Multi-objective Evolutionary Learning for Near Pareto-Optimal Optimization of Solar Deployment

Cooper Sigrist
csigrist@umass.edu
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Archimedes Li
archimedesli@umass.edu
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Alice Zhang
az2880@columbia.edu
Columbia University
New York, New York, USA

Adam Lechowicz
alechowicz@cs.umass.edu
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Noman Bashir
nbashir@mit.edu
Massachusetts Institute of Technology
Cambridge, Massachusetts, USA

Pichsinee Lertsaroj
plertsaroj@umass.edu
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Ryan Bahlous-Boldi
ryanbb@mit.edu
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Mohammad Hajiesmaili
hajiesmaili@cs.umass.edu
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Abstract

Existing residential rooftop photovoltaic (PV) installations in the United States are inequitable, as they are concentrated in high-income neighborhoods, and carbon-inefficient because they are often not located in electric grids dominated by fossil-fuel generators. Prior work, however, shows that prioritizing socioeconomic equity can also significantly increase the carbon efficiency of new installations. In this paper, we formalize the problem of site selection for rooftop PV installations as a multi-objective optimization problem, with metrics including *energy generation*, *carbon offsetting*, and *demographic equity*. We introduce a novel method called Evolutionary Value Assignment (EVA) that uses a neural network trained via evolutionary learning to select ideal sites for deployment. We evaluate our proposed approach in a case study using a dataset of U.S. solar generation and demographic information. Compared to projections of current installation trends, our method improves Carbon Efficiency by 43%, Income Equity by 41%, and Racial Equity by 24%, while increasing Energy Generation Potential by up to 10%. Therefore, our optimized placement can achieve the estimated carbon offset needed for *net-zero* emissions from electricity generation earlier than current deployment trends.

ACM Reference Format:
Cooper Sigrist, Archimedes Li, Alice Zhang, Adam Lechowicz, Noman Bashir, Pichsinee Lertsaroj, Ryan Bahlous-Boldi, and Mohammad Hajiesmaili. 2025. Multi-objective Evolutionary Learning for Near Pareto-Optimal Optimization of Solar Deployment. In *The 12th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '25), November 20–21, 2025, Golden, Colorado, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3736425.3770115

BuildSys '25, Golden, Colorado, USA

ACM ISBN 979-8-4007-1945-5/25/11
https://doi.org/10.1145/3736425.3770115

1 Introduction

The threat of climate change is becoming more apparent, marked by frequent flash floods and extreme heat waves even in areas previously considered immune to its effects [11, 37]. As a result, many governments are setting ambitious targets to reduce carbon emissions from economic activity [8]. To support these targets, particularly in the electricity generation sector, they are administering various incentive programs for residential households, such as tax credits and feed-in tariffs [4, 22]. These incentives, combined with rapidly declining costs, have made solar photovoltaic (PV) installations a popular approach for households to lower their energy bills while reducing their environmental impact.

Ideally, solar PV would be deployed in locations where it produces the most energy, offsets the most carbon-intensive power generation, and promotes socioeconomic and demographic equity. However, existing installations in the United States are not ideally located according to these metrics, despite solar PV being modular and easily deployable. While current installations are in areas with high Energy Generation Potential, they are concentrated in regions with relatively clean electric grids, as detailed in Section 3. This clean energy mix is not always from utility-scale solar; Arizona, for instance, has a high number of solar PV installations, but its grid is green due to significant nuclear and hydropower capacity [9]. Moreover, prior work has shown that existing solar deployments are concentrated in high-income neighborhoods, and continuing this trend would perpetuate socioeconomic and demographic inequities [33].

The current trends are a complex function of local climate, government incentives, economic factors, electric grid infrastructure, and political ideology. For example, as we discuss in Section 3.2, Kansas has few solar installations despite having high Energy Generation, a short payback period, and a fossil-fuel-heavy grid. However, not all locations are as straightforward, and these deployment objectives often conflict. Therefore, altering the current trajectory of solar PV installations toward ideal locations is a challenging task that requires effective policy and incentive design. A prerequisite for creating such policies is to identify locations that successfully balance these trade-offs. A simple approach can be to reduce the problem to a single objective using a linear weighted combination of the desired attributes. However, as we discuss in this paper, solutions derived from such combinations, while optimal for a specific set of weights, do not generalize well and are difficult to apply in practice. Therefore, a new approach is needed to achieve comparable performance while ensuring practical utility.

In this paper, we formulate the site-selection problem for residential rooftop PV installations as a multi-objective optimization problem that ranks sites (in our case, ZIP codes) using a super-objective function across Energy Generation Potential, Carbon Offset Potential, Racial Equity, and Income Equity metrics. To do so, we employ an evolutionary learning approach [34] that has been successful in other domains for efficiently solving similarly large-scale and complex multi-objective problems. Given a set of attributes for a site—including average sunlight hours, carbon offset per panel, demographics, and realized potential—our method assigns a score to each site indicating its desirability for new installations. Policymakers can then use this score to design incentives and devise strategies to promote solar PV installations in ideal locations.

Our approach of using evolutionary algorithms offers distinct benefits over a baseline method such as mixed-integer linear programming (MILP). A solution like MILP treats this as a site-selection problem, providing a set of ZIP codes that are considered optimal. However, not all ZIP codes in that set are equally desirable for new installations, and a targeted policy design that recognizes those differences is more practical. While a MILP-based solution likely could be extended to assign scores, we hypothesize that an evolutionary solution does so natively, consistently assigns high scores to sites near the Pareto-optimal front, and generalizes well to unseen ZIP codes. In evaluating our hypothesis, we make the following contributions.

(1) We formalize the rooftop solar siting problem as a multi-objective optimization problem. To do this, we create four distinct, numerical objectives (Energy Generation Potential, Carbon Offset Potential, Income Equity, and Racial Equity) that capture the contradictory goals of the solar siting task.
(2) We introduce a new algorithm, EVA (Evolutionary Value Assignment), for solving the solar siting problem and demonstrate that it is a significant improvement on both current solar siting practices and prior work. Our algorithm estimates a single super-objective using evolutionary learning. To our knowledge, this is the first work which uses this specific approach.
(3) We demonstrate the efficacy of our approach using a dataset that spans the continental U.S. [33]. We show that EVA is able to select sites for solar installations which improve on the currently estimated projection (i.e., the "*Status-Quo*") by up to 43% in Carbon Offset, 41% in Racial Equity, 24% in Income Equity, and 10% in Energy Generation.
(4) Using our MILP solutions we estimate a pareto-frontier of the tradeoff between Carbon Offset and Energy Generation. We show that our EVA method produces a model near this frontier, demonstrating its near-optimality.

2 Related Work

In this section, we review prior work on multi-objective optimization and the rooftop solar siting task in particular.

Equitability in Solar Installation Placement. The placement of solar installations can have far-reaching societal impacts. For instance, Dokshin and Thiede [7], O'Shaughnessy et al. [27], and Sunter et al. [36] study and quantify how existing adoption of rooftop PV in the U.S. is inequitable with respect to race and income. Beyond adoption, Crago et al. [5] study household financial returns due to solar PV, which often differ based on the business model adopted by the installation company—they find similar disparities along lines of race and income. These studies identify high upfront costs as one key reason for these disparities, which are further skewed by subsidies for rooftop solar that disproportionately benefit high-income homeowners [19]. Taken together, these distributional studies highlight that the existing Status-Quo of residential PV dissemination in the United States may exacerbate inequities along income or racial lines.

In the closest work to ours, Sigrist et al. [33] highlighted the inequitability and carbon-inefficiency of existing and future residential PV adoption in the U.S. They proposed simple "greedy algorithms" to remedy the problem (e.g., placement policies that prioritize areas for solar development based on race or income) and highlight the potential of optimized placement strategies. In our work, we build on these findings to apply a principled and holistic multi-objective lens to the problem – rather than myopically focusing on maximizing e.g., income equity, we focus on developing effective and practical optimization techniques for the solar siting problem that hope to simultaneously optimize multiple metrics such as equity, Carbon Offset, and Energy Generation. In doing so, we demonstrate how "a rising tide can lift all boats," identifying opportunities where e.g., optimizing for one metric can have positive effects on other metrics.

Multi-objective Optimization. Optimizing multiple objectives is a classic problem that arises in several application domains, and there have been several approaches proposed over the years – Sharma and Kumar [32] present a comprehensive overview. In a classic optimization sense, multi-objective problems can be solved using a mixed-integer linear program (MILP), simulated annealing [2], and particle swarm optimization [28].

However, due to the difficulty in optimizing multiple objectives simultaneously with gradient descent (as it is hard to specify a loss function that takes into account a good trade-off between the objectives), multi-objective optimization is often performed using evolutionary computation, such as the Non-Dominated Sorting Genetic Algorithm (NSGA) [6]. The high-level idea of these approaches is to explore the possible search space of solutions (i.e., neural network architectures and weights) in directions that are biased towards Pareto-optimality (defined formally in Section 3.3), which yields solutions that are "as good as possible" on all objectives simultaneously.

Another track of research considers solving multi-objective optimization problems by learning (or selecting) a weight for each objective, effectively reducing it to a single-objective optimization problem. Conventional Weighted Aggregation (CWA) [12] simply uses a selected weight for each objective at the beginning of a search run. While simple, this method can only explore a single point on the Pareto front during each run. To address this, Evolutionary Dynamic Weighted Aggregation (EDWA) [17] proposes learning weights for each objective to explore different areas of the Pareto front adaptively during search time. We build on these works by learning a mapping from features (i.e., attributes, see Section 4.3) to a single *score* that effectively acts as a single, combined objective.

Several studies apply multi-objective optimization to problems that concern the deployment of PV, including HVAC system operation [40], photovoltaic system design [18, 30, 35], supply chain networks [26], generation forecasting [39], and PV module cooling [31], among others. However, to the best of our knowledge, all of these works use multi-objective optimization to balance between *technical* objectives (e.g., energy production, lifespan), rather than the socioeconomic and geographic factors that are considered in the solar siting problem that we study.

3 Background and Preliminaries

In this section, we provide background and preliminaries on the energy grid, residential and rooftop photovoltaics, multi-objective optimization, and other relevant topics that inform the development of our approach.

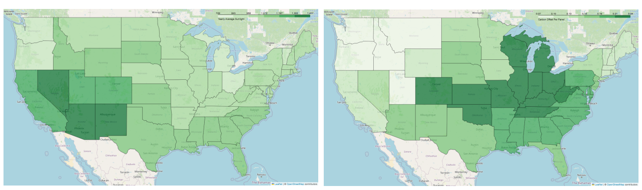

Energy Generation Potential Carbon Offset Potential

Figure 1: Energy generation Potential per installation does not correlate well with the Carbon Offset Potential per installation for the continental U.S.. A darker color (green) represents a high value for either of the potentials and vice versa.

3.1 Residential and Rooftop Solar PV

Amongst carbon-free Energy Generation sources, solar photovoltaics (PV) are perhaps the most mature and widespread technology. Photovoltaic deployments on building rooftops have seen significant growth over the last decade – from 2021 to 2022, the global installed capacity of these deployments jumped by 49% [10]. These include residential installations that deploy panels on the roof of, e.g., single- or multi-family buildings. Beyond the benefits of reduction in carbon emissions, these deployments reduce demand on the broader electric grid and offer opportunities for building owners to benefit financially, by selling electricity for instance. Currently, most rooftop PV in the U.S. has been deployed on an individual (i.e., unplanned) basis. However, there is mounting interest in models such as community solar [29] and targeted incentives [27] that orchestrate rooftop PV deployment in a more planned manner.

3.2 Electric Grid Diversity

For decarbonization goals, an ideal solar PV installation is located in a region with high Energy Generation Potential (i.e., abundant sunlight) as well as high Carbon Offset Potential (i.e., a grid supplied by carbon-intensive generation sources). However, due to the presence of utility-scale renewables and low-carbon generation in grids throughout the U.S., Energy Generation Potential and Carbon Offset Potential are not highly correlated – see Figure 1. States with the highest Energy Generation Potential (dark green color on the left map, primarily the West Coast of the U.S.) have a very low Carbon Offset Potential (light color on the right map). This implies that new PV installations in these states are less impactful. On the other hand, the states in Midwest U.S. (e.g., Missouri, Indiana, etc.) have the highest Carbon Offset Potential, as their existing grid generation mix primarily consists of fossil-fuel-based energy sources with high carbon emissions. However, their climates are not as readily suited for PV compared to some other states. This diversity in electric grids and the resulting disparity creates a tension between Energy Generation Potential and Carbon Offset Potential, which a site selection algorithm must navigate tactfully.

Moreover, even when these objectives align, as is the case in Kansas with the 11th highest Carbon Offset per installation and 14th highest Energy Generation per installation, that is no guarantee of adoption. Kansas, for example, has the 7th *lowest* proportion of its energy generation from Solar – 18th lowest when normalized to possible installations in the state. As explored in previous works [33], this may be due to relative costs of installation (i.e., average payback period in Kansas is 10 years) or political affiliation (Kansas voted the 21st most conservative in the 2024 presidential election[24]).

3.3 Multi-Objective Optimization

Multi-objective optimization generalizes classic optimization problems by optimizing (i.e., maximizing, minimizing) $n > 1$ objectives simultaneously. Formally, for some decision variable(s) θ and objective functions, $\{f_1, f_2, \cdots, f_n\}$, where each objective function $f_i(\theta) : \theta \rightarrow \mathbb{R}$ quantifies the quality of θ along some dimension of interest, a multi-objective optimization attempts to find "the best" solution (or series of solutions) under all n objectives. As some objectives in the set $\{f_1, f_2, \cdots, f_n\}$ may contradict each other, it may be impossible to find a solution that is simultaneously optimal for each objective considered individually. Thus, a solution to a multi-objective problem is considered optimal if it is *Pareto-optimal* (i.e., no other solution performs better on every single objective than the current solution – see below).

Pareto-optimality. A multi-objective problem often does not admit a unique optimal solution but rather a set of solutions that optimally trade-off between objectives, known as the *Pareto front*. A classic goal of multi-objective optimization is to find solutions that lie on the Pareto front, deemed *Pareto-optimal* solutions. Formally, a solution is Pareto-optimal over a set of solutions and objectives if there is no other solution that *dominates* it. A solution dominates another if it is *at least* as good on all objectives, and *strictly better* on at least one objective. Letting $\theta \in \mathcal{S}$ denote a solution θ in the feasible solution set $\mathcal{S}$, θ is Pareto-optimal over a set of objectives, $\mathcal{F}$, if and only if:

$$\nexists\, \theta' \in \mathcal{S} : (\forall f_i \in \mathcal{F} : f(\theta') \geq f_i(\theta)) \wedge (\exists f_j \in \mathcal{F} : f_j(\theta') > f_j(\theta)).$$

Linearly Aggregating Multiple Objectives. Assuming that all objectives are directionally aligned (i.e., all maximization or all minimization), a standard method to reduce n objectives into a single function is conventional weighted aggregation (CWA). This method balances different objectives by giving each a weight and summing them up to a single scalar, i.e.,

$$F(\theta) = \sum_{i=1}^{n} w_i f_i(\theta),$$

where w_i is a user-provided weight that is fixed throughout the optimization. CWA assumes some prior knowledge of the objectives (such as the order of significance or the target optimal value). When prior knowledge about the objectives is unknown, a variant, dynamic weighted aggregation (DWA), may be used. This class of methods is characterized by automated learning of the weights, $w_1, \ldots, w_n$ through, e.g., a grid search of possible weights.

We note that a linear combination of objectives is often not sufficiently expressive for certain problems. For example, consider two students who are evaluated based on their final grade in two courses (i.e., two objectives). Under any linear weighted combination of these objectives (final grades), a final grade of 100 in the higher-weighted course and 40 in the lower-weighted course would be considered as good as or better than a score of 70 in each course. However, in most cases, a final grade of 40 is a *failing grade*, which might be considered substantially worse for e.g., the student's progression. In Section 5, we describe our main approach, EVA, which is able to learn these *non-linear relationships* for different objectives.

3.4 Evolutionary Learning

As discussed in the example above, common approaches such as CWA or DWA are easy to define and combine many objectives into a single objective, but have been shown to be limited in the portion of the Pareto front they can represent [13, 25]. Motivated by the non-linear structure of some multi-objective optimization problems, it is thus natural to consider something more expressive, such as learning a neural network. However, classical supervised learning is generally insufficient for learning Pareto-optimal solutions without strong prior knowledge on the structure of the problem [20]. To address this, we follow literature on *evolutionary learning* and genetic programming, which have seen success in multi-objective settings [14–16]. At a high-level, the approach we propose learns a nonlinear scoring function (or ranking function) in the form of a neural network – in our context, this function assigns scores to each potential solar site (higher scores indicating more favorable locations). The nonlinearity of the neural network allows these scores to be more expressive than a linear combination of the objectives – for instance, the scoring function is able to learn representations that better capture the structure of the student grade problem described in Section 3.3.

We introduce some terminology and briefly detail the evolutionary process that is used for learning here – see Section 5 for a formal description of our approach. For the purpose of evolution, each of the multiple objectives are represented by a *fitness* function (e.g., absolute or normalized value for this objective). The learning goal is to find a neural network which achieves high overall fitness (i.e., across all objectives, particularly Pareto-optimal) – to do so, we begin with a *population* (i.e., collection) of *networks* (a specific case of the more general *genome* used in evolution literature), which are a combination of a neural network architecture and a set of parameters (i.e., weights). To this population a *selection* step is applied to identify and select high-performing networks – these selected neural networks are then *mutated* by randomly adding, removing, or modifying nodes/connections in each network to generate a new population, referred to as the next *generation*. This process repeats for either a fixed number of generations or until a single network reaches a desired fitness score. See Section 5 for details on the mapping of each of these concepts to our approach for the solar siting problem.

4 Problem Formulation

In this section, we formalize the rooftop solar siting problem as a multi-objective optimization problem and define each of the objectives we use as metrics in subsequent sections.

4.1 The Multi-objective Rooftop Solar Siting Problem

Let $n \geq 1$ denote the number of installations to be built, and $L \geq 1$ denote a list of locations (e.g., ZIP codes) by $z_1, \ldots z_L$, where each z_l has a maximum capacity (i.e., available rooftop space) of c_l and k additional attributes (e.g., median income or average yearly sunlight) denoted by $z_l^{(1)}, \ldots, z_l^{(k)}$. A solution to the multi-objective rooftop solar siting problem is an assignment, $\mathcal{A}\colon \mathbb{R}^{L\times k} \to \mathbb{R}^L$, that maps each location (including its attributes) to a number of installations to be built. That is, $\mathcal{A}(z_i)$ is the number of installations that assignment $\mathcal{A}$ allocates to ZIP code z_i. A "best" solution to this problem should be Pareto-optimal over the space of considered siting strategies $\mathcal{S}$ and $j \geq 1$ objectives given by $\mathcal{F} = f_1(\mathcal{A}), \ldots, f_j(\mathcal{A})$, while not violating the capacity constraint, c_l, of any location. Formally, a solution, $\mathcal{A}^*$, is Pareto-optimal if and only if:

$$\begin{gathered} \nexists \mathcal{A} : \forall f_i \in \mathcal{F} : f_i(\mathcal{A}) \geq f_i(\mathcal{A}^*) \\ \mathcal{A}^*(z_l) \leq c_l : \forall l \in [L]. \end{gathered}$$

4.2 Objectives

As outlined in the introduction and related work, there are several natural considerations in the build-out of an energy system such as rooftop PV installations. Energy generation, carbon impact, and even socio-economic equity are all affected by the type and location of new installations. With this consideration, we formally define four *maximization* objectives that we will consider in our case study (see Section 6):

Energy Generation Potential. This objective estimates the total *new* energy generated after the addition of installations designated by a given assignment. In our case study, we use Energy Generation estimates for a given ZIP code provided by the Project Sunroof dataset [21], described in Section 6.1. This objective is denoted as f_{EG}.

Carbon Offset Potential. This objective estimates the total amount of avoided carbon dioxide emissions (in kg) due to the new locations of rooftop PV adhering to a given placement. In our case study, this estimation is calculated using the methodology of the SunSight dataset [33] (see Section 6.1) – this takes estimates of non-baseload carbon emissions (and equivalent greenhouse gas emissions) per kWh from EPA eGRID data [38] and multiplies them by the yearly average Energy Generation Potential taken from Project Sunroof. This objective is denoted as f_{CO}.

Income Equity. This objective measures the distribution of installations across income lines. In our case study, we define Income Equity as one minus the difference between the proportion of installations installed in high income neighborhoods versus the proportion of installations in low income neighborhoods. Formally, we define Income Equity, f_{IE}, as:

$$f_{\mathsf{IE}} = 1 - \frac{|\#\text{installations}_{\text{high income}} - \#\text{installations}_{\text{low income}}|}{\#\text{installations}_{\text{total}}}.$$

Racial Equity. This objective measures distribution of installations across lines of self-reported race. In our case study, we define it as one minus the difference between the proportion of installations in below-median Black population neighborhoods versus those in above-median Black population neighborhoods. Racial Equity, f_{RE}, is defined as:

$$f_{\mathsf{RE}} = 1 - \frac{|\#\text{installations}_{\text{low Black}} - \#\text{installations}_{\text{high Black}}|}{\#\text{installations}_{\text{total}}}.$$

We choose the Black Population Proportion as a representative case because it has been shown to be the most skewed demographic for rooftop PV [5, 33, 36].

Composite Performance Ratio (CPR). While co-optimizing the four objectives above, we also define an aggregate objective used in training our model and for statistical analysis. This aggregate objective is defined as the sum of each objective's improvement over the "Status-Quo" projection strategy described in Section 6.2, with each term expressed as a *ratio* of the achieved value to the Status-Quo value. We note that the CPR metric does not capture some real-world considerations. In particular, our ratio-to-Status-Quo approach effectively penalizes objectives that the Status-Quo projection already does well on, and those that change less as the assignment changes. However, we find that defining this aggregate metric is highly effective in training a model that outperforms a target baseline such as the Status-Quo (see Section 5.2 for details).

4.3 Normalized Attributes

Attributes are measurable properties of each location that are known apriori and can be used to inform an installation assignment. In our formulation, we define each attribute as a normalized quantity that falls in the range of $[0, 1]$. Formally, for each ZIP code, z_i, each one of its attributes, $z_i^{(k)}$ is normalized to $\tilde{z_i}^{(k)}$ as follows:

$$\tilde{z_i}^{(k)} = \frac{z_i^{(k)} - \min_{l\in[L]}\left(z_l^{(k)}\right)}{\max_{l\in[L]}\left(z_l^{(k)}\right) - \min_{l\in[L]}\left(z_l^{(k)}\right)}.$$

In our case study, these attributes (for each ZIP code) are taken primarily from the Project Sunroof dataset and American Community Survey dataset aggregated over a 5 year period from 2016 to 2020 (abbreviated to ACS5), which are included in the SunSight dataset (see Section 6.1). The five attributes we consider are: Energy Generation Potential per Panel, Carbon Offset Potential per Panel, Black population proportion, median income, and realized potential (as a percentage).

5 Methodology

In this section, we outline the detailed methodology of our EVA approach. We first define the EVA method in Section 5.1 before describing the technical details and challenges of fine-tuning our approach in Section 5.2.

5.1 Evolved Value Assignment

Our main method, Evolved Value Assignment (EVA), uses an evolutionary framework (introduced in Section 3.4) to learn a *scoring function* represented by a neural network denoted by $V_g : \mathbb{R}^5 \to \mathbb{R}$, which maps a location's normalized attributes, $\tilde{z}_i^{(1)}, \tilde{z}_i^{(2)}, \cdots, \tilde{z}_i^{(5)}$ (as described in Section 4.3), to a score that represents the *quality* of location z_i. Using V_g, we construct a installation assignment, $\mathcal{A}_g$, that places solar installations so as to maximize the total quality of the selected sites. Formally:

$$\mathcal{A}_g = \arg\max_{\mathcal{A}} \left(\sum_{l\in[L]} V_g(\tilde{z}_i^{(1)}, \tilde{z}_i^{(2)}, \cdots, \tilde{z}_i^{(5)}) * \mathcal{A}(z_l) \right).$$

This inferencing process to create an installation assignment is illustrated in Figure 2.

Evolutionary Framework. Our method, Evolutionary Value Assignment (EVA), is a variant of the **N**euro**e**volution of **A**ugmenting **T**opologies (NEAT) algorithm, which uses principles of evolution to optimize both the weights and the structure (topology) of a neural network [34].

(1) A population of neural networks are randomly initialized. For our setting, these are constrained to networks with five inputs (i.e., each of the attributes) and one output (i.e., the output score).
(2) Each of these networks is evaluated on the four objectives described in Section 4.2.
(3) A subset of these networks are selected via their evaluation on the four objectives and our *selection method* (see below). These selected networks then continue to the next step.
(4) The selected networks are used to create a new population via our *reproduction method* (see below). This new population is then considered the basis of a new *generation*, and the process repeats from step (2) onwards.
(5) After a chosen number of generations, g, we choose the final model to be the network with the highest *CPR score*. CPR score is further discussed in Section 7.3.

Now, we further elaborate on these steps.

Evaluation. Using each network's associated installation assignment, $\mathcal{A}_g$, each network is given a fitness vector, $\mathcal{F}_g \in \mathbb{R}^4$, with one entry for each of the four objectives described in Section 4.2. Formally:

$$\mathcal{F}_g = \left[f_{\mathsf{CO}}(\mathcal{A}_g),\ f_{\mathsf{EG}}(\mathcal{A}_g),\ f_{\mathsf{IE}}(\mathcal{A}_g),\ f_{\mathsf{RE}}(\mathcal{A}_g)\right].$$

To calculate $\mathcal{F}_g$, we use the SunSight dataset's simulator (Section 6.1), simulating a placement of two million panels – this number represents a tripling of the total rooftop solar PV compared to what exists in the SunSight dataset, and is based on industry projections for the number of installations in 2034 [33].

Selection. Our selection method is a variant of *lexicase selection* [16] modified to suit our use case. This modified version is based on the following steps. Selection starts by uniformly randomly choosing one objective, f, out of the set of four $f_{\mathtt{CO}}, f_{\mathtt{EG}}, f_{\mathtt{IE}}, f_{\mathtt{RE}}$ that has not yet been used during this iteration of the selection step. Next, we filter out a fraction of networks that have the lowest fitness on this specific objective – this fraction is the *objective selectivity* that can be manually or programmatically tuned for each objective function. These two steps (choosing an objective and filtering the population on it) repeat until all of the objectives have been chosen once. After selecting through each of the four objectives, the remaining population continues to the reproduction step. The size of this remaining population is controlled by a *survival threshold* hyperparameter.

Reproduction. Amongst the selected networks, *each unique pair* undergoes a process known as *crossover* to produce a new network, and each resulting network is then *mutated*.

Given a unique pair of networks, a *gene* represents either a node or connection in either network, where *homologous genes* are present in both networks and *disjoint genes* are found in only one. Crossover combines homologous genes by randomly selecting a copy from either network, while disjoint genes are taken from the fitter network (i.e. the network with a higher CPR score).

After this new network is created, it mutates. Let $i \longrightarrow j$ represent the connection from node i to j. Each *possible* mutation occurs probabilistically and independently:

- Add a node ($P = 0.2$): A random connection $i \longrightarrow j$ is deleted, a new node n is created, and two new connections ($i \longrightarrow n$ with weight 1 and $n \longrightarrow j$ with the weight of the original connection $i \longrightarrow j$) are created. *(Occurs at most once)*
- Delete a node ($P = 0.2$): Select a random node j. For all of j's incoming neighbors i, remove connection $i \longrightarrow j$. For all of j's outgoing neighbors k, remove connection $j \longrightarrow k$. *(Occurs at most once)*
- Add a connection ($P = 0.5$): Select random nodes i, j such that i is in an earlier layer than j. Create connection $i \longrightarrow j$ with a random weight drawn from a mean-zero normal distribution with $\sigma = 1$. *(Occurs at most once)*
- Delete a connection ($P = 0.5$): Select random, valid connection $i \longrightarrow j$ and remove it. *(Occurs at most once)*
- Change a weight or bias ($P = 0.8$): The weight or bias is updated by adding a sample from a mean-zero normal distribution with $\sigma = 0.5$. *(Occurs at most once for each connection or node)*

To implement crossover and mutation, we used the neat-python library [23]. We tuned the hyperparameters mentioned in the next section (Section 5.2), leaving other hyperparameters set to default values (for instance, the parameters of the mutation step). Our implementation also uses *elitism*, which means that the network with the highest CPR score (see Section 4.2) that is also selected is added to the new population without reproduction or mutations.

5.2 Hyperparameter Tuning on EVA

In our implementation of EVA, we tune the following hyperparameters: *survival threshold*, *population size*, *number of generations*, *training installation count*, and the *objective selectivity* for each objective. Below, we describe each hyperparameter.

Population Size. *Population size* is the number of neural networks that compete during each generation. A larger population tends to produce better final results, but exponentially increases the training time. For this reason, we used a population size of 30 networks for the subsequent results, unless specified otherwise.

Survival Threshold. *Survival threshold* is the proportion of the population that remains after the selection step. A low survival threshold has strong selective pressure, potentially converging to a high-quality solution faster at the expense of "genetic diversity" (i.e., less of the possible neural network search space is able to be explored). Conversely, a higher threshold has lower selective pressure but greater genetic diversity. We tune this hyperparameter over the range $[0.05, 0.5]$, choosing a final threshold of 0.3 for the results presented here. We chose this threshold as it led to the model having the highest CPR while strictly dominating Status-Quo.

Number of Generations. *The number of generations* is the number of iterations of the selection, reproduction, and mutation steps taken (i.e., during training) before a final model is chosen – this is analogous to the number of epochs in a supervised learning setting. We trained models for six generations, with each generation offering some improvement over the last. However, later generations offered smaller increases than earlier generations. While there is a potential for marginal performance improvements, which is highly desirable, computational limitations constrained us from testing with even more generations for the purpose of this work.

Training Installation Count. *Training installation count* is the number of installation assignments simulated during the evaluation step. Our presented model used a training installation count of $2 \cdot 10^6$ installations, because $2 \cdot 10^6$ is approximately the number of solar installations that must be added to the dataset's ZIP codes to reach *net-zero emissions* [33]. Beyond the presented model in the paper, we also tested training with a smaller installation count of $5 \cdot 10^5$ to reduce the computational cost, but the model performance degraded significantly compared to training on $2 \cdot 10^6$ installations.

Objective Selectivity. *Objective selectivity* controls the fraction of networks filtered out by lexicase for each objective. Given selectivity parameters, v_i, for each objective, f_i, and a survival threshold of h, the fraction of individuals selected s_i is given by:

$$s_i = h^{p_i}\,, \ \ \mathtt{for}\ \ p_i = \left(\frac{v_i}{\sum_{j=1}^{4} v_j}\right)$$

Put simply, a *higher* selectivity parameter for an objective means that it is under higher *selective pressure* (i.e., networks that perform poorly on that objective are more likely to be eliminated earlier during lexicase selection).

One of the primary goals of EVA is to outperform existing baseline strategies (e.g., see Section 6.2) in all objectives. One way to achieve this is by changing the selectivity hyperparameters for each objective so that the model properly balances each objective. To tune these parameters, our method iteratively evaluates a scoring

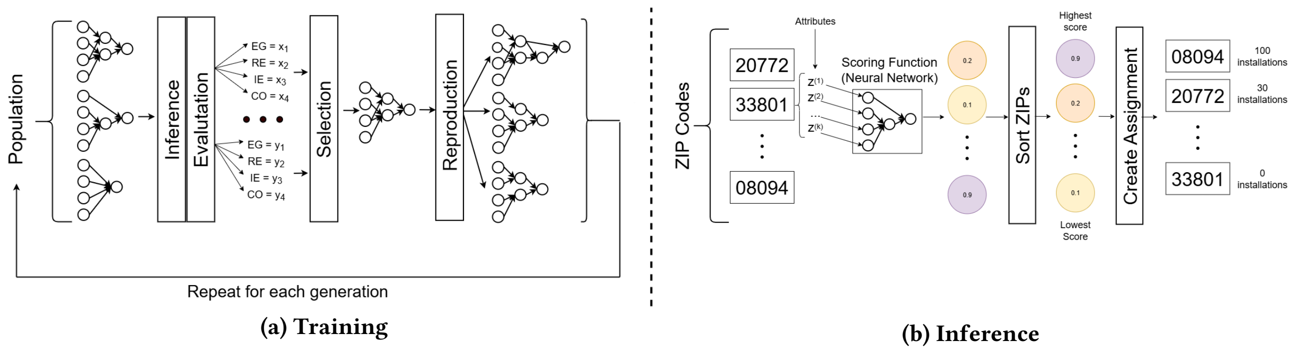

(a) Training

(b) Inference

Figure 2: Visualizing the Evolutionary Value Assignment (EVA) method, both training *(a)* and inference *(b)*. In the training stage, the EVA method performs evolutionary learning iterations for each of g generations (see Section 5.1 for details). The goal of the training stage is to obtain a *scoring function* (represented by a neural network) that can be used to score potential locations for PV siting (i.e., ZIP codes). In (*b*), we illustrate how these scores are used to create assignments of installations to locations.

function trained from scratch on the current selectivity parameter configuration, then increases the relative selectivity of any objective that falls short of the target baseline. In the next iteration, the increased selectivity increases the likelihood the new model will outperform the baseline in that objective. We detail the tuning method as follows:

(1) Let $v_{j,f}$ represent the selectivity of objective f for iteration j. Also, define all objective selectivity parameters for an iteration as $\vec{v}_j = [v_{j,f}]_{f\in\mathcal{F}}$.
(2) Let C_f be the objective value obtained by the Status-Quo projection (see Section 6.2) for each objective $f \in \mathcal{F}$.
(3) Initialize $v_{1,f} \leftarrow 1$ for each objective $f \in \mathcal{F}$.
(4) Perform the following iteration t times, or until the model outperforms the Status-Quo on all objectives.
   (a) For iteration j, create a new neural network scoring function using EVA with the objective selectivity parameters $\vec{v}_j$. Evaluate on two million panels and let $E_{j,f}$ be the model's projection for each objective $f \in \mathcal{F}$.
   (b) For each $f \in \mathcal{F}$, calculate the ratio between the performance of the currently selected model and the target baseline $R_{j,f} = E_{j,f}/C_f$. If $R_{j,f} < 1$, update the selectivity for the next iteration: $v_{j+1,f} \leftarrow v_{j,f}/R_{j,f}$.

Our results (in Section 7) use the "best" hyperparameters that we found using these procedures. In the next section, we detail the setup and baseline strategies for our case study.

6 Experimental Setup

This section presents our experimental setup, including the dataset we use and the baseline strategies we implement for comparison.

6.1 Dataset

We use data, visualization, and simulation tools from the SunSight dataset [33] for rooftop solar installation analysis. This dataset includes a toolkit that consists of a simulation testbed and visualization library, and the underlying data used is a curated combination of Google's Project Sunroof [21] and 5-year American Community Survey (i.e., U.S. Census, or ACS5) [3] data at a ZIP code granularity.

In our experiments, we simulate placing up to two million installations in the locations (i.e., ZIP codes) that are included in the solar data of the Project Sunroof dataset. The *attributes* defined in Section 4.3 are provided at a ZIP code granularity via the SunSight dataset's combination of the ACS5 and Project Sunroof datasets. The objective functions defined in Section 4.2 are calculated using the same ZIP code granularity data that combines the Project Sunroof and ACS5 datasets, giving estimates of quantities such as Carbon Offset, yearly solar Energy Generation, and demographics.

6.2 Baseline Strategies

Our baseline strategies include a *Status-Quo* strategy that projects current installation trends into the future, a round robin strategy presented in prior work [33], and an optimal Mixed Integer Linear Programming (MILP) solution to act as an upper bound. Particularly, these MILP solutions represent the best possible solutions of all possible CWA and DWA approaches.

Status-Quo Projection. The *Status-Quo* projection baseline estimates the current trajectory of U.S. rooftop PV installations in terms of the objectives defined in Section 4.2 for a number of additional installations (beyond existing installations as of 2025) up to two million installations. This simulates the case where the distribution of installations across the U.S. remains the same as additional installations are built, and evaluates each objective *after* new installations are factored into the mix. To simulate this, we first calculate the proportion of new small-scale solar installed in each state from March 2024 to March 2025 using the EIA small-scale PV dataset[1]. We also calculate the intra-state proportion of existing installations using the Project Sunroof dataset. For any given number of *new installations (i.e., to be added)*, we then allocate these installations to ZIP codes such that both proportions are maintained to give a reasonable estimate of the current trajectory.

We note that this does not account for any shifts in the distribution of new solar installations or significant future shifts in the (utility-scale) carbon intensity of different grid regions, but this does serve as an estimate of the current trajectory of rooftop solar.

Round Robin. The simple *Round Robin* strategy described in [33], is used as a baseline comparison. This algorithm alternates assignments between greedily placing in the ZIP codes with the highest Carbon Offset, highest Energy Generation, lowest median income, and highest Black population proportion sequentially.

Mixed Integer Linear Programming (MILP). Given a particular weighted linear combinations of our four objectives, we can formulate the solar siting problem as an MILP and obtain an instance-optimal solution. This allows us to approximate the Pareto-frontier of our multi-objective optimization.

In general, an MILP-based solution is not feasible at scale due to its information requirements – in particular, a MILP solver is able to observe *realizations* of the four objective values for any fine-grained assignment of installations to locations. This is in contrast to our EVA model, which only relies on five location attributes (see Section 4.3) that are reasonably known apriori. Furthermore, an MILP solution can specify a precise number of installations for each site, while our EVA model is constrained to produce a ranking (scoring) of locations, which is more valuable in practice (e.g., for incentivization policies). However, an MILP solution serves as a useful benchmark since this "full information" setting can be considered a target that our solution should approach.

MILP supports both discrete (panel counts) and continuous (auxiliary) decision variables while optimizing the linear aggregate objective subject to linear constraints. Since the equity objectives are defined as absolute deviations in panel allocation across demographic lines and are thus non-linear, we introduce auxiliary variables a_{RE} and a_{IE} constrained to equal the absolute deviation at optimality. These auxiliary variables are then penalized in the final objective function. Formally, the objective function is:

$$f(\mathcal{A}) = \left(\sum_{l \in [L]} \mathcal{A}(z_l) \left(w_1 \frac{\tilde{z}_i^{(EG)}}{n} + w_2 \frac{\tilde{z}_i^{(CO)}}{n} \right) \right) - w_3 \frac{a_{IE}}{n} - w_4 \frac{a_{RE}}{n}$$

We conduct an extensive grid search over the space of possible per-objective weights (w_1, w_2, w_3, w_4) and map out an approximate Pareto front using our results (see Section 3.3).

7 Experimental Results

This section presents the outcomes of our proposed approaches against the baseline strategies described in the previous section. We use the four objectives from Section 4.2 as the metrics for evaluation. The optimal MILP solutions are also analyzed to understand the nature of the tradeoff in our siting problem.

7.1 Status-Quo Comparison

In this section, we provide a detailed examination of EVA's performance, particularly as it relates to the Status-Quo projection. In Figure 3, we plot results for EVA at several numbers of "additional installations" – this reveals how the objectives of interest change relative to a Status-Quo at different scales or points in the future.

A notable observation is that EVA achieves a significantly higher Carbon Offset score (42% higher than Status-Quo) compared to its Energy Generation score (10% higher than Status-Quo). Carbon Offset first dramatically increases (for small numbers of additional installations) because EVA is able to fill out "low-hanging outliers"

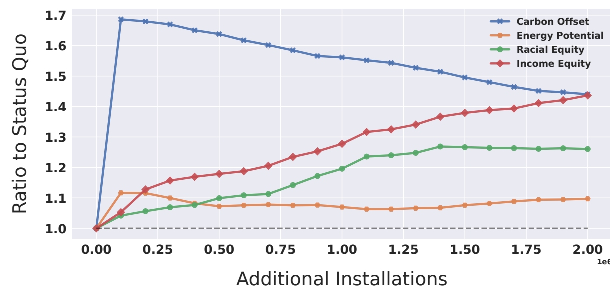

Figure 3: The performance of our EVA model on each objective as a ratio to Status-Quo. The dashed line at $y = 1$ represents the performance of the Status-Quo projection. Results are recorded every $100,000$ installations.

in e.g., the most carbon intensive grid regions – as these grid regions become saturated with installations, the Carbon Offset ratio to Status-Quo slightly decreases as the prospective panel siting locations become less-and-less carbon intense. We note that using the Status-Quo as a baseline exaggerates this difference between Energy Potential and Carbon Offset because the Status-Quo projection performs poorly in Carbon Offset and current installations are known to be skewed toward high Energy Generation [33]. However, it is notable that EVA still outperforms the Status-Quo on both metrics. Furthermore, its outperformance on Racial Equity (24% higher than the Status-Quo) and Income Equity (41% higher than the Status-Quo) also draws attention to the potential for significant improvement over our current trajectory.

The map in Figure 4 visualizes the scores given by EVA to each ZIP code in the dataset on a map, which reveals a preference for ZIP codes surrounding the *coal belt* of the U.S. (i.e., Appalachian states and Midwestern states such as West Virginia and Indiana), and some interesting disparities across states. These scores roughly correlate with the Carbon Offset Potential by state (e.g., see Figure 1) and are likely explained by differences in state-level solar adoption and incentives (e.g., state and utility subsidies) [33]. Another notable region with high scores from EVA is the Southwestern states, particularly Arizona and New Mexico. The scores in these states correlate to their high Energy Generation Potential (see Figure 1). Alternatively, Massachusetts and Washington, states with high number of installations and low Energy Generation Potential, are assigned the lowest scores by our model.

Key Takeaway. *EVA outperforms Status-Quo in Carbon Offset, Energy Potential, Racial Equity, and Income Equity (i.e., all objectives) for all installation counts evaluated. Notably, since net-zero carbon targets require fewer than this number of installations [33], we show that an EVA-based assignment would reach net-zero carbon with fewer installations (up to 41% fewer, taking years off of net-zero trajectories) while also achieving higher Energy Generation, additionally improving distributional measures relative to the Status-Quo.*

7.2 Pareto-optimality

In this section, we examine the Pareto front of the solar siting problem by aggregating the results of all MILP solutions generated by the grid search process described in Section 6.2. In doing so, we characterize the relationship between this estimated Pareto front,

Objective	Average Relative Weight of... Dominating Solutions	Average Relative Weight of... Pareto-optimal Solutions	Objective Ratio-to-Status-Quo Correlation w/ CPR (Pearson r)	Average % Contribution to CPR
Carbon Offset	0.493	0.492	0.888	19.8%
Energy Generation	0.103	0.155	0.656	25.6%
Racial Equity	0.210	0.200	0.511	25.6%
Income Equity	0.195	0.154	0.809	28.9%

Table 1: Statistics of all MILP-based solutions across objectives. Dominating solutions exhibit ratio-to-Status-Quo values > 1 for all objectives, and the Pareto-optimal solutions (amongst MILP solution set) are defined formally in Section 3.3. We also report the *correlation* between the ratio-to-Status-Quo of an objective and its contribution to the CPR score Finally, we report the *average weight* (i.e., w_i) given to each objective by a specific class of solutions. Across both solution classes, we report the *average contribution* (as a percentage) to the CPR score (defined in Section 4.2) – a higher value means that a higher percentage of the CPR score (on average) comes from good performance on this particular objective.

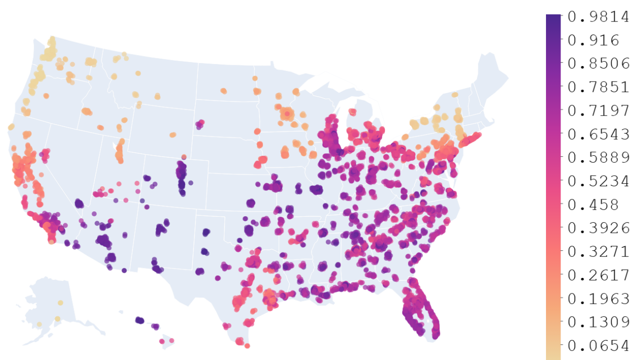

Figure 4: Map of the scores for ZIP codes generated by our model. Higher scoring ZIP codes are picked first for siting.

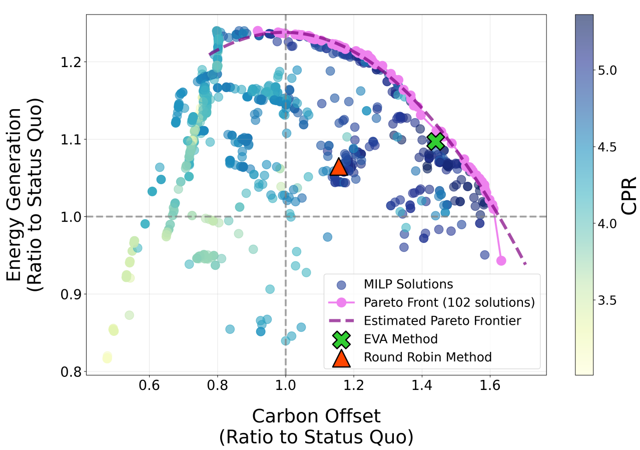

Figure 5: The Carbon Offset and Energy Generation tradeoff for all strategies. Each point represents a potential solution, colored by its CPR (defined in Section 4.2). The Pareto front of all MILP solutions is drawn in purple, estimated using a quadratic line of best fit. The green "X" marks the solution found by the EVA method. The dashed vertical and horizontal gray lines represent the Status-Quo reference point (1.0, 1.0).

our EVA-based model, and other baselines. In Figure 5, we plot this estimated Pareto front of the MILP model over the Carbon Offset and Energy Generation objectives (i.e., f_{CO} and f_{EG}, respectively). We observe that the previously studied Round Robin strategy lies "behind the front", implying that it is sub-optimal for siting, in terms of Carbon Offset and Energy Generation, while EVA is near the front, meaning that it is close to Pareto-optimal.

Key Takeaway. *The EVA-based solution lies close to the MILP-generated Pareto frontier for Carbon Offset and Energy Generation. Only one MILP solution (out of 1,942) dominates the EVA-generated solution across all objectives. This demonstrates the near-optimality of our EVA model – EVA works with more limited information (and is thus more able to generalize), but still performs as well as these clairvoyant MILP baselines.*

7.3 Comparison of Objectives

In order to better understand how each objective contributes to overall performance, we analyze our set of MILP solutions by computing pairwise Pearson correlation coefficients (r) that quantify the strength of alignment between the relative performance of each objective (Carbon Offset, Energy Generation, Racial Equity, and Income Equity) and the Composite Performance Ratio (CPR, see Section 4.2). Large correlation coefficient r-values between an objective and CPR suggest that objective may be important and correlated with the other objectives. In Table 1, we observe that Carbon Offset's ratio-to-Status-Quo term most highly correlates with CPR ($r = 0.888$), followed by Income Equity ($r = 0.809$).

Recall the definition of the CPR score (see Section 4.2) – each objective is represented by a ratio to the Status-Quo projection. To observe any differences in scale (i.e., if a single objective's ratio-to-Status-Quo term is much larger than other objectives on average), we also report the average *contribution* of each objective (as a percentage) towards CPR. Interestingly, while Carbon Offset and Income Equity are both highly correlated with CPR, they significantly differ in their average contribution to CPR – namely, Carbon Offset contributes 19.8% to the CPR on average, while Income Equity contributes 28.9% on average. This suggests an interesting dynamic borne from the definition of CPR – for example, if a certain solution doubles the Carbon Offset and simultaneously increases Income Equity from 0.2 to 0.6, the corresponding ratios to the Status-Quo are 2 and 3 for Carbon Offset and Income Equity, respectively, meaning that Income Equity contributes 50% more to the CPR score in this case. In this sense, the design of the CPR score itself influences our results. Since Carbon Offset still exhibits the best correlation

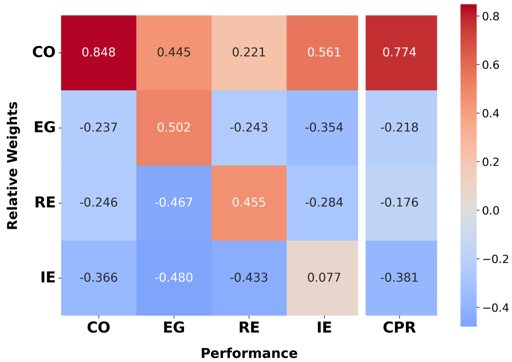

Figure 6: This heatmap displays the Pearson's r correlations between the relative weights assigned to each objective (rows) and the performance achieved on each objective as well as the CPR (columns). Each cell in the grid is colored by its correlation, ranging from red for positive correlation to blue for negative correlation.

with CPR despite these scoring dynamics, the rightmost columns of Table 1 suggest that Carbon Offset is particularly important for increases in CPR.

We investigate this further in Figure 6, which visualizes the Pearson r correlation coefficients between the *relative weights* of each objective (on average) and performance in each objective. Note that in the MILP formulation, the relative weight assigned to an objective is exactly interpreted as the *importance* that the solver places on this objective. The final column in Figure 6 also visualizes these same correlations between relative weights and the overall CPR score. As expected, the relative weights for any given objective correlate positively with its own performance ratio (e.g., the diagonal of red (positive) cells). Out of all objectives, the weighting of Carbon Offset exhibits the strongest correlation with its own performance ($r = 0.848$). However, perhaps most surprisingly, this figure also shows that the Carbon Offset weighting correlates positively with *every other performance metric*, and we find that this synergy is unique to Carbon Offset. This means that, perhaps paradoxically, placing *more* importance on the Carbon Offset objective actually improves performance on all objectives (up to a certain point). For other objectives, their relative weightings correlate negatively with performance on other objectives. This suggests that Carbon Offset is the most "efficient" single objective to focus on – i.e., not only does it strongly correlate with its own performance ratio, it also positively correlates with the other objectives of interest in our setting. We note that this statistical analysis is limited by two factors: *i)* we only consider MILP solutions that linearly aggregate the multiple objectives and *ii)* we only consider the Pearson correlation statistic, which is itself linear. Thus, these results may not fully capture complex interactions or nonlinear dependencies between objectives and overall performance.

Key Takeaway. *Through statistical analysis of our MILP solution set that explores the search space of the problem, we investigate how our four objectives interact with one another on our dataset of the continental U.S. Interestingly, we find that Carbon Offset stands out as the singular objective that both exhibits the strongest correlations with overall performance (i.e., CPR score), as well as correlating strongly and positively with all other objectives, such as Energy Generation.*

8 Conclusion

In this work, we define and implement a training algorithm called EVA, which applies NEAT-based evolutionary neural networks with a variant of lexicase selection for a multi-objective solar siting problem. Our work demonstrates that a model produced using this method substantially outperforms both existing Status-Quo projections and simpler algorithms considered by prior work. Our model improves on these baselines across multiple key objectives, including Carbon Offset, Racial Equity, Income Equity, and Energy Generation. Our proposed EVA approach leverages the expressiveness of neural networks to *learn* a function, distilling multiple objective values into a single score.

Simulating the solar siting problem using available data sets of solar generation and demographic data for the continental U.S., we find that our NEAT-based strategy generates placements for new rooftop PV installations that improve overall carbon emissions offset by 42% and significantly address distributional equity concerns while also increasing the expected Energy Generation by 10%. These positive results highlight the importance of site selection for new solar panel installations. Using projections, we estimate that an optimized strategy can accelerate progress toward goals such as net-zero electricity generation by several years.

Acknowledgments

This research is supported by National Science Foundation grants 2045641 and 2325956, and the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0024386.

Disclaimers

References

[1] U.S. Energy Information Administration. 2025. Electric Power Monthly. https://www.eia.gov/electricity/monthly/. [Accessed 29-07-2025].

[2] Khalil Amine. 2019. Multiobjective Simulated Annealing: Principles and Algorithm Variants. *Advances in Operations Research* 2019 (May 2019), 1–13. doi:10.1155/2019/8134674

[3] U.S. Census Bureau. 2022. American Community Survey 2017-2021 5-Year Data Release.

[4] Christine Crago and Ilya Chernyakhovskiy. 2017. Are policy incentives for solar power effective? Evidence from residential installations in the Northeast. *Journal of Environmental Economics and Management* (2017).

[5] Christine L. Crago, Emma Grazier, and Dwayne Breger. 2023. Income and racial disparities in financial returns from solar PV deployment. *Energy Economics* 117 (Jan. 2023), 106409. doi:10.1016/j.eneco.2022.106409

[6] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2002), 182–197. doi:10.1109/4235.996017

[7] Fedor A. Dokshin and Brian C. Thiede. 2023. Revised estimates of racial and ethnic disparities in rooftop photovoltaic deployment in the United States. *Nature Sustainability* 6 (2023), 752 – 755. https://api.semanticscholar.org/CorpusID:259398453

[8] Earth.Org. 2023. *3 Leading Countries in Climate Policy*. https://earth.org/countries-climate-policy/ Accessed on 1 August 2025.

[9] Electricity Maps. 2025. Electricity Map. https://www.electricitymap.org/map.

[10] SolarPower Europe. 2023. Solar for the masses as global solar rooftops skyrocket by 50%. https://www.solarpowereurope.org/press-releases/new-report-solar-for-the-masses-as-global-solar-rooftops-skyrocket-by-50-2

[11] European Environment Agency. 2024. *Extreme weather: floods, droughts and heatwaves*. https://www.eea.europa.eu/en/topics/in-depth/extreme-weather-floods-droughts-and-heatwaves Accessed on 1 August 2025.

[12] P. J. Fleming. 1985. Computer aided design of regulators using multiobjective optimization. *Annual Review in Automatic Programming* 13 (1985), 47–52. doi:10.1016/0066-4138(85)90462-8

[13] Ellen Hidemi Fukuda and L. Mauricio Graña Drummond. 2014. A Survey On Multiobjective Descent Methods. *Pesquisa Operacional* 34 (2014), 585–620. https://api.semanticscholar.org/CorpusID:120307972

[14] Thomas Helmuth, Nicholas Freitag McPhee, and Lee Spector. 2016. Lexicase Selection for Program Synthesis: A Diversity Analysis. In *Genetic Programming Theory and Practice XIII*, Rick Riolo, W.P. Worzel, Mark Kotanchek, and Arthur Kordon (Eds.). Springer International Publishing, Cham, 151–167. doi:10.1007/978-3-319-34223-8_9 Series Title: Genetic and Evolutionary Computation.

[15] Thomas Helmuth, Edward Pantridge, and Lee Spector. 2019. Lexicase selection of specialists. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, Prague Czech Republic, 1030–1038. doi:10.1145/3321707.3321875

[16] Thomas Helmuth, Lee Spector, and James Matheson. 2015. Solving Uncompromising Problems With Lexicase Selection. *IEEE Transactions on Evolutionary Computation* 19, 5 (Oct. 2015), 630–643. doi:10.1109/TEVC.2014.2362729

[17] Yaochu Jin, Markus Olhofer, and Bernhard Sendhoff. 2001. Dynamic Weighted Aggregation for evolutionary multi-objective optimization: why does it work and how?. In *Proceedings of the 3rd Annual Conference on Genetic and Evolutionary Computation* (San Francisco, California) *(GECCO'01)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1042–1049.

[18] Choongwan Koo, Taehoon Hong, Minhyun Lee, and Jimin Kim. 2016. An integrated multi-objective optimization model for determining the optimal solution in implementing the rooftop photovoltaic system. *Renewable and Sustainable Energy Reviews* 57 (May 2016), 822–837. doi:10.1016/j.rser.2015.12.205

[19] V. Lekavičius, V. Bobinaitė, A. Galinis, and A. Pažėraitė. 2020. Distributional impacts of investment subsidies for residential energy technologies. *Renewable and Sustainable Energy Reviews* 130 (Sept. 2020), 109961. doi:10.1016/j.rser.2020.109961

[20] Suyun Liu and Luís Nunes Vicente. 2019. The stochastic multi-gradient algorithm for multi-objective optimization and its application to supervised machine learning. *ArXiv* abs/1907.04472 (2019). https://api.semanticscholar.org/CorpusID:195874172

[21] Google LLC. [n. d.]. Project Sunroof. https://sunroof.withgoogle.com.

[22] Daniel C. Matisoff and Erik P. Johnson. 2017. The comparative effectiveness of residential solar incentives. *Energy Policy* 108 (Sept. 2017), 44–54.

[23] Alan McIntyre, Matt Kallada, Cesar G. Miguel, Carolina Feher de Silva, and Marcio Lobo Netto. [n. d.]. *neat-python*.

[24] WBUR Newsroom. 2024. Map: State-by-state results in the 2024 presidential election. *WBUR News* (5 November 2024). http://wbur.org/news/2024/11/05/results-president-2024-map-states-harris-trump

[25] P. Ngatchou, A. Zarei, and A. El-Sharkawi. 2005. Pareto Multi Objective Optimization. In *Proceedings of the 13th International Conference on, Intelligent Systems Application to Power Systems*. 84–91. doi:10.1109/ISAP.2005.1599245

[26] Maryam Nili, Seyed Mohammad Seyedhosseini, Mohammad Saeed Jabalameli, and Ehsan Dehghani. 2021. A multi-objective optimization model to sustainable closed-loop solar photovoltaic supply chain network design: A case study in Iran. *Renewable and Sustainable Energy Reviews* 150 (Oct. 2021), 111428. doi:10.1016/j.rser.2021.111428

[27] Eric O'Shaughnessy, Galen Barbose, Ryan Wiser, Sydney Forrester, and Naïm Darghouth. 2020. The impact of policies and business models on income equity in rooftop solar adoption. *Nature Energy* 6, 1 (Nov. 2020), 84–91. doi:10.1038/s41560-020-00724-2

[28] K. E. Parsopoulos and M. N. Vrahatis. 2002. Particle swarm optimization method in multiobjective problems. In *Proceedings of the 2002 ACM Symposium on Applied Computing* (Madrid, Spain) *(SAC '02)*. Association for Computing Machinery, New York, NY, USA, 603–607. doi:10.1145/508791.508907

[29] Michael Peters, Shane Fudge, Angela High-Pippert, Vincent Carragher, and Steven M. Hoffman. 2018. Community solar initiatives in the United States of America: Comparisons with – and lessons for – the UK and other European countries. *Energy Policy* 121 (Oct. 2018), 355–364. doi:10.1016/j.enpol.2018.06.022

[30] Hussein Mohammed Ridha, Chandima Gomes, Hashim Hizam, Masoud Ahmadipour, Ali Asghar Heidari, and Huiling Chen. 2021. Multi-objective optimization and multi-criteria decision-making methods for optimal design of standalone photovoltaic system: A comprehensive review. *Renewable and Sustainable Energy Reviews* 135 (Jan. 2021), 110202. doi:10.1016/j.rser.2020.110202

[31] Mohammad Hassan Shahverdian, Ali Sohani, Hoseyn Sayyaadi, Saman Samiezadeh, Mohammad Hossein Doranehgard, Nader Karimi, and Larry K.B. Li. 2021. A dynamic multi-objective optimization procedure for water cooling of a photovoltaic module. *Sustainable Energy Technologies and Assessments* 45 (June 2021), 101111. doi:10.1016/j.seta.2021.101111

[32] Shubhkirti Sharma and Vijay Kumar. 2022. A Comprehensive Review on Multi-objective Optimization Techniques: Past, Present and Future. *Archives of Computational Methods in Engineering* 29, 7 (Nov. 2022), 5605–5633. doi:10.1007/s11831-022-09778-9

[33] Cooper Sigrist, Adam Lechowicz, Jovan Champ, Noman Bashir, and Mohammad Hajiesmaili. 2025. Lost in Siting: The Hidden Carbon Cost of Inequitable Residential Solar Installations. In *Proceedings of the Sixteenth ACM International Conference on Future Energy Systems* (Rotterdam, Netherlands) *(e-Energy '25)*. Association for Computing Machinery. arXiv:2501.13868 [cs.CE] https://arxiv.org/abs/2501.13868

[34] Kenneth O. Stanley and Risto Miikkulainen. 2002. Evolving Neural Networks Through Augmenting Topologies. *Evolutionary Computation* 10, 2 (2002), 99–127. http://nn.cs.utexas.edu/?stanley:ec02

[35] Allan R. Starke, José M. Cardemil, Rodrigo Escobar, and Sergio Colle. 2018. Multi-objective optimization of hybrid CSP+PV system using genetic algorithm. *Energy* 147 (March 2018), 490–503. doi:10.1016/j.energy.2017.12.116

[36] Deborah A. Sunter, Sergio Castellanos, and Daniel M. Kammen. 2019. Disparities in rooftop photovoltaics deployment in the United States by race and ethnicity. *Nature Sustainability* 2, 1 (Jan. 2019), 71–76.

[37] UCAR Center for Science Education. 2023. *Climate Change: Regional Impacts*. https://scied.ucar.edu/learning-zone/climate-change-impacts/regional Accessed on 1 August 2025.

[38] U.S. Environmental Protection Agency. 2015. eGRID2012 Summary Tables. https://www.epa.gov/sites/default/files/2015-10/documents/egrid2012_summarytables_0.pdf Available from: https://www.epa.gov/sites/default/files/2015-10/documents/egrid2012_summarytables_0.pdf [Accessed 2024-07-23]..

[39] Jianzhou Wang, Yilin Zhou, and Zhiwu Li. 2022. Hour-ahead photovoltaic generation forecasting method based on machine learning and multi objective optimization algorithm. *Applied Energy* 312 (April 2022), 118725. doi:10.1016/j.apenergy.2022.118725

[40] Man Wang, Jiangfeng Wang, Pan Zhao, and Yiping Dai. 2015. Multi-objective optimization of a combined cooling, heating and power system driven by solar energy. *Energy Conversion and Management* 89 (Jan. 2015), 289–297. doi:10.1016/j.enconman.2014.10.009