do not return, and charging decisions no longer affect the energy requirements of future users. Consequently, logging station requests alone would be sufficient to generate training data, without the need to track external charging or driving behavior. In addition, this setting enables direct comparison with state-of-the-art optimization-based EV charging control approaches. Other interesting scenarios for evaluating the approach include cases with time-varying power limits and situations where users can access both AC and DC charging stations. Furthermore, the approach will be extended to address scenarios where not only the charging power but also the number of charging stations is limited. As the evaluation so far has been carried out only in simulation under simplifying assumptions, implementing and testing the approach in a real-world setting represents another promising direction for future work. Finally, since a general gap in the literature on fair EV charging is a lack of user studies, another interesting topic for future work would be to evaluate the perceived fairness and interpretability of the evolved policies in a user study.

CRediT authorship contribution statement

Steffen Limmer: Conceptualization, Investigation, Methodology, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing. **Angus Kenny:** Conceptualization, Investigation, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Tapabrata Ray** Conceptualization, Investigation, Supervision, Writing – review & editing. **Felix Lanfermann** Conceptualization, Project administration, Visualization, Writing – review & editing. **Hemant Kumar Singh** Conceptualization, Investigation, Writing – review & editing. **Andrea Castellani** Conceptualization, Writing – review & editing

Declaration of competing interest

Steffen Limmer, Felix Lanfermann and Andrea Castellani are employees of Honda Research Institute Europe, which filed a patent application disclosing the approach described in this paper.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used a local large language model (Llama 3.3) in order to improve language and readability. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Appendix A. Charging Control Algorithms

Algorithm 1: Charging control in the GP-P variant.

Input: list $\mathcal{N}_c$ of connected EVs, available energy $E^{\text{avail}} = P^{\text{lim}} \cdot \Delta t$, current timestep t

for EV in $\mathcal{N}_c$ **do**
  $EV.alloc = 0$;
  $X1, \ldots, XK$ = features(EV,t);
  $EV.v$ = gp_model($X1, \ldots, XK$);
end
$\mathcal{N}^{\text{sort}}$ = sort_decreasing_v($\mathcal{N}_c$);
for EV in $\mathcal{N}^{\text{sort}}$ **do**
  if $E^{\text{avail}} > 0$ **then**
    $EV.alloc = \min\{E^{\text{avail}}, \min\{EV.Cap - EV.Bat, EV.P^{\max} \cdot \Delta t\}\}$;
    $E^{\text{avail}} = E^{\text{avail}} - EV.alloc$;
end

Algorithm 2: Charging control in the GP-S variant.

Input: list $\mathcal{N}_c$ of connected EVs, available energy $E^{\text{avail}} = P^{\text{lim}} \cdot \Delta t$, current timestep t

$n^{\text{to_charge}} = 0$;
for EV in $\mathcal{N}_c$ **do**
  $EV.alloc = 0$;
  if $EV.Bat < EV.Cap$ **then**
    $X1, \ldots, XK$ = features(EV,t);
    $EV.v = |\text{gp_model}(X1, \ldots, XK)|$;
    $n^{\text{to_charge}}$ += 1;
  else
    $EV.v = 0$;
  end
end
$v^{\text{sum}} = \sum_{EV \in \mathcal{N}_c} EV.v$;
$E^{\text{left}} = E^{\text{avail}}$;
while $E^{\text{left}} > 0$ *&* $n^{\text{to_charge}} > 0$ **do**
  $n^{\text{to_charge}} = 0$;
  $E^* = E^{\text{left}}$;
  for EV in $\mathcal{N}_c$ **do**
    if $EV.v > 0$ **then**
      $e = \min\{\frac{EV.v}{v^{\text{sum}}} \cdot E^{\text{left}}, \min\{EV.Cap - EV.Bat - EV.alloc, EV.P^{\max} \cdot \Delta t\}\}$;
      E^* -= e;
      $EV.alloc$ += e;
      if $EV.Bat + EV.alloc < EV.Cap$ *&* $EV.alloc < EV.P^{\max} \cdot \Delta t$ **then**
        $n^{\text{to_charge}}$ += 1;
      else
        $EV.v = 0$;
      end
  end
  $v^{\text{sum}} = \sum_{EV \in \mathcal{N}_c} EV.v$;
  $E^{\text{left}} = E^*$;
end

References

[1] International Energy Agency, Global EV Outlook 2025, 2025. URL: https://www.iea.org/reports/global-ev-outlook-2025.

[2] G. Fridgen, M. Thimmel, M. Weibelzahl, L. Wolf, Smarter charging: Power allocation accounting for travel time of electric vehicle drivers, Transportation Research Part D: Transport and Environment 97 (2021) 102916. doi:https://doi.org/10.1016/j.trd.2021.102916.

[3] S. Limmer, Dynamic pricing for electric vehicle charging – A literature review, Energies 12 (2019). doi:10.3390/en12183574.

[4] C. Jin, J. Tang, P. Ghosh, Optimizing electric vehicle charging with energy storage in the electricity market, IEEE Trans. on Smart Grid 4 (2013) 311–320.
[5] J. R. Koza, Genetic programming: On the programming of computers by means of natural selection, MIT Press (1992).
[6] Y. Mei, Q. Chen, A. Lensen, B. Xue, M. Zhang, Explainable artificial intelligence by genetic programming: A survey, IEEE Transactions on Evolutionary Computation 27 (2023) 621–641. doi:10.1109/TEVC.2022.3225509.
[7] S. Limmer, T. Rodemann, Combination of charging policies for fair and efficient EV charging under limited capacity, in: 2024 22nd International Conference on Intelligent Systems Applications to Power Systems (ISAP), 2024, pp. 1–6. doi:10.1109/ISAP63260.2024.10744246.
[8] Q. Wang, X. Liu, J. Du, F. Kong, Smart charging for electric vehicles: A survey from the algorithmic perspective, IEEE Communications Surveys & Tutorials 18 (2016) 1500–1517. doi:10.1109/COMST.2016.2518628.
[9] O. Sadeghian, A. Oshnoei, B. Mohammadi-ivatloo, V. Vahidinasab, A. Anvari-Moghaddam, A comprehensive review on electric vehicles smart charging: Solutions, strategies, technologies, and challenges, Journal of Energy Storage 54 (2022) 105241. doi:https://doi.org/10.1016/j.est.2022.105241.
[10] Z. Zhao, C. K. Lee, X. Yan, H. Wang, Reinforcement learning for electric vehicle charging scheduling: A systematic review, Transportation Research Part E: Logistics and Transportation Review 190 (2024) 103698. doi:https://doi.org/10.1016/j.tre.2024.103698.
[11] S. Ø. Ottesen, Techno-economic Models in Smart Grids: Demand Side Flexibility Optimization for Bidding and Scheduling Problems, Ph.D. thesis, NTSU, Norwegian University of Science and Technology, 2017.
[12] D. T. Phan, J. Xiong, S. Ghosh, A distributed scheme for fair EV charging under transmission constraints, in: 2012 American Control Conference (ACC), 2012, pp. 1053–1058. doi:10.1109/ACC.2012.6315622.
[13] I. K. A. Aswantara, K. S. Ko, D. K. Sung, A centralized EV charging scheme based on user satisfaction fairness and cost, in: 2013 IEEE Innovative Smart Grid Technologies-Asia (ISGT Asia), 2013, pp. 1–4. doi:10.1109/ISGT-Asia.2013.6698730.
[14] M. Zeballos, A. Ferragut, F. Paganini, Proportional fairness for EV charging in overload, IEEE Transactions on Smart Grid 10 (2019) 6792–6801. doi:10.1109/TSG.2019.2911231.
[15] A. Hussain, P. Musilek, Fairness and utilitarianism in allocating energy to EVs during power contingencies using modified division rules, IEEE Transactions on Sustainable Energy 13 (2022) 1444–1456. doi:10.1109/TSTE.2022.3161897.
[16] D. Danner, H. de Meer, Quality of service and fairness for electric vehicle charging as a service, Energy Informatics 4 (2021) 1–20.
[17] O. Frendo, N. Gaertner, H. Stuckenschmidt, Real-time smart charging based on precomputed schedules, IEEE Transactions on Smart Grid 10 (2019) 6921–6932.
[18] O. Frendo, N. Gaertner, H. Stuckenschmidt, Improving smart charging prioritization by predicting electric vehicle departure time, IEEE Transactions on Intelligent Transportation Systems 22 (2021) 6646–6653. doi:10.1109/TITS.2020.2988648.
[19] O. Frendo, J. Graf, N. Gaertner, H. Stuckenschmidt, Data-driven smart charging for heterogeneous electric vehicle fleets, Energy and AI 1 (2020) 100007. doi:https://doi.org/10.1016/j.egyai.2020.100007.
[20] M. Shariatzadeh, C. Henggeler Antunes, M. A. Lopes, Charging scheduling in a workplace parking lot: Bi-objective optimization approaches through predictive analytics of electric vehicle users' charging behavior, Sustainable Energy, Grids and Networks 39 (2024) 101463. doi:https://doi.org/10.1016/j.segan.2024.101463.
[21] A. Khele, C. Jiang, H. Wang, Fairness-aware optimization of vehicle-to-vehicle interaction for smart EV charging coordination, in: 2023 IEEE/IAS 59th Industrial and Commercial Power Systems Technical Conference (I&CPS), 2023, pp. 1–9. doi:10.1109/ICPS57144.2023.10142094.
[22] Y. Zhou, N. Maxemchuk, X. Qian, C. Wang, The fair distribution of power to electric vehicles: An alternative to pricing, in: 2014 IEEE International Conference on Smart Grid Communications (SmartGridComm), 2014, pp. 686–691. doi:10.1109/SmartGridComm.2014.7007727.
[23] G. Tsaousoglou, J. S. Giraldo, P. Pinson, N. G. Paterakis, Fair and scalable electric vehicle charging under electrical grid constraints, IEEE Transactions on Intelligent Transportation Systems 24 (2023) 15169–15177. doi:10.1109/TITS.2023.3311509.
[24] E. H. Gerding, A. Perez-Diaz, H. Aziz, S. Gaspers, A. Marcu, N. Mattei, T. Walsh, Fair online allocation of perishable goods and its application to electric vehicle charging, in: Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19, International Joint Conferences on Artificial Intelligence Organization, 2019, pp. 5569–5575. doi:10.24963/ijcai.2019/773.
[25] T. Panayiotou, M. Mavrovouniotis, G. Ellinas, On the fair-efficient charging scheduling of electric vehicles in parking structures, in: 2021 IEEE International Intelligent Transportation Systems Conference (ITSC), 2021, pp. 1627–1634. doi:10.1109/ITSC48978.2021.9565024.
[26] T. Panayiotou, G. Ellinas, Fair-enough charging of electric vehicles, in: 2024 IEEE 99th Vehicular Technology Conference (VTC2024-Spring), 2024, pp. 01–06.
[27] O. Ardakanian, C. Rosenberg, S. Keshav, Distributed control of electric vehicle charging, in: Proceedings of the Fourth International Conference on Future Energy Systems, Association for Computing Machinery, New York, NY, USA, 2013, pp. 101–112. doi:10.1145/2487166.2487178.
[28] R. Carvalho, L. Buzna, R. Gibbens, F. Kelly, Critical behaviour in charging of electric vehicles, New Journal of Physics 17 (2015) 095001. doi:10.1088/1367-2630/17/9/095001.
[29] A. Al Zishan, M. Moghimi Haji, O. Ardakanian, Reputation-based fair power allocation to plug-in electric vehicles in the smart grid, in: 2020 ACM/IEEE 11th International Conference on Cyber-Physical Systems (ICCPS), 2020, pp. 63–74. doi:10.1109/ICCPS48487.2020.00014.
[30] D. Dreucci, Y. Yu, G. R. Chandra Mouli, A. Shekhar, P. Bauer, Centralised distribution grid congestion management through EV charging control considering fairness and priority, Applied Energy 384 (2025) 125417. doi:https://doi.org/10.1016/j.apenergy.2025.125417.
[31] D. Potoglou, R. Song, G. Santos, Public charging choices of electric vehicle users: A review and conceptual framework, Transportation Research Part D: Transport and Environment 121 (2023) 103824. doi:https://doi.org/10.1016/j.trd.2023.103824.
[32] S. Grimmelikhuijsen, Explaining why the computer says no: Algorithmic transparency affects the perceived trustworthiness of automated decision-making, Public Administration Review 83 (2023) 241–262. doi:https://doi.org/10.1111/puar.13483.
[33] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, IEEE Transactions on Evolutionary Computation 6 (2002) 182–197. doi:10.1109/4235.996017.

[34] Å.L. Sørensen, K. Lindberg, I. Sartori, I. Andresen, Analysis of residential EV energy flexibility potential based on real-world charging reports and smart meter data, Energy and Buildings 241 (2021) 110923. doi:https://doi.org/10.1016/j.enbuild.2021.110923.
[35] M. Weiss, K. C. Cloos, E. Helmers, Energy efficiency trade-offs in small to large electric vehicles, Environmental Sciences Europe 32 (2020) 1–17.
[36] H. Bekker, 2020 (full year) Germany: Best-selling electric cars by brand and model, https://www.best-selling-cars.com/germany/2020-full-year-germany-best-selling-electric-cars-by-brand-and-model, 2021. Accessed: April, 2025.
[37] EV Database, EV Database, https://ev-database.org/, 2024. Accessed: April, 2025.
[38] F.-A. Fortin, F.-M. De Rainville, M.-A. Gardner, M. Parizeau, C. Gagné, DEAP: Evolutionary algorithms made easy, Journal of Machine Learning Research 13 (2012) 2171–2175.
[39] J. Blank, K. Deb, pymoo: Multi-objective optimization in Python, IEEE Access 8 (2020) 89497–89509.
[40] F. Biscani, D. Izzo, A parallel global multiobjective framework for optimization: pagmo, Journal of Open Source Software 5 (2020) 2338. doi:10.21105/joss.02338.
[41] K. Shang, H. Ishibuchi, L. He, L. M. Pang, A survey on the hypervolume indicator in evolutionary multiobjective optimization, IEEE Transactions on Evolutionary Computation 25 (2021) 1–20. doi:10.1109/TEVC.2020.3013290.
[42] G. Yu, Y. Jin, M. Olhofer, A method for a posteriori identification of knee points based on solution density, in: Proc. IEEE CEC, 2018, pp. 1–8. doi:10.1109/CEC.2018.8477885.
[43] T. Schmitt, M. Hoffmann, T. Rodemann, J. Adamy, Incorporating human preferences in decision making for dynamic multi-objective optimization in model predictive control, Inventions 7 (2022). doi:10.3390/inventions7030046.
[44] A. Meurer, C. P. Smith, M. Paprocki, O. Čertík, S. B. Kirpichev, M. Rocklin, A. Kumar, S. Ivanov, J. K. Moore, S. Singh, T. Rathnayake, S. Vig, B. E. Granger, R. P. Muller, F. Bonazzi, H. Gupta, S. Vats, F. Johansson, F. Pedregosa, M. J. Curry, A. R. Terrel, v. Roučka, A. Saboo, I. Fernando, S. Kulal, R. Cimrman, A. Scopatz, Sympy: Symbolic computing in Python, PeerJ Computer Science 3 (2017) 1–23. doi:10.7717/peerj-cs.103.