applied sciences

MDPI

Article

Explainable GeoAI for Solar Photovoltaic Site Selection: An Interpretable Gradient Boosting Framework

Batuhan Kilic [1,*], Atakan Bilgili [2], Mert Gurturk [3] and Yasin Ferit Uguz [2]

[1] Department of Civil Engineering, Firat University, Elazığ 23119, Türkiye
[2] Department of Geomatic Engineering, Yildiz Technical University, Istanbul 34220, Türkiye; atakanb@yildiz.edu.tr (A.B.); ferit.uguz@std.yildiz.edu.tr (Y.F.U.)
[3] Department of Civil Engineering, Adiyaman University, Adiyaman 02040, Türkiye; mgurturk@adiyaman.edu.tr
* Correspondence: batuhankilic@firat.edu.tr

Abstract

Site selection for solar power is largely based on expert-weighted GIS multi-criteria methods. We present an interpretable machine learning framework that learns siting preferences from existing installations in Adıyaman, Elazığ, and Malatya, eastern Türkiye. Cells of 166 ground-mounted plants and an equal number of undeveloped cells were collected under an identical restriction mask, and plants were held out as whole units. Fourteen attributes on a 30 m grid were condensed to eleven by association analysis, mutual information, and tree-based backward selection. CatBoost, XGBoost, and LightGBM, tuned by plant-grouped cross-validation, reached an accuracy of about 0.72 in repeated plant-grouped cross-validation, with no persistent difference between them, and transfer between provinces was weak. In all resampled folds, Shapley Additive exPlanations ranked distance to substations as most important, followed by distance to airports. Global horizontal irradiance varied little within each province and had little influence. The result is a relative suitability score rather than a calibrated probability. The high and very high classes cover 14.0% of the candidate surface (10,971 km^2) and contain 76.1% of the held-out plant cells, compared to 27.3% for the Best–Worst Method benchmark. Within the sampled region, infrastructure rather than the solar resource explains where plants have been built.

Keywords: solar photovoltaic; site selection; GeoAI; machine learning; Shapley Additive exPlanations; explainable artificial intelligence (XAI); CatBoost

1. Introduction

The first Global Stocktake under the Paris Agreement, adopted at COP28 in Dubai in December 2023, called on Parties to triple global renewable energy capacity, double the annual rate of energy efficiency improvement by 2030, and transition away from fossil fuels in energy systems [1]. In November 2026, Türkiye will host and chair COP31 in Antalya [2]. Its own commitments are correspondingly ambitious: carbon neutrality by 2053, a 41% reduction in greenhouse gas emissions below a business-as-usual trajectory by 2030 [3], and installed solar photovoltaic (PV) capacity of 32.9 GW by 2030 and 52.9 GW by 2035 under the National Energy Plan [4]. By the end of August 2026, solar accounted for 22% of Türkiye's 126.9 GW installed generating capacity, approximately 28 GW, making it the second-largest source after hydropower [5]. Sustaining this trajectory therefore shifts the central constraint from ambition to land.

Academic Editor: Jose D. Marcos

Received: 7 October 2026
Revised: 23 November 2026
Accepted: 24 November 2026
Published: 26 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

Achieving capacity targets of this scale is fundamentally a land-allocation problem. Utility-scale PV requires flat, well-irradiated terrain with access to roads and transmission infrastructure, competing with agriculture, settlement expansion, and conservation objectives [6,7]. This competition intensifies as development progresses because the most favorable parcels are typically developed first. Plant location determines annual energy yield, grid-connection costs, agricultural land displacement, and exposure to natural hazards. Site selection is therefore the critical stage at which national capacity targets become actionable projects.

Over the last two decades, multi-criteria decision making (MCDM) integrated with geographic information systems (GIS) has been the dominant approach to solar PV site selection. The most widely used methods are the analytic hierarchy process (AHP) and its fuzzy extensions, often combined with weighted linear combination, distance-based methods such as TOPSIS, or outranking methods such as ELECTRE [8,9]. Weighting schemes have progressively evolved, including Fermatean fuzzy [10] and Pythagorean fuzzy [11] formulations, while applications have expanded to continental scales [12] and humid, forest-dominated environments [13]. In Türkiye, such approaches have been applied in Karaman [14], Kayseri [15], Kahramanmaraş [16], Erzurum [17], İzmir [18], and Gaziantep [19].

GIS-MCDM approaches nevertheless have several well-documented limitations. First, criterion weights depend on expert judgment: different panels or aggregation operators can produce different suitability maps from the same input data. Sensitivity analysis measures map stability under weight perturbation [20] but not the validity of the weights themselves. Second, linear additive aggregation assumes constant compensation among criteria and separable effects, assumptions that are often unrealistic. Third, expert-weighted suitability maps are difficult to validate against observed development patterns [21].

Recent machine learning (ML) approaches provide an alternative by inferring criterion importance from the observed spatial distribution of existing installations. Sun et al. [21] used China's national inventory of large-scale PV installations to train multilayer perceptron, random forest, and XGBoost models stratified by land cover type, deriving feature influence from Shapley additive explanations (SHAP) values and feature importance measures. Sachit et al. [22] trained random forest, support vector machine, and multilayer perceptron models on more than 55,000 operational wind and solar plants worldwide, deriving decision weights from the fitted models rather than expert panels. The approach has since been applied in Australia [23], Adıyaman [24], the Cholistan Desert [25], nationally in China [26,27], Japan [28], South Africa [29], India [30], China in combination with transmission-corridor planning [31], and the Middle East with a TabNet deep-learning model [32]. Related studies have linked neural and tree-based learners to MCDM outputs [33,34] and extended the approach to hybrid solar–wind installations [35].

Three limitations remain evident in this literature. First, the range of algorithms is narrow, with most studies relying on random forests and shallow feed-forward neural networks. Gradient boosting has only recently been adopted, with XGBoost [36] used in four studies [21,25,29,31], LightGBM [37] in one [31], and CatBoost [38] only as a baseline for a deep-learning model [32]. Site-selection datasets commonly contain categorical predictors such as land use and land cover (LULC), which the three libraries handle differently, so their relative performance on such data remains an open empirical question.

Second, model selection is typically based on accuracy or AUC from a single held-out partition, with the highest-scoring model reported as the best [21,22,25,31]. Differences of one or two percentage points on a single partition may, however, arise

from resampling variability rather than genuine differences between algorithms, and when SHAP values from the selected model are used as feature weights [26,31], this uncertainty is propagated into the suitability map. Dietterich [39] identified statistical procedures for classifier comparison, among them McNemar's test [40] and Cochran's Q test [41] for classifiers evaluated on a common test set; neither appears in the literature reviewed here. Both tests assume independent test cases, whereas cells drawn from the same installation are not independent observations, so the unit of analysis also needs attention.

Third, model explanation is generally performed at the global level. SHAP [42] is commonly used to produce a single feature ranking for the study area, which is then interpreted as a weight vector [26,31]. Some studies also report dependence or partial dependence analyses [22,27–29], and one reports interaction values [31]. Local attribution (i.e., decomposing the prediction for an individual location) has only been demonstrated once, by Sachit et al. [22], and at a global mapping resolution of 10 km, where a single explained pixel covers roughly 100 km^2 and does not correspond to an identifiable siting decision. Because land acquisition, permitting, and stakeholder negotiation occur site by site, explanations should be available at the same spatial scale as the siting decision. The stability of the resulting rankings across resamples is also rarely examined.

The present study addresses these limitations in Adıyaman, Elazığ, and Malatya, three neighboring provinces of eastern Türkiye. The region provides a challenging test because multiple constraints interact: much of the area is inaccessible because of ridge-and-valley topography, while the most productive agricultural land occurs in alluvial basins that are also favorably sloped, and the East Anatolian Fault Zone crosses all three provinces. Distance to active faults therefore needs to be considered jointly with terrain, land use, and infrastructure. Such interactions are where non-linear learners may outperform additive weighting and where site-level explanations become most informative.

In this context, 14 geospatial features harmonized to a common 30 m grid are used to train CatBoost, XGBoost, and LightGBM classifiers. Features are screened using association measures suited to mixed data types, a generalized variance inflation factor, and mutual information, and the final subset is selected by sequential backward selection with a tree-based learner. Positive and negative samples are drawn only from land remaining after the same restriction mask is applied, and installations are assigned to the training or test partition as whole units, so that the models learn the distinction between developed and undeveloped sites on otherwise feasible land and are evaluated on installations they have not seen. The reviewed ML studies of solar PV site selection are summarized in Table 1, which also positions the present study within the existing literature.

Table 1. ML studies of solar PV site selection, with the present study positioned against them.

Study	Region	Machine Learning Model(s)	Statistical Comparison of Models	Level of Model Explanation
Sun et al. [21]	China (national)	RF; XGBoost; MLP, fitted separately for each land cover type	None reported; models ranked by AUC	Global only (SHAP summary plot and feature importance)
Sachit et al. [22]	Global (onshore, 60° N–60° S)	RF; SVM; MLP	None reported; models ranked by accuracy, kappa, and AUC	Global and local (SHAP summary, dependence, and force plots at 10 km resolution)
Cattani [23]	Australia (91 PV plants)	RF; benchmarked against ANN, GAM,	None reported; models ranked by cross-	Global only (permutation feature

Study	Region	Machine Learning Model(s)	Statistical Comparison of Models	Level of Model Explanation
		SVM, and boosted trees	validated RMSE	importance)
Bostancı et al. [24]	Adıyaman Province, Türkiye	SVM with RBF kernel (100 models across C and γ)	Not applicable; single algorithm family	None (no post hoc attribution)
Ashraf et al. [25]	Cholistan Desert, Pakistan	RF; XGBoost; MLP	None reported; models ranked by AUC	Global only (SHAP importance)
Liao et al. [26]	China (national)	MLP for feature weighting; isotonic k-means for threshold determination	None reported; compared against conventional weighting by metric improvement	Global only (SHAP weights)
Sun et al. [27]	China (national)	RF (twelve annual models, 2011–2022)	Not applicable; single algorithm across annual models	Global and dependence (partial dependence plots)
Tao et al. [28]	Aichi Prefecture, Japan	MaxEnt (presence-only); RF (presence–absence)	None reported; models ranked by AUC	Global and dependence (response curves, partial dependence plots)
Mohlehli and Nhamo [29]	Gauteng Province, South Africa	XGBoost regressor; k-means clustering for zoning	Not applicable; single model	Global and dependence (SHAP summary and dependence plots)
Shivgan et al. [30]	India (371 sites)	RF classifier; k-means clustering; PCA	Not applicable; single classifier	Global only (SHAP attribution)
Zhou et al. [31]	China (national)	XGBoost; RF; LightGBM	None reported; models ranked by accuracy and AUC	Global and interaction (SHAP importance and interaction values)
Khazael et al. [32]	Middle East (612 PV stations)	TabNet tuned with Optuna; benchmarked against two deep-learning and seven classical models, including CatBoost	None reported; models ranked by accuracy and AUC	Global only (SHAP importance and summary plots)
This study	Adıyaman, Elazığ, and Malatya Provinces, Türkiye	CatBoost; XGBoost; LightGBM	Plant-level Cochran's Q and McNemar's tests; plant-cluster bootstrap; repeated plant-grouped cross-validation with corrected *t*-test	Global, dependence, and local at 30 m (SHAP bar, beeswarm, dependence, and force plots), with stability of the rankings across folds and models

The objective Is to develop an interpretable ML framework that identifies where ground-mounted solar PV development is likely in Adıyaman, Elazığ, and Malatya, and to assess how far its results can be relied on. The study makes four contributions. First, CatBoost, XGBoost, and LightGBM are tuned under one protocol and evaluated on held-out installations. Second, the classifiers are compared with procedures that respect the dependence between cells of the same plant: plant-level Cochran's Q and McNemar's

tests, a plant-cluster bootstrap, and repeated plant-grouped cross-validation. Third, model behavior is explained at the global, dependence, and local levels at 30 m resolution, and the stability of the SHAP rankings is examined across resampled folds and across the three models. Fourth, the suitability map is validated on held-out installations and benchmarked against a conventional GIS-MCDM map produced with the Best–Worst Method (BWM) [43,44].

2. Materials and Methods

2.1. Study Area

The research was conducted in the provinces of Elazığ, Adıyaman, and Malatya in eastern Türkiye, at the transition between the Eastern and Southeastern Anatolia regions (Figure 1). Elazığ lies in the Upper Euphrates region (38°17′–39°11′ N, 38°30′–40°21′ E), at an elevation of 1067 m, and is characterized by predominantly elevated, irregular terrain fragmented by major reservoirs and lakes. Adıyaman lies in the Middle Euphrates region (37°25′–38°11′ N, 37°00′–39°00′ E), with the Taurus Mountains to the north and extensive plains to the south. Malatya adjoins both, to the west of Elazığ and the north of Adıyaman. The provinces differ in climate and solar resources. Elazığ has a mean temperature of 13.2 °C, annual precipitation of 420.3 mm, and 7.1 h of daily sunshine, compared with 17.4 °C, 709.1 mm, and 7.8 h in Adıyaman and 13.8 °C, 383.6 mm, and 7.6 h in Malatya [45]. Annual solar irradiation is highest in Adıyaman (1884.5 kWh/m^2), followed by Malatya (1810.9 kWh/m^2) and Elazığ (1788.4 kWh/m^2) [46,47]. Licensed capacity does not follow this order: Elazığ has eight operating plants totaling approximately 250 MWe, Adıyaman fourteen plants totaling about 156 MWe, and Malatya seven plants totaling about 146 MWe [48]. These figures refer to licensed facilities, whereas the inventory of Section 2.2 records installation footprints, so it also includes unlicensed installations, and one facility may be represented by several polygons. Solar-resource availability alone therefore does not explain investment patterns, consistent with the theory–practice divide reported by Shen et al. [49].

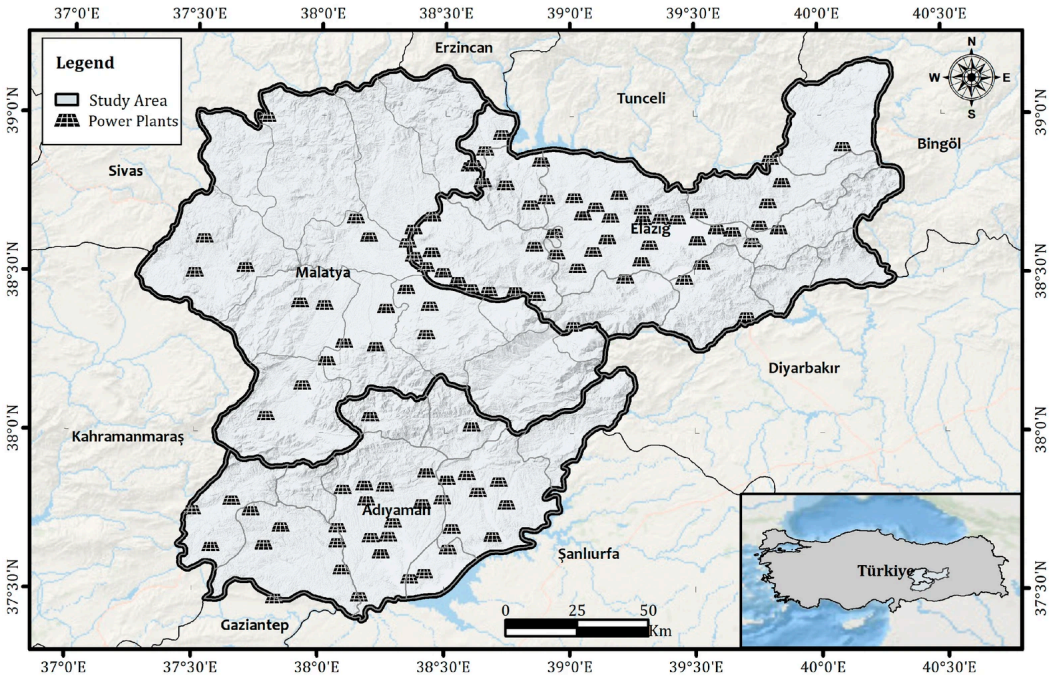

Figure 1. Study area and available solar PV plants.

The provinces also differ in their land-use, geological, and seismic constraints. In

Adıyaman, agricultural expansion associated with the Southeast Anatolia Project, restrictions under Law No. 5403 on converting agricultural land, and Mount Nemrut National Park constrain solar development [50]. In Elazığ and Malatya, major constraints include the Keban and Karakaya reservoirs, Lake Hazar, and steep terrain. All three provinces are crossed by the East Anatolian Fault Zone and were affected by the Mw 6.8 Sivrice–Elazığ earthquake in 2020 or the Mw 7.8 Pazarcık–Mw 7.6 Elbistan earthquake sequence in 2023. Together, these differences make the three provinces a useful setting for examining solar site selection under interacting resource, terrain, land-use, and seismic constraints.

2.2. Solar PV Inventory

The positive labels were obtained from operational ground-mounted solar PV plants in the three provinces. Each installation reflects a developer's decision that has already satisfied several constraints, including land availability, grid connection, and financial viability, and existing installations are therefore widely used as reference data for suitability modeling [21,28]. They document the locations where development has occurred, which is a revealed siting preference, rather than a normative assessment of the optimal locations for plant construction.

The inventory was made in three stages. First, locations of existing installations were obtained from the OpenStreetMap (OSM) database using the Overpass Turbo interface, by querying the tags "power = plant", "power = generator", and "generator:source = solar" [51]. Second, each resulting polygon was visually compared to high-resolution satellite imagery, and rooftop or inoperative records were removed. Third, installations absent from OSM were manually digitized from satellite imagery, as existing solar inventories are incomplete and satellite imagery has proven effective in filling such gaps [52]. Site boundaries were digitized as polygons, so that the true extent of each array is retained.

The digitized inventory in Elazığ and Adıyaman started with 473 ground-mounted PV polygons (median area 0.06 ha). After a visual check with satellite imagery, where a few polygons were corrected, added, or removed, the inventory was clipped to the study area and a minimum area criterion of 0.5 ha was applied, to exclude small self-consumption arrays, which follow a different siting logic from utility-scale plants; 46 polygons digitized in Malatya were then added. The last set consists of 184 polygons (138 in Elazığ–Adıyaman, and 46 in Malatya) with an area of 1185.5 ha, from 0.51 ha to 134.00 ha (median 2.84 ha), represented by 13,203 cells at 30 m resolution. Following application of the restriction mask from Section 2.3, 18 polygons (9.8%) lay fully within restricted land, and overall, 2745 cells (20.8%) were on restricted land (Table S1). The constraint that eliminated the largest number of polygons was elevation (7 polygons), followed by active faults (4), airports (3), and protected areas (2). Slope and proximity to water each removed one, and the land cover categories, substations, and the irrigated agricultural land layer did not eliminate any.

Positive cells were chosen only from the parts of each installation on the candidate surface: the positive class (1) contains 166 plants with 10,458 cells (79.2% of the inventory cells; median 27 cells, ~2.4 ha, per plant). Throughout this paper, "plant" denotes one digitized installation polygon. The inventory records the footprint of each plant, not its licensed status. In each province, the same number of cells was randomly drawn from the candidate surface and assigned to the negative class (0), resulting in a total of 20,916 labeled cells (Table S2). Thus, both classes run through the same restriction mask. They cover the same range on each masked feature and contain the same four land cover classes (Table S3). Negative cells are undeveloped background, not confirmed unsuitable

land, and some of them are probably suitable sites that have not been developed.

2.3. Conditioning Features and Restricted Zones

Two categories of spatial elements determine where a ground-mounted facility can be built. Restriction features decide whether development is allowed at a site, whereas assessment features describe which of the permitted sites are more attractive. Some factors appear in both roles under different parameterizations.

Fourteen assessment features were identified from solar PV siting research and checked against Turkish planning and environmental regulations. They are grouped into four thematic clusters, and Table 2 lists their sources and native resolutions.

Table 2. Conditioning features used as model predictors.

ID	Category	Feature	Format	Source
F1	Environmental and social	LULC	Raster/10 m	Zanaga et al. [53]
F2		Distance to protected zones	Vector	TUCBS [54]
F3		Distance to water bodies	Vector/raster 100 m	CORINE Land Cover [55]
F4		Distance to airports	Vector	OpenStreetMap [56]
F5		Distance to active fault lines	Vector	General Directorate of Mineral Research and Exploration of Türkiye [57]
F6		Distance to settlements	Vector/raster 100 m	CORINE Land Cover [55]
F7	Economic and infrastructural	Distance to main roads	Vector	OpenStreetMap [56]
F8		Distance to energy transmission lines (ETL)	Vector	OpenStreetMap [56]
F9		Distance to substations	Vector	OpenStreetMap [56]
F10	Climatic	Global horizontal irradiation	Raster/200 m	Global Solar Atlas [47]
F11		Mean annual air temperature	Raster/200 m	Global Solar Atlas [47]
F12	Topographic	Slope	Raster/30 m	Zink et al. [58]
F13		Aspect	Raster/30 m	Zink et al. [58]
F14		Elevation	Raster/30 m	Zink et al. [58]

Restricted zones are areas where a utility-scale solar facility cannot be built or can be built only under severe limitations, for environmental, social, technical, or regulatory reasons; they were removed from the study region before sampling and modeling (Table 3). The rules were rasterized and overlaid in ArcGIS Pro (version 3.5.3; Esri, Redlands, CA, USA), which was used for all spatial processing and mapping, so that any cell violating at least one rule was excluded. The mask combines the land cover classes of F1 with the irrigated agricultural land layer of the Turkish National Geographic Information System (TUCBS) [54], which is used for restriction only, and with the protected area polygons of the same system. Under Law No. 5403 [50] and the Regulation on the Protection and Use of Agricultural Land [59], solar power plants may be installed only on dry marginal agricultural land. Cropland (ESA WorldCover class 40) was there-

fore kept in the candidate surface, and irrigated agricultural land was excluded separately. The irrigated layer was obtained from the openly available TUCBS Web Map Service and digitized into vector polygons; within the study area (2,897,887 ha) it contains 649 polygons covering 12,711 ha after merging overlapping polygons, about 0.44% of the area (0.22% in Elazığ, 0.02% in Adıyaman, and 0.85% in Malatya). Irrigated parcels that are not mapped in this layer could not be excluded. Protected areas designated under Laws No. 2873 [60] and 2863 [61] were excluded as polygons without a buffer because the legislation prescribes no fixed distance and requires the opinion of the relevant authority. The remaining thresholds follow the siting literature [62,63].

Table 3. Restricted areas and thresholds applied in candidate site screening.

Feature	Threshold	Rationale
LULC	Forest, wetland, water surfaces, and built-up areas (ESA WorldCover classes 10, 50, 80, 90, 95)	Legal and ecological incompatibility
Irrigated agricultural land	Inside layer polygons, no buffer	Solar plants are allowed only on dry marginal agricultural land according to Turkish legislation
Protected areas	Inside layer polygons, no buffer	No fixed buffer in legislation; opinion of the relevant authority is required
Water bodies	<250 m	Riparian protection; flood and shoreline instability
Airports	<3000 m	Glare and approach corridors
Active faults	<500 m	Surface rupture and structural risk
Substations	<250 m	Existing installations and safety clearance
Slope	>25%	Constructability and shading geometry
Elevation	>1500 m	Accessibility, snow loading, and winter operability

Both classes were sampled after the mask had been applied because the comparison of interest is between land on which a plant was built and land on which one could have been built but was not. Sampling the whole study area would teach the classifier to separate feasible from infeasible land, a distinction the restriction rules already make.

2.3.1. Environmental and Social Features

These features reflect the environmental cost and social acceptability of land conversion. For utility-scale PV, the main environmental concern is land-use displacement rather than emissions because arrays occupy a continuous area for two to three decades (Figure 2).

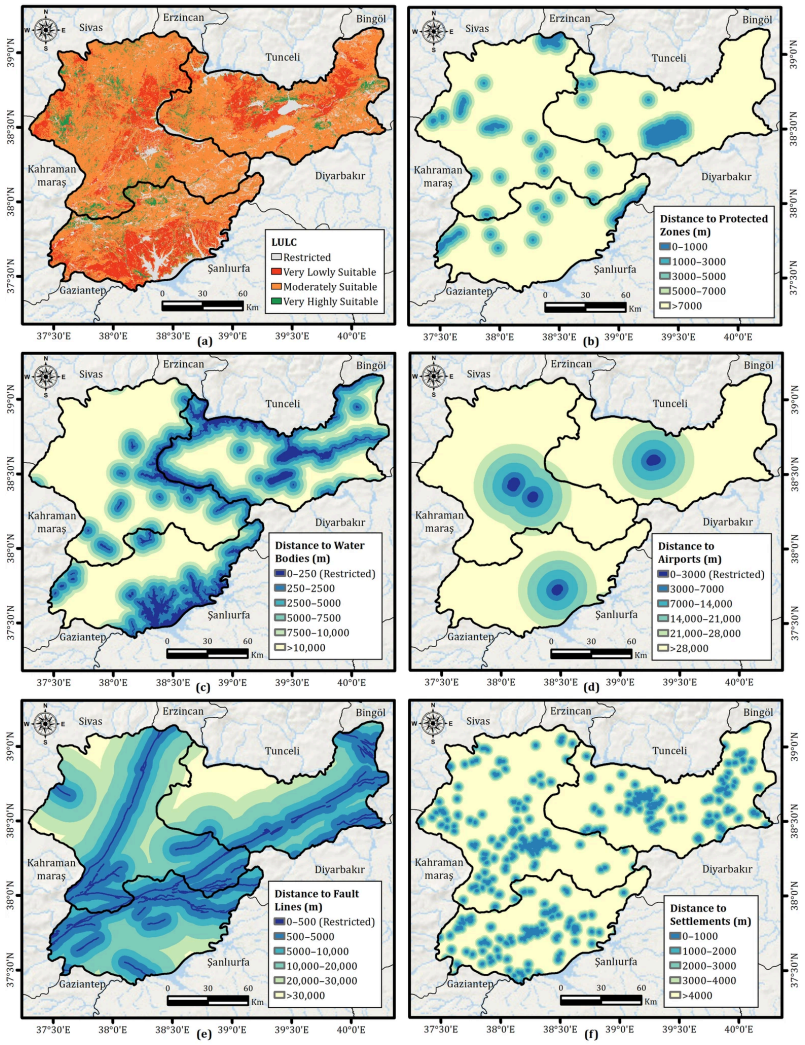

Figure 2. Spatial distribution of the environmental and social features: (**a**) LULC; (**b**) distance to protected zones; (**c**) distance to water bodies; (**d**) distance to airports; (**e**) distance to fault lines; (**f**) distance to settlements.

Land use and land cover (LULC) (F1). Land cover represents legal acceptability and opportunity cost [63], since Turkish soil conservation law reserves irrigated and designated agricultural land for farming and directs energy projects to marginal areas (Figure 2a). The layer is the ESA WorldCover 2021 product at 10 m [53], resampled to the 30 m grid. After application of the restriction mask, four classes remain in both the positive and the negative samples: cropland, grassland, shrubland, and bare or sparse vegetation. Because the product postdates most installations, the class recorded inside plant footprints may partly reflect the plants themselves (Section 4).

Distance to protected zones (F2). National parks, nature conservation areas, wildlife development areas, and designated cultural or archeological sites are delimited to prevent direct encroachment and to reduce the indirect effects of access routes, barriers, and maintenance traffic [64]; Adıyaman hosts the World Heritage-listed Mount Nemrut National Park (Figure 2b). The polygons were taken from TUCBS [54].

Distance to water bodies (F3). Utility-scale PV plants are generally kept away from reservoirs, lakes, rivers, and wetlands to protect aquatic and riparian habitats, to avoid land prone to flooding and shoreline instability, and to limit construction costs on saturated ground. The Keban, Karakaya, and Atatürk reservoirs and Lake Hazar occupy a considerable share of the study area (Figure 2c). Water bodies were taken from CORINE

Land Cover [55].

Distance to airports (F4). Specular reflection from module surfaces can create glint and glare that interfere with pilot visibility during approach and with control tower sightlines, and civil aviation rules set minimum distances around airports and approach paths (Figure 2d).

Distance to active faults (F5). The study area lies in the East Anatolian Fault Zone and was struck by the Sivrice earthquake of 24 January 2020 and the Kahramanmaraş doublet of 6 February 2023. Mounting systems, inverter stations, and transformer foundations are sensitive to ground acceleration, and surface rupture or liquefaction can damage the cabling network even when the modules remain intact [65] (Figure 2e).

Distance to settlements (F6). Proximity to settlements lowers access and connection costs and places generation close to demand, but it can also raise land values, competition with residential and agricultural uses, and objections to visual impact and construction traffic [62] (Figure 2f). Settlements were taken from CORINE Land Cover [55] because the built-up class of the 10 m land cover product also contains roads.

2.3.2. Economic and Infrastructural Features

These features describe the cost of connecting a site to existing infrastructure [66]. The cost of a road or a power line is approximately proportional to its length, whereas the energy output of the plant is independent of that distance, so accessibility may weigh more than the solar resource in areas with consistently good irradiation (Figure 3).

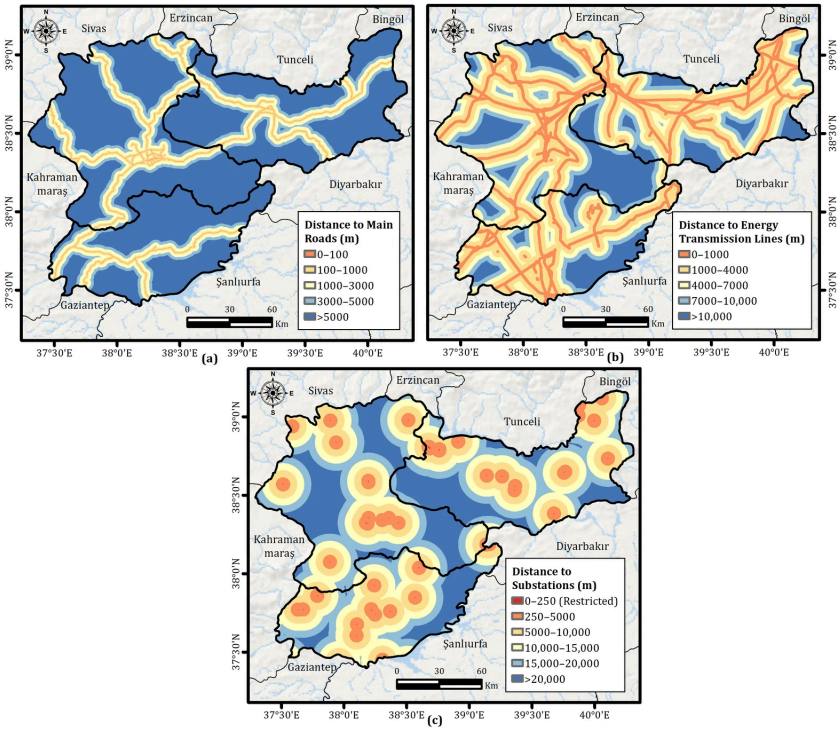

Figure 3. Spatial distribution of the economic and infrastructural features: (**a**) distance to main roads; (**b**) distance to energy transmission lines; (**c**) distance to substations.

Distance to main roads (F7). Road access affects the transport of modules, mounting frames, transformers, and heavy equipment during construction and the later costs of maintenance and module cleaning (Figure 3a).

Distance to energy transmission lines (ETL) (F8). The grid connection is often the

largest site-specific cost of a utility-scale project, with long connecting lines increasing construction costs, extending permitting through additional right-of-way negotiations, and increasing resistive losses [62] (Figure 3b). The power lines were taken from OpenStreetMap [56] and used without filtering by the voltage tag, i.e., all mapped line features were included independent of whether a voltage value was recorded. Of the 433 line features (5986 km; 34.5 kV, 154 kV, and 380 kV) in the OSM export used for the study area, 282 features had a voltage of 34.5 kV or higher, 7 features were tagged below 34.5 kV (7.7 km), and 138 features had no voltage tag (955 km). Substation busbars were not included. The feature therefore does not distinguish voltage classes, and may include both lower-voltage distribution lines as well as transmission lines.

Distance to substations (F9). Distance to a substation represents how far electricity must be carried from the plant to an existing grid connection point; connecting to an existing substation reduces the need for additional switching infrastructure and can accelerate interconnection (Figure 3c).

2.3.3. Climatic Features

The climatic features describe how much energy reaches the array and how efficiently it is converted (Figure 4).

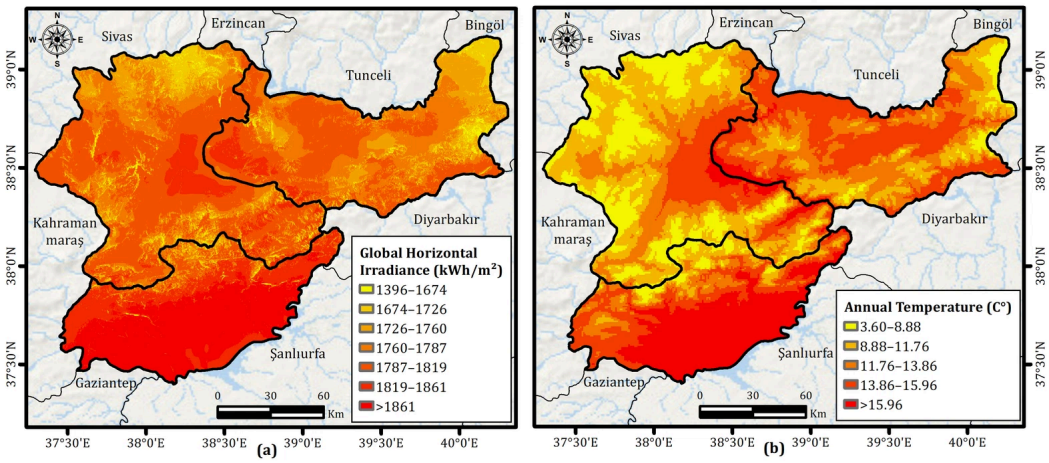

Figure 4. Spatial distribution of the climatic features: (**a**) global horizontal irradiance; (**b**) annual temperature.

Global horizontal irradiance (GHI) (F10). GHI is the key determinant of energy production from a fixed-tilt array and, consequently, of the revenue available to offset capital costs [67]. Adıyaman lies in the high-resource zone of Southeastern Anatolia, and Elazığ and Malatya lie in the mid-range of Eastern Anatolia (Figure 4a).

Mean annual air temperature (F11). For crystalline silicon modules, power output decreases as cell temperature rises, typically by 0.3–0.45% per °C above standard test conditions [68], so warm, well-irradiated areas may not outperform cooler areas with slightly lower irradiation (Figure 4b).

2.3.4. Topographic Features

Terrain influences the configuration of arrays, the cost of site preparation, and accessibility, and is expected to matter in a study area with more than 2 km of elevation change (Figure 5).

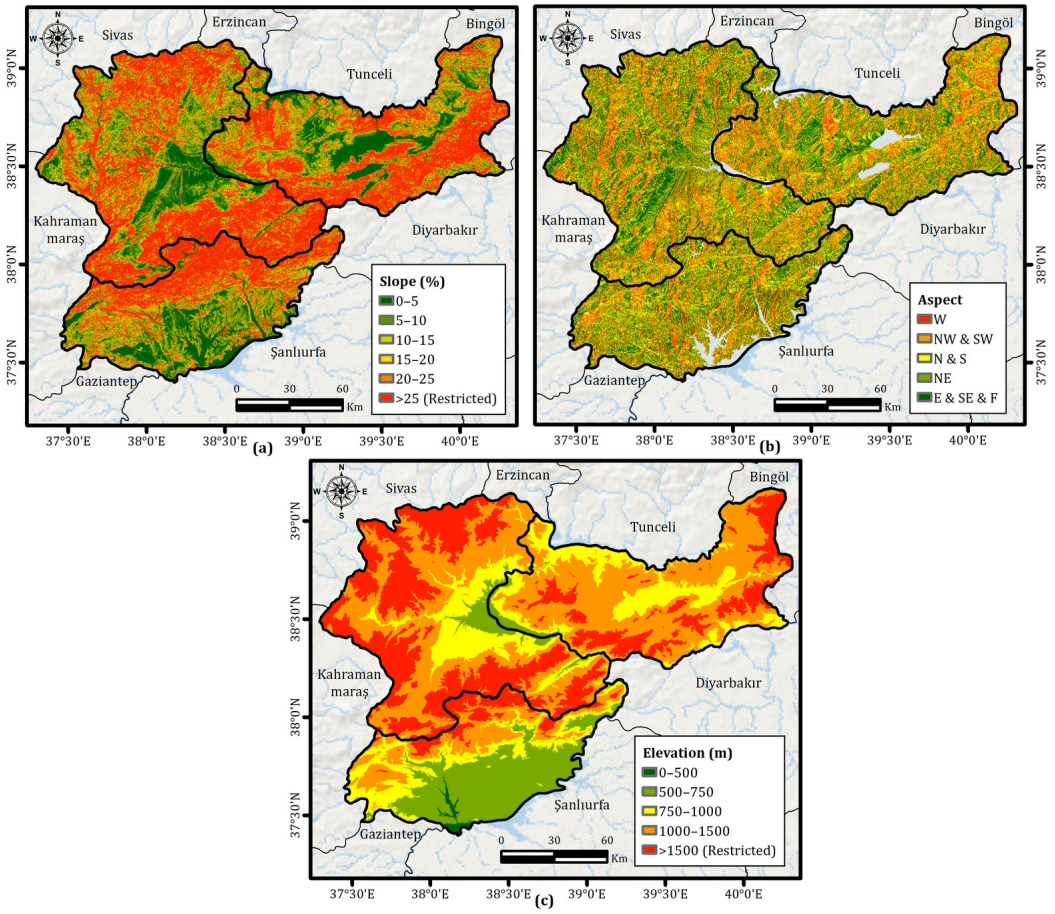

Figure 5. Spatial distribution of the topographic features: (a) slope; (b) aspect; (c) elevation.

Slope (F12). Steep terrain requires more extensive earthworks and foundation engineering, complicates road construction, increases erosion risk, and affects the row spacing needed to prevent inter-row shading [62] (Figure 5a).

Aspect (F13). At the latitudes of the study area, south-facing slopes receive substantially more solar energy than north-facing slopes, while flat terrain generally provides adequate conditions [69]. Aspect enters the analysis as eight compass classes; Figure 5b displays them in grouped form.

Elevation (F14). Higher ground is generally farther from transport routes and utility networks, receives more snow and harsher winter weather, and is more often under ecological protection, whereas lower air density slightly increases atmospheric transmissivity; on balance, low to mid-range elevations are favored (Figure 5c).

Each feature layer was rasterized to a common resolution of 30 m, and the pixel values were extracted at every sample location and joined with the class label. The resulting data frame was the single input to all three ML algorithms.

2.4. Methodology

2.4.1. Problem Definition and Workflow

Solar PV site selection is formulated here as a supervised binary classification task. The three provinces are represented as a regular grid of 30 m cells, and the analysis is confined to the cells that remain once the restriction rules of Section 2.3 have been applied. Each remaining cell s is described by a vector of fourteen conditioning features:

$$x(s) = (x_1(s), x_2(s), \ldots, x_{14}(s)) \in \mathbb{R}^{14} \quad (1)$$

Each cell is also assigned a binary label $y(s)$, which equals one if the cell lies within an operating PV facility and zero if it was drawn from the remaining candidate land (Section 2.2). The positive class records where solar development has taken place, and the negative class records where it has not, although it was permitted; the model therefore learns a revealed siting preference and not physical or normative suitability.

The aim is to learn a function f that maps the feature vector of a cell to a relative suitability score:

$$p(s) = f(x(s)) \in [0, 1] \quad (2)$$

The score $p(s)$ lies between zero and one and is kept continuous because PV suitability varies by degree: a well-irradiated site far from the grid is not ruled out but is less attractive than an equally sunny site close to a substation. Because the model is trained on a balanced 1:1 sample in which the negative class is undeveloped background, $p(s)$ orders cells by their similarity to developed sites and is not a probability of development; its calibration is assessed in Section 3.2.

The scores were then grouped into five suitability classes, from very low to very high, using the Jenks natural breaks method, which places the class boundaries so that variance is minimized within classes and maximized between them. A cell is assigned to class k according to the following:

$$C(s) = \begin{cases} k, & \text{if } T_{k-1} \le p(s) < T_k, \quad k = 1, \ldots, 4 \\ 5, & \text{if } T_4 \le p(s) \le T_5 \end{cases} \quad (3)$$

where $C(s)$ is the suitability class of cell s, $T_0 = 0$ and $T_5 = 1$ are the bounds of the score, and T_1, …, T_4 are the break values. The breaks depend on the score distribution of the study area, so the classes are ranks of relative suitability within it and not physically defined categories.

The methodology comprises five phases (Figure 6): (1) dataset construction and partitioning of the installations into training and test sets; (2) feature screening and selection; (3) model tuning, evaluation, and statistical comparison; (4) interpretation with SHAP; and (5) suitability mapping, validation on held-out installations, and comparison with a BWM-based map. All label-dependent steps use the training installations only; the test installations are used once, for the final evaluation.

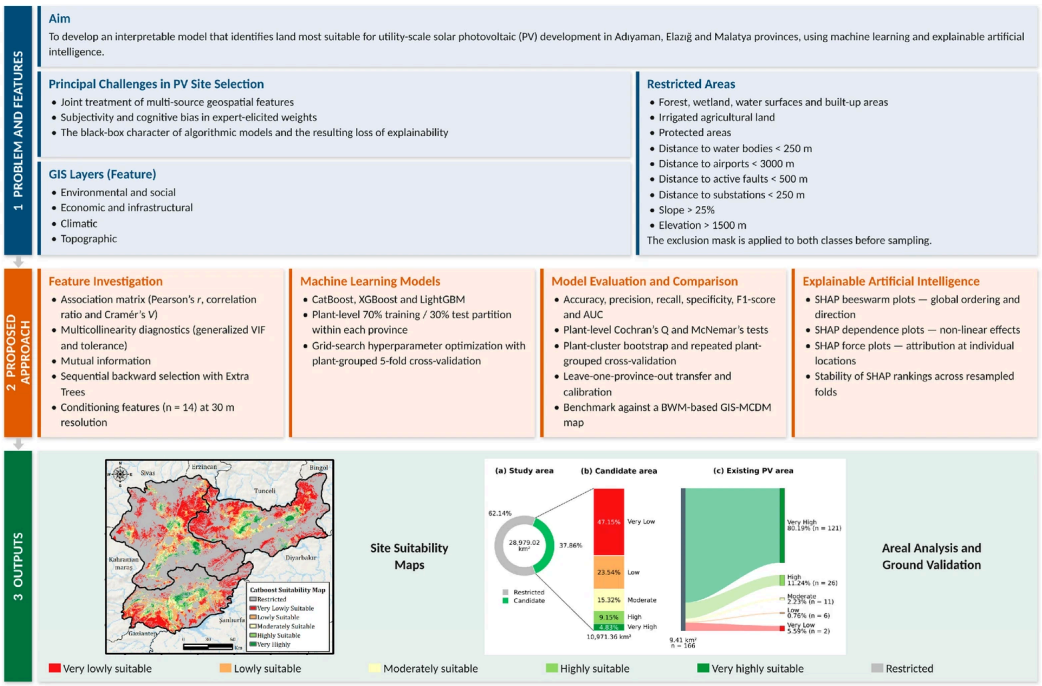

Figure 6. Proposed workflow of the study.

2.4.2. Investigation of Conditioning Features

Redundancy among the features was analyzed prior to modeling. Since the feature set contains both continuous and categorical variables, pairwise association was measured using Pearson's r for two continuous features, the correlation ratio η between a categorical (LULC, aspect) and a continuous feature, and the bias-corrected Cramér's V between the two categorical features. An absolute value above 0.7 was taken to indicate redundancy. The generalized variance inflation factor (GVIF) [70] was used to check multicollinearity, where a categorical feature is entered as a set of dummy variables and the value is rescaled by their number to be comparable with an ordinary VIF. The importance of each feature to the class label was scored using mutual information [71], computed on the raw feature values with LULC and aspect defined as discrete variables, and averaged over ten random seeds.

The final subset was selected by sequential backward selection with an Extra Trees classifier [72] (100 trees, maximum depth 4, minimum leaf size 20), a non-linear tree ensemble like the final classifiers. Starting from the features left after the redundancy screening, the feature whose removal resulted in the highest accuracy was eliminated at each step until only one feature remained. The subset with the highest mean accuracy was chosen. Each subset was scored by 3 × 5 repeated stratified plant-grouped cross-validation on the training installations, so that cells of one plant never fall on both sides of a fold. LULC and aspect entered as one-hot blocks, which were kept or removed as a whole.

2.4.3. ML Models

Gradient boosting builds an additive ensemble of shallow decision trees, each fitted to the errors left by the trees before it [73], and can therefore capture non-linearity, threshold effects, and interactions without specifying them in advance. Three implementations are compared: XGBoost [36], LightGBM [37], and CatBoost [38,74]. They share this basis but differ in how trees are grown, how splits are evaluated, and how complexity is penalized; their formulations are given in Text S1.

Installations were assigned to the training or test partition as complete polygons within each province (70/30 based on cell count), so that cells from the same plant do not appear in both partitions [75]. This led to 136 training plants (7321 cells) and 30 held-out plants (3137 cells); as the division prioritizes cells over plants, the distribution at the plant level is 82% and 18%. Neighboring polygons may belong to the same development, and polygons were not merged into sites before grouping. The distance from each held-out polygon to the nearest training polygon was therefore calculated, and the validation is also reported separately for held-out polygons with and without a training polygon nearby (Section 3.4). Negative cells (sampled independently) were partitioned at the cell level in the same ratio, leading to a test partition with 6274 out of 20,916 cells (30.0%; Table S2). The grouped hold-out eliminates the intra-plant dependence, which is the main cause of autocorrelation. Moreover, geographically blocked designs also separate training and test regions and are suitable for testing performance in unsampled areas [75,76]. Thus, the transfer between provinces was analyzed separately using a leave-one-province-out approach, where the optimized models were trained on two provinces and tested on the other.

All three classifiers were fine-tuned in the same manner, using an exhaustive grid search with 5-fold stratified cross-validation grouped by plant groups during training, evaluated by accuracy, with the random seed set to 42. Early stopping was not used, and therefore the test installations did not contribute to the tuning process. The final models were refitted on all training installations with the selected parameters. The search space for all three classifiers was the same (Table 4). Row sampling was set to 0.7. For CatBoost, the minimum leaf size was not tuned as it has no effect on symmetric trees. For XGBoost, the column sampling was 1.0 as it has an internal random sampling mechanism that produced different results on different operating systems. Its 70% row sample was thus obtained externally with a platform-independent generator. The only categorical feature retained was LULC, which was used in its original categorical form in all three classifiers, i.e., using ordered target statistics in CatBoost and categorical splits in XGBoost and LightGBM. All evaluations were done in Python (version 3.12.16; Python Software Foundation, Wilmington, DE, USA) with the scikit-learn library (version 1.9.1) [77] and the XGBoost (version 3.4.1), LightGBM (version 4.7.0; Microsoft Corporation, Redmond, WA, USA), and CatBoost (version 1.2.10; Yandex, Moscow, Russia) frameworks.

Table 4. Hyperparameter search space and selected values of the three classifiers.

Hyperparameter	Search Space	CatBoost	XGBoost	LightGBM
Learning rate	0.02, 0.05	0.05	0.02	0.05
Maximum tree depth	2, 3, 4	4	3	3
Number of trees	200, 500	200	500	500
L2 penalty	10, 80	80	10	80
Minimum leaf size	20, 80	not tuned	80	20
Column sampling	0.7, 1.0	1.0	fixed at 1.0	1.0
Row sampling	fixed	0.7	0.7	0.7
Candidates evaluated	–	48	48	96

2.4.4. Model Evaluation and Statistical Comparison

A false positive, in which unsuitable land is classified as suitable, carries planning and investment risk, whereas a false negative overlooks developable land; the two errors are not interchangeable. The models were therefore assessed with accuracy, precision, re-

call (sensitivity), specificity, and the F1-score, computed from the confusion matrix of true positives (TPs), true negatives (TNs), false positives (FPs), and false negatives (FNs) (Table 5), together with the area under the receiver operating characteristic curve (AUC).

Table 5. Evaluation metrics used to assess classification performance.

Metric	Formula	Interpretation
Accuracy	$\frac{TP+TN}{TP+TN+FP+FN}$	The proportion of all cells that are correctly classified.
Precision	$\frac{TP}{TP+FP}$	The proportion of cells predicted as suitable that are in fact suitable. It is the relevant measure when the cost of a false positive is high.
Recall	$\frac{TP}{TP+FN}$	The proportion of genuinely suitable cells that the model identifies. It indicates how much of the suitable land the model would overlook.
Specificity	$\frac{TN}{TN+FP}$	The proportion of unsuitable cells correctly excluded. It reflects how reliably the model screens out land that should not be suitable.
F1-score	$\frac{2 \times Precision \times Recall}{Precision+Recall}$	The harmonic mean of precision and recall. It is appropriate where false positives and false negatives both carry consequences.

The alteration of one or two percentage points in a single test partition may just be due to sample distribution [39], and cells originating from the same installation are not independent test cases. The classifiers were therefore evaluated in three ways that consider this dependence. At the plant level, first, Cochran's Q test [41] and pairwise McNemar's tests [40] with continuity correction and Bonferroni adjustment were used. Each held-out plant is a test case, defined as detected if at least fifty percent of its cells are classified as suitable, and each negative test cell, sampled independently, is a test case (3167 units). Because negative cells make up 99% of these units, the tests mainly compare specificity; they are reported as supplementary evidence, like the cell-level tests, and the comparison rests on the two procedures that follow. Second, a plant-cluster bootstrap with 5000 resamples, where plants were resampled as whole blocks and negative cells individually, was used to provide 95% confidence intervals for each accuracy and each pairwise difference. Third, a 5 × 5 repeated stratified plant-grouped cross-validation was performed across all 166 plants, while keeping the tuned hyperparameters fixed, and was evaluated using the corrected resampled *t*-test [78], which accounts for overlap among training sets.

Robustness was assessed by (i) refitting the tuned models with ten random seeds, (ii) analyzing the variation of cross-validation accuracy across all grid candidates, (iii) refitting with LULC one-hot encoded to ensure uniform inputs for the three classifiers, and (iv) performing a leave-one-province-out test. Calibration of the scores on the held-out installations was assessed with reliability curves, the Brier score, and the expected calibration error (ECE) over ten bins.

2.4.5. Global and Local Feature Attribution

A suitability map is useful to planners only if the reasons behind it can be understood. Conventional feature importance measures report how strongly a variable influences the model but not in which direction, nor how that influence varies between locations. The fitted models were therefore interpreted using SHAP [42], which preserves the sign of each contribution and provides explanations both for individual predictions and for the dataset as a whole [79,80]. SHAP expresses each prediction as an additive decomposition of feature contributions:

$$f(x) = \phi_0 + \sum_{i=1}^{M} \phi_i \quad (4)$$

where $f(x)$ is the model output for sample x, ϕ_0 the base value, M the number of features, and ϕ_i the SHAP value of the i-th feature, i.e., its average marginal contribution over the orderings in which the features can enter the model. The global influence of a feature is its mean absolute SHAP value across the N samples, which is expressed as follows:

$$|SHAP|_i = \frac{1}{N}\sum_{j=1}^{N} |\phi_{ij}| \quad (5)$$

where ϕ_{ij} is the SHAP value of feature i for sample j. Bar and beeswarm plots summarize the global ordering and direction of the feature effects, dependence plots show how the contribution of a feature varies across its range, and force plots decompose the prediction for an individual cell. SHAP values describe associations learned from where plants were built and are not causal effects; several predictors act as proxies for factors that were not observed, and attribution can shift between correlated features. The stability of the attributions was therefore examined: in each of the 25 folds of the repeated plant-grouped cross-validation, every model was refitted, SHAP values were computed on the validation plants, and the share of each feature in the summed mean absolute SHAP value and its rank were recorded. Agreement between rankings, across folds and across models, was measured with Kendall's τ.

2.4.6. Suitability Mapping and Ground Validation

The final CatBoost model, fitted on the training installations, was applied to all cells of the candidate surface, and the scores were classified into five classes as described in Section 2.4.1. The area of each class was calculated, and the map was validated by the share of plant cells falling in each class, reported separately for training and held-out installations; only the latter measures predictive skill. All plant cells considered lie on the candidate surface.

As a conventional benchmark, a GIS-MCDM suitability map was produced from the same eleven features retained by the selection (Section 3.1) with the BWM [43,44]. Criterion weights were elicited through a questionnaire completed by three experts in GIS-based spatial planning, PV project development and grid connection, and environmental and geotechnical assessment (Table S4). Each expert compared the four main categories, the criteria within each category, and the classes of each criterion against the best and the worst elements. Individual weights were combined by the geometric mean with equal expert weights and normalized, and global weights were obtained by multiplying category and criterion weights (Table S5). The input-based consistency ratio was below its threshold in every comparison set (Table S6), so all judgments were accepted

without revision. The weighted overlay was computed within the same restriction mask as the ML-based map, and the scores were classified into five classes by natural breaks. Agreement between the two maps was measured with Spearman's ρ and Cohen's κ.

3. Results and Discussion

3.1. Feature Investigation

The association matrix (Figure 7) shows a single pair above the 0.7 threshold: elevation and mean annual temperature (r = −0.93), which also had the highest GVIF values (8.38 and 9.98; Figure S1). Temperature was removed, leaving thirteen features. Among the remaining pairs, the strongest associations were between GHI and elevation (r = −0.67), distance to main roads and distance to ETL (0.57), and distance to main roads and distance to substations (0.51). LULC was most associated with slope (η = 0.44), all associations of aspect were 0.17 or lower, and LULC and aspect were nearly independent (V = 0.05).

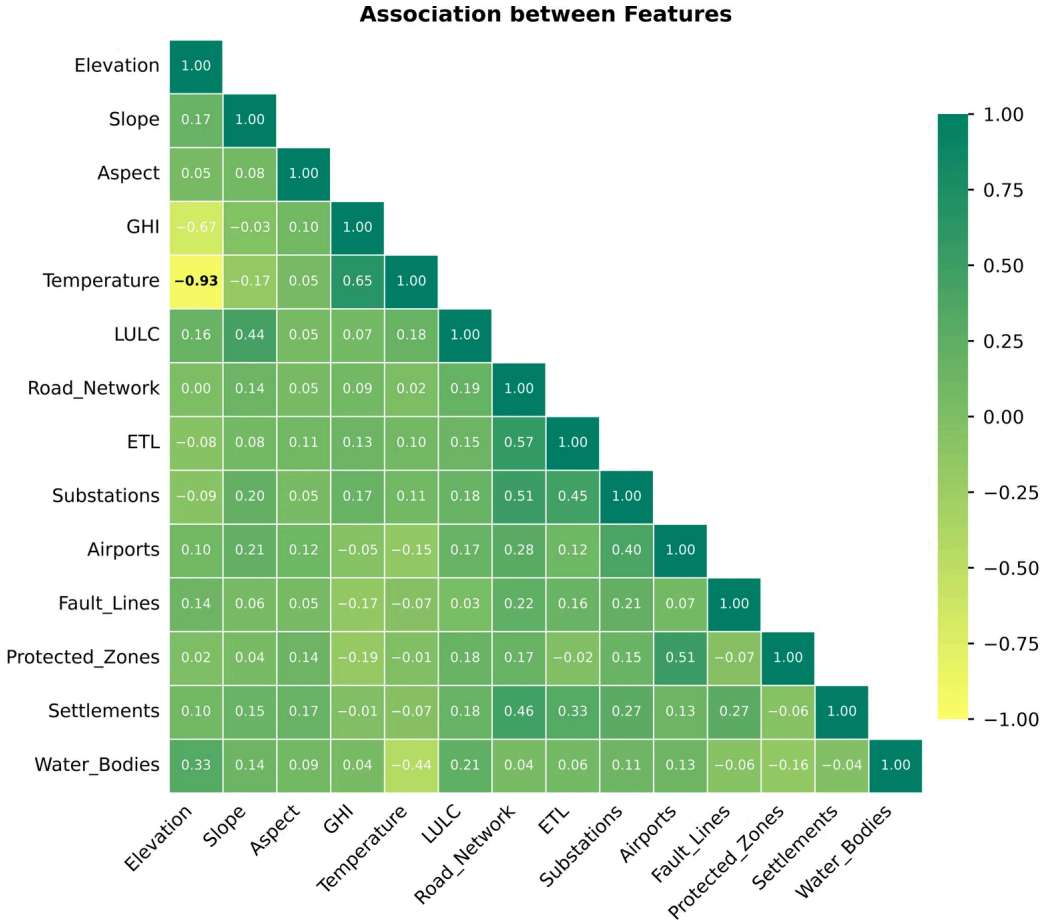

Figure 7. Association between the fourteen conditioning features: Pearson's r for pairs of continuous features, the correlation ratio η for the LULC and aspect rows, and Cramér's V for LULC–aspect.

The highest values of mutual information were for proximity to airports (0.325), GHI (0.257), proximity to substations (0.220), proximity to protected zones (0.215), and proximity to active faults (0.211) and the lowest values were for LULC (0.031), slope (0.026), and aspect (0.017) (Figure S2). Aspect and slope were removed in the backward selection, and the plant-grouped cross-validation accuracy improved from 0.706 with thirteen features to 0.735 with eleven (Figure S3 and Table S7). Low slope relevance is

consistent with its range capped by the restriction mask to 25%. The graph is flat from seven to twelve features (0.726–0.735, standard error around 0.013), meaning the exact subset is not clearly defined. Accuracy dropped to 0.658 when GHI was removed. The fault and settlement distances were maintained. The eleven features carried forward are elevation, GHI, LULC, and distances to main roads, ETL, substations, airports, active faults, protected zones, settlements, and water bodies.

3.2. ML Model Evaluation

Table 4 lists the selected hyperparameters. Across all grid candidates, the cross-validation accuracy ranged from 0.67 to 0.75 (Table S8 and Figure S4). Training accuracy (0.91–0.96) was well above the plant-grouped cross-validation accuracy (0.74–0.75) for all three classifiers (Table 6), indicating a generalization gap between plants: the models fit the training installations more closely than they predict new ones. For comparison, the shallow Extra Trees classifier used in the feature selection reached 0.735 on the same features.

Table 6. Performance of the three classifiers: accuracy on the training installations, in plant-grouped cross-validation (CV), and on the held-out installations, with the remaining metrics computed on the held-out installations.

Model	Training Accuracy	CV Accuracy	Test Accuracy	Precision	Recall	Specificity	F1-Score	AUC
CatBoost	0.9436	0.7423	0.8151	0.9060	0.7032	0.9270	0.7918	0.8352
XGBoost	0.9103	0.7463	0.7741	0.8598	0.6551	0.8932	0.7436	0.8358
LightGBM	0.9581	0.7451	0.7631	0.8994	0.5926	0.9337	0.7145	0.8542

On the 30 held-out installations (6274 cells), CatBoost reached an accuracy of 0.815, XGBoost 0.774, and LightGBM 0.763, with AUC values of 0.835, 0.836, and 0.854, respectively (Table 6 and Figure 8). All three classifiers were conservative: specificity (0.89–0.93) and precision (0.86–0.91) were high, whereas recall was 0.70 for CatBoost, 0.66 for XGBoost, and 0.59 for LightGBM, so most errors were plant cells classified as unsuitable. At the plant level, CatBoost and XGBoost each detected 21 of the 30 held-out plants and LightGBM 19 (Table S9). These values come from a single partition of 30 plants and depend in part on how close the held-out plants are to training plants (Section 3.4).

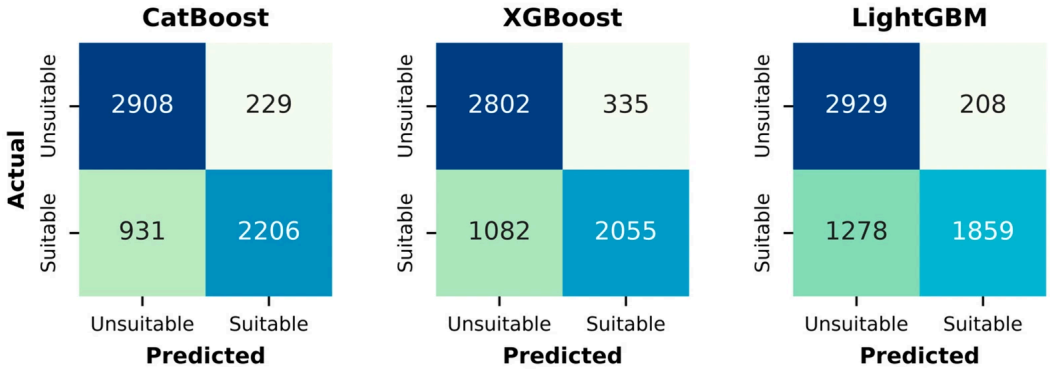

Figure 8. Confusion matrices of the CatBoost, XGBoost, and LightGBM classifiers on the held-out installations (n = 6274 cells).

Cell-level tests rejected equal performance for every pair (Cochran's Q = 223.0, $p < 0.001$; Figure S5), but they treat the cells of one plant as independent. The picture changes when plants are used as the unit of analysis (Figure 9). Cochran's Q remained

significant (Q = 119.4, $p < 0.001$), but CatBoost and LightGBM were not different (adjusted $p = 0.25$) and XGBoost differed from both (adjusted $p < 0.001$; Table S10). XGBoost accepted more undeveloped cells as suitable (specificity 0.893 versus 0.927 and 0.934), but did not detect fewer plants. As negative cells form 99% of the test units, these tests reflect specificity more than the detection of plants, and the bootstrap and the repeated cross-validation are taken as the main evidence. The plant-cluster bootstrap gave a 95% confidence interval for CatBoost accuracy of 0.65–0.92; the interval for the pairwise difference included zero for CatBoost–XGBoost (0.041; −0.010 to 0.111) and XGBoost–LightGBM (0.011; −0.040 to 0.066), but narrowly excluded zero for CatBoost–LightGBM (0.052; 0.002 to 0.096) (Table S11). In the plant-grouped cross-validation over all 166 plants, the accuracies were 0.717 ± 0.037, 0.716 ± 0.049, and 0.716 ± 0.034, and the corrected t-test showed no difference ($p \geq 0.95$; Figure S6, Table S12 and S13). The repeated cross-validation accuracy of about 0.72 is therefore the more representative estimate of performance on unseen plants, and the hold-out values of 0.76–0.82 lie within the bootstrap interval.

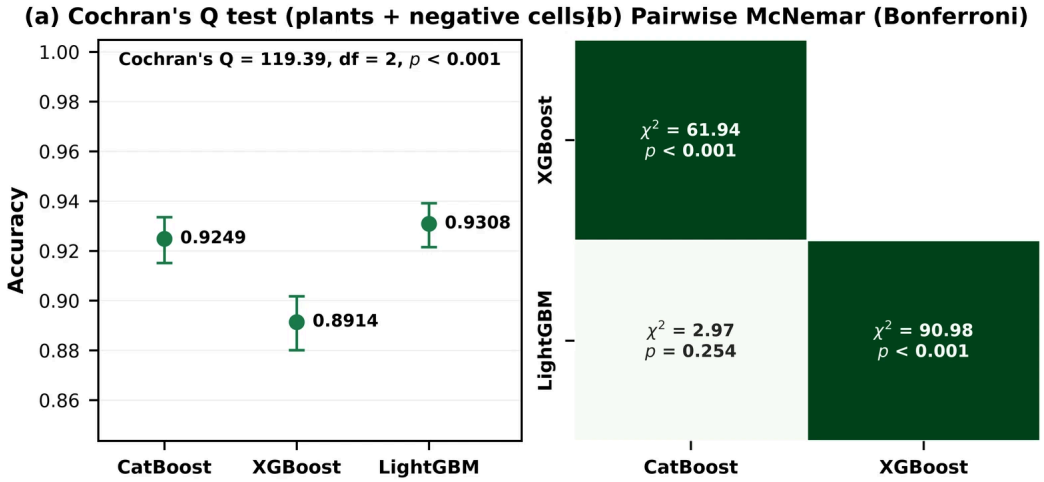

Figure 9. Statistical comparison of the classifiers at the plant level (30 held-out plants and 3137 negative test cells): (**a**) accuracy with 95% confidence intervals and the result of Cochran's Q test; (**b**) pairwise McNemar's tests with Bonferroni-adjusted *p*-values.

The following validations confirm the same interpretation. CatBoost had a test accuracy of 0.822 ± 0.008, XGBoost 0.762 ± 0.007, and LightGBM 0.771 ± 0.009 across ten random seeds (Table S14). The one-hot encoding of LULC did not affect the accuracy of LightGBM (0.763) and changed the accuracy of CatBoost and XGBoost by 0.002 and 0.004, respectively (Table S15), suggesting that the differences among the classifiers are not due to the categorical encoding. When the unit of analysis is the plants, no classifier is consistently better than the others. CatBoost was chosen for interpretation and mapping as it had the highest accuracy on the held-out installations and consistently performed well across the ten seeds. The SHAP results for XGBoost and LightGBM are shown in the Supplementary Materials.

Performance was uneven between provinces (Table S16). CatBoost reached accuracies of 0.895 in Elazığ and 0.906 in Malatya but 0.592 in Adıyaman, where a single plant holding 571 of the 863 positive test cells was missed by all three classifiers. In the leave-one-province-out test, accuracy fell to 0.57–0.68 and AUC to 0.57–0.81 (Table S17). The models therefore transfer poorly to a province on which they were not trained, and the results should be applied only within the sampled region.

The scores were not well calibrated: on the held-out installations, the Brier score was

0.174–0.186 and the ECE 0.157–0.185 (Figure S7 and Table S18). The model output is therefore treated as a relative suitability score and the classes derived from it as ranks.

3.3. ML Model Interpretation

3.3.1. Global Explanations

Figure 10 shows the global explanation of the CatBoost model on the held-out cells. Distance to substations had the largest mean absolute SHAP value (0.63), followed by distance to airports (0.40), distance to main roads (0.35), LULC (0.27) and distance to settlements (0.21); the remaining six features lay between 0.12 and 0.17. In the beeswarm plot, short distances to substations, main roads, and settlements push the score upward, and long distances push it downward.

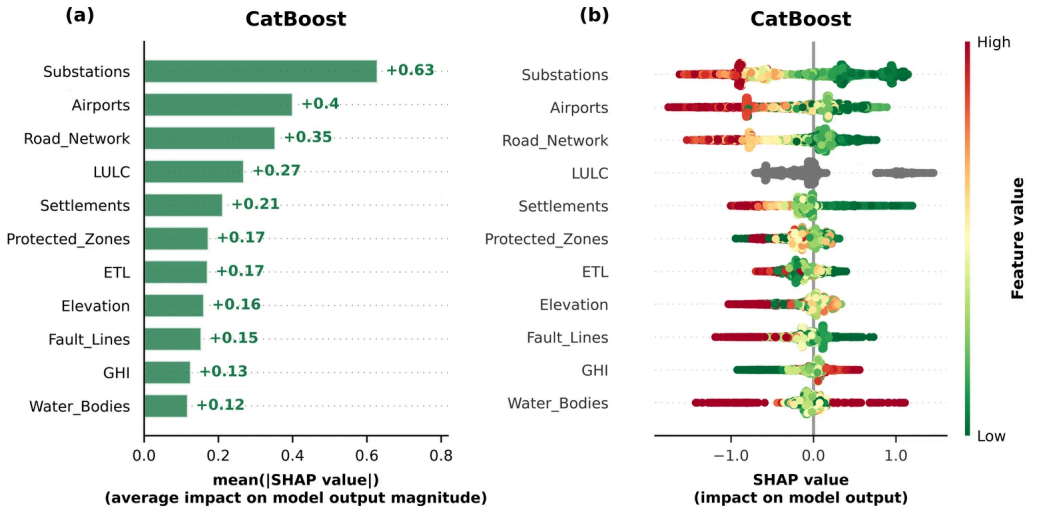

Figure 10. Global explanation of the CatBoost model on the held-out cells: (**a**) mean absolute SHAP values of the eleven features and (**b**) SHAP beeswarm plot. In the beeswarm plot, each point represents one cell, colors denote the value of the feature from low (green) to high (red), and gray points denote the categorical feature LULC.

The ranking was stable at the top and less stable below (Figure 11 and Table S19). Across the 25 resampled folds, distance to substations ranked first in every fold for CatBoost and LightGBM and in 24 of 25 for XGBoost, with a mean share of 21.6% of the summed attribution in CatBoost, and distance to airports was second on average in all three models (14.1% in CatBoost). The relative ranks of the other features varied largely, from third to eleventh for distance to protected zones and distance to active faults, for example. The Kendall's τ ranged from 0.60 to 0.65 within folds of the same model and 0.78 to 0.89 between models (Table S20). Therefore, only the leading features are interpreted, and the order of the others is not considered a finding. The global and dependence plots of XGBoost and LightGBM are given in Figure S8–S13.

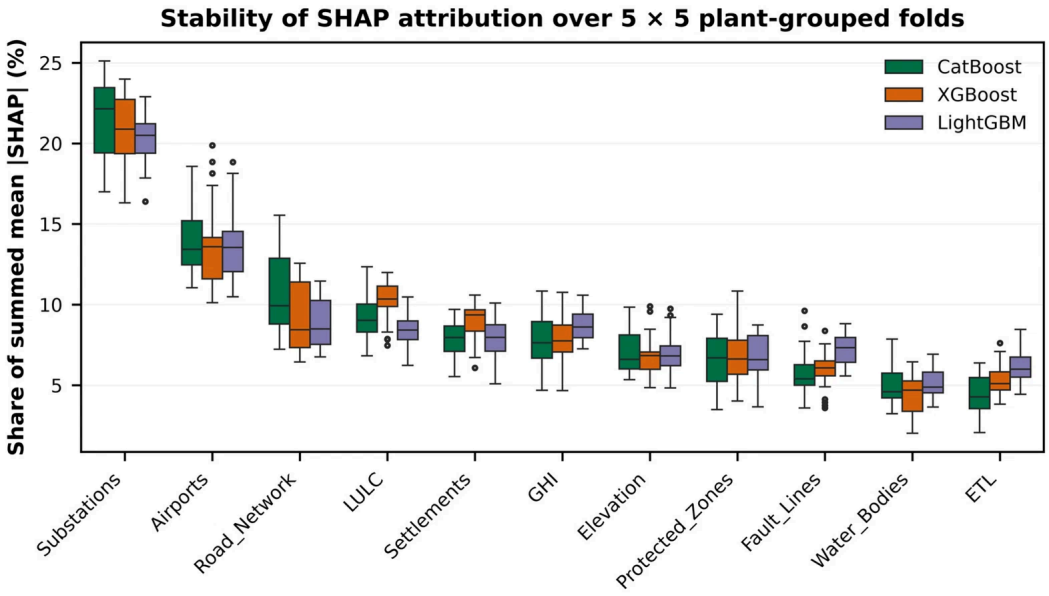

Figure 11. Stability of the SHAP attributions over the 25 folds of the repeated plant-grouped cross-validation: share of each feature in the summed mean absolute SHAP value for CatBoost, XGBoost, and LightGBM.

GHI accounted for 7.7% of the CatBoost attribution and on average was sixth out of eleven (third to ninth by fold). This modest role depends on three properties of the data. The candidate surface spans a relatively narrow range of irradiance (coefficient of variation 2.5%; Table S21 and Figure S14), the GHI layer is at a native resolution of 200 m, coarser than the 30 m sampling grid, and the majority of the variation is between provinces: the province explains 74% of the GHI variance in the candidate surface, while the within-province coefficient of variation is only ~1.2–1.3% (1.28% in Elazığ, 1.17% in Adıyaman, and 1.30% in Malatya). Furthermore, the second climatic feature, temperature, was excluded due to collinearity with the elevation. Nonetheless, GHI was informative: it had the second-highest mutual information and its removal reduced the accuracy of the backward selection (Section 3.1). The result does not imply that the solar resource is unimportant but that it varies too little within each province to distinguish between developed and undeveloped land. Expert-weighted studies ranked GHI first in Türkiye [15,16], but data-driven studies find the pattern seen here: vegetation and distance to the grid were the strongest predictors of PV location in China [21,27], technical and economic factors dominated a global model [22], and in the Middle East, a TabNet–SHAP analysis ranked solar radiation low and infrastructure proximity high, which its authors attributed to the uniformly high irradiance of the region [32].

Two of the distance features should be interpreted as proxies. The airports are located close to the provincial centers, so distance to airports reflects proximity to the main urban and infrastructure cores, not an effect of aviation. Distance to ETL had the smallest share (4.4%; ranks 8–11): the layer does not differentiate voltage classes (Section 2.3.2) and is related to distance to main roads and to substations (Figure 7), so grid access is represented in the model mainly by distance to substations.

3.3.2. Dependence Analysis

Dependence plots of the six most important features are shown in Figure 12, and dependence plots of the remaining features are shown in Figure S15. The impact of distance to substations is positive up to about 5 km, close to zero between 5 and 10 km, and negative beyond 10 km. Distance to main roads follows a similar pattern, with positive

effects up to about 5 km and strongly negative effects (around −1.0 to −1.5) beyond about 12 km. Distance to settlements has a strongly positive effect within about 1 km and a negative effect beyond about 7 km. Distance to airports is mixed within 20 km, close to zero till about 50 km, and negative beyond. For LULC, the largest positive contribution is from bare or sparsely vegetated land (+0.8 to +1.5), the negative contribution is from cropland (−0.2 to −0.7), and grassland and shrubland values are near zero or slightly negative. The contribution of distance to protected zones is weak and irregular.

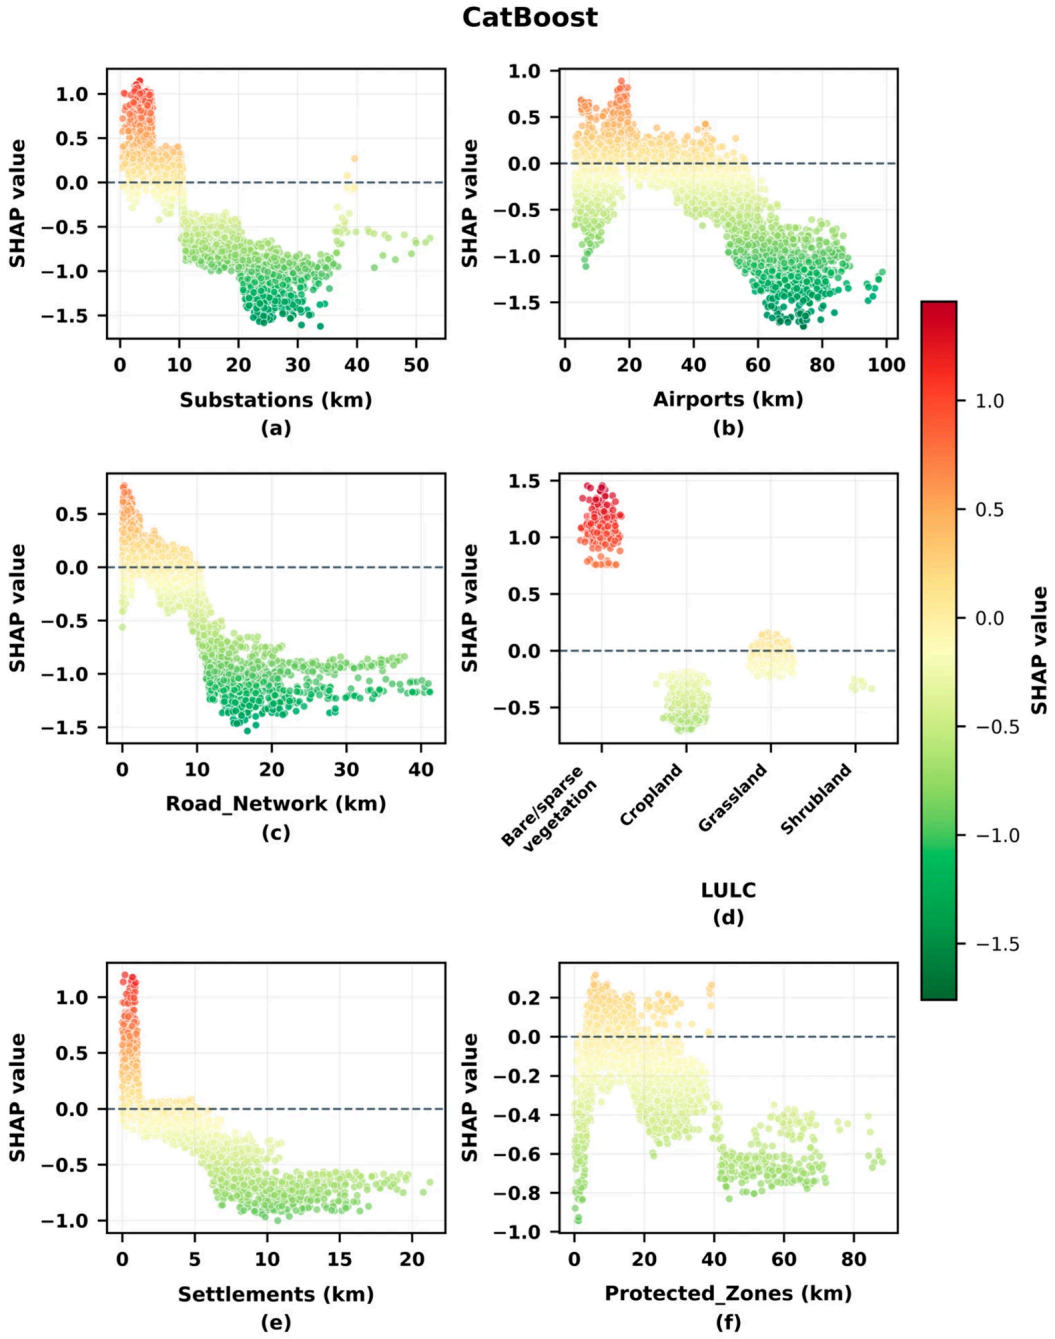

Figure 12. SHAP dependence plots of the CatBoost model for (**a**) distance to substations, (**b**) distance to airports, (**c**) distance to main roads, (**d**) LULC, (**e**) distance to settlements, and (**f**) distance to protected zones.

For the remaining features, GHI contributes negatively below approximately 1780 kWh/m^2, has a contribution near zero between 1780 and 1870 kWh/m^2, and contributes

positively above. The upper transition is near the provincial value of Adıyaman (1884.5 kWh/m^2) and above those of Malatya (1810.9 kWh/m^2) and Elazığ (1788.4 kWh/m^2) [46,47], so some of the GHI effect may distinguish the resource regimes of provinces rather than sites within them. Distance to ETL varies irregularly with values within ±0.6 and no threshold. These are ranges where the learned association flips sign, not engineering thresholds, and their position is dependent on the sample and the completeness of the OSM layers.

The positive contribution of bare or sparsely vegetated land should be regarded with caution. This class accounts for 18% of the positive cells and 5% of the negative cells. Since pre-construction land cover was not used, it is not possible to exclude that some of this difference is a reflection of how the land cover product labels ground within plant footprints (Section 4).

3.3.3. Local Explanations

Figure 13 decomposes the scores of two correctly classified held-out cells with median scores in their class. The base value is about −0.15 on the log-odds scale. For the positive cell (score 0.675), distance to airports (7.39 km), distance to substations (4.39 km), elevation (1156 m), distance to main roads (3.95 km), and distance to water bodies (11.41 km) raise the score, while its land cover, cropland, lowers it. For the negative cell (score 0.086), distance to settlements (10.68 km), distance to substations (16.66 km), GHI (1772.9 kWh/m^2), distance to protected zones (17.84 km), and distance to ETL (1.42 km) lower the score, and only the short distance to a main road (1.81 km) raises it. These decompositions show which features led to a given score at the 30 m scale at which parcels are assessed. Like the global results, these are associations and should be verified with site information before use.

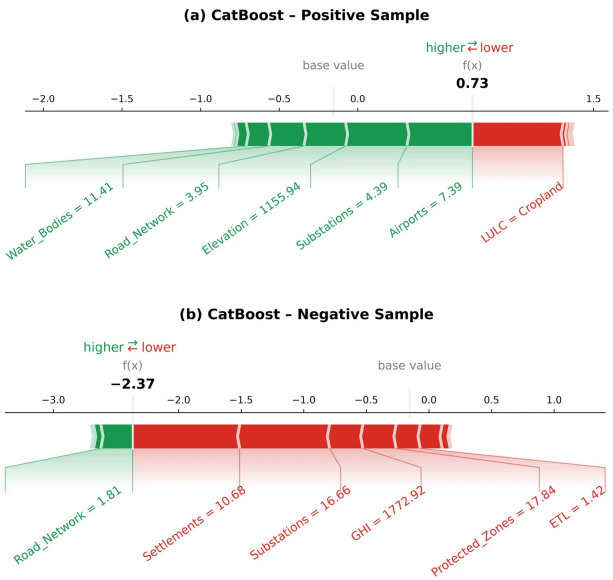

Figure 13. SHAP force plots of the CatBoost model for (**a**) a positive sample and (**b**) a negative sample. Values are expressed in log-odds; green segments increase and red segments decrease the output relative to the base value, and f(x) denotes the final model output.

3.4. Site-Based Suitability Assessments

Of the 28,979.02 km^2 study area, 62.14% is restricted and 37.86% (10,971.36 km^2) forms the candidate surface (Figure 14a). The candidate surface is skewed toward the lower classes (Figure 14b): 47.15% is very low, 23.54% low, 15.32% moderate, 9.15% high,

and 4.83% very high, with class breaks at scores of 0.097, 0.224, 0.390, and 0.606 (Table S22). The two upper classes together cover 14.0% of the candidate surface.

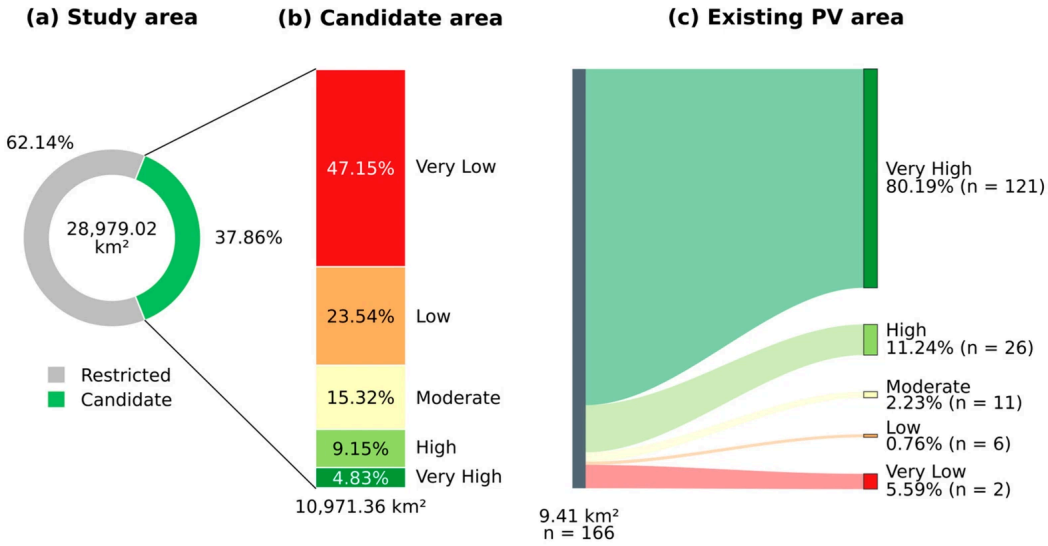

Figure 14. Areal distribution of the suitability assessment: (**a**) share of the study area removed by the restriction rules and remaining as candidate surface; (**b**) distribution of the candidate surface among the five suitability classes; (**c**) distribution of the existing PV area on the candidate surface among the suitability classes (n = number of plants).

The CatBoost map (Figure 15a) shows that high and very high suitability is concentrated in compact zones around the provincial centers and along the main infrastructure corridors. The mountainous and remote parts of the candidate surface are graded low or very low. The BWM-based map (Figure 15b) shows a different structure, with linear bands of high suitability along the transmission lines.

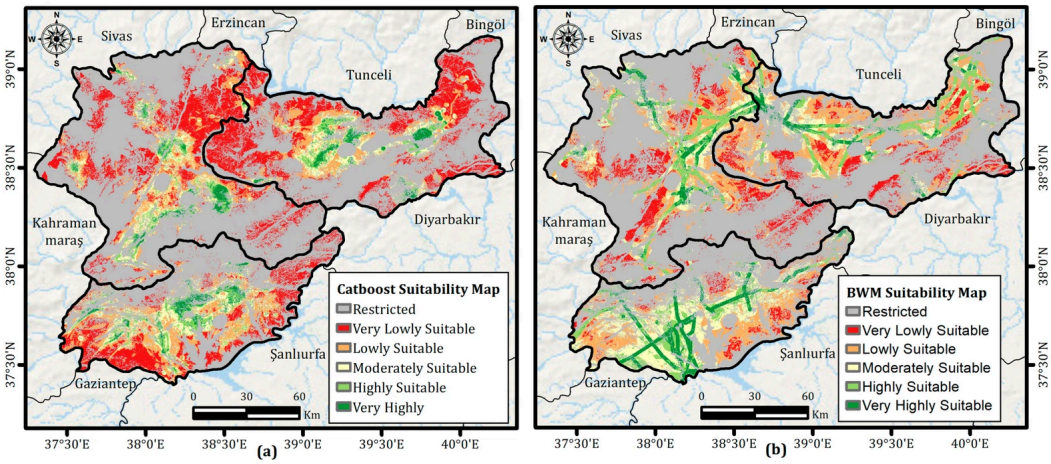

Figure 15. Solar PV suitability maps of Adıyaman, Elazığ, and Malatya: (**a**) CatBoost model; (**b**) BWM benchmark. Restricted areas defined in Table 3 are shown in gray.

The CatBoost map was overlaid with existing plants (Figure 16 and Table S23). In total, 76.1% of the cells of the 30 held-out plants are in the high or very high classes (55.2% in very high), 4.1% in moderate, 1.1% in low, and 18.7% in very low, while these two classes cover 14.0% of the candidate surface. This share is weighted by cells. With each plant weighted equally, the mean share of a held-out plant's cells in the high or very high class is 73.7%, and 22 of the 30 plants have at least half of their cells in these classes (

Table S24). The fit, not predictive skill, is reflected in the share of 98.0% for the training plants and 91.4% for all plants (Figure 14c). The held-out share by province is 94.6% in Elazığ, 91.3% in Malatya, and 31.2% in Adıyaman (Figure S16). The lower value is because one plant in Adıyaman (571 cells, 66% of that province's held-out cells) is entirely in the very low class; without that plant, 93.1% of held-out cells are high or very high. Proximity to training plants also matters. Twelve of the 30 held-out plants (513 cells) lie within 500 m of a training plant, and nine of them (404 cells) within 250 m; merging polygons closer than 250 m would give 130 sites, six of which contain both training and held-out polygons. For these twelve plants the cell-weighted share is 99.2% and CatBoost classified 92.8% of their cells as suitable, whereas for the 18 plants farther than 500 m from any training plant, the share is 71.6% (56.6% with plants weighted equally) and 65.9% of the cells were classified as suitable (Table S25). The validation is therefore more favorable for plants that adjoin training plants. These shares refer to plant cells on the candidate surface. The 2745 inventory cells on restricted land (20.8% of the inventory) are not on the mapped surface.

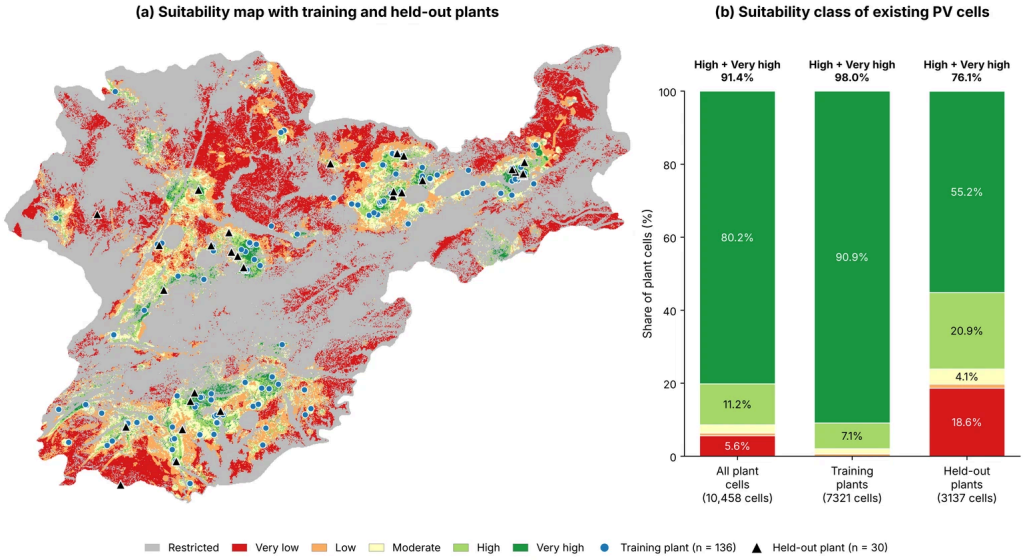

Figure 16. Ground validation of the CatBoost suitability map: (**a**) suitability classes with training and held-out plants; (**b**) distribution of plant cells among the suitability classes for all, training, and held-out plants.

The BWM-derived map showed a weak correspondence with the ML-derived map (Spearman's $\rho = 0.27$; Cohen's $\kappa = 0.07$). It also scored lower for existing plants: 41.9% of plant cells were in its high and very high classes compared to 91.4% for the ML-based map (27.3% compared to 76.1% for held-out plants). The BWM surface was dominated by proximity to ETL, whose nearest class had the largest single class weight (0.13; Table S5). Conversely, the learned model gave most weight to distance to substations. The expert weighting describes where plants would be expected under stated priorities, and the learned score where they have been constructed. This comparison is not a test of physical suitability, which would require independent yield and cost information (Section 4).

4. Limitations and Future Research Directions

Several limitations of the data and design point to further work:

First, the positive class consists of installations that developers have actually built,

so their locations reflect not only physical conditions but also the incentive programs, licensing requirements, and interconnection rules in force when each project was permitted. The score is a revealed preference within a particular historical context, not a normative or physical statement of suitability, and it was not validated against independent yield, cost, or feasibility information, which is not available at 30 m for the region. Negative cells are undeveloped background, and the sensitivity of the results to the negative-sampling strategy was not tested. Re-training with updated inventories or fitting annual models as in Sun et al. [27] would show how the siting frontier moves with policy.

Second, the land cover layer is a post-construction product. Restricting both classes to the candidate surface removes land cover classes that occur in one class only, but pre-construction land cover with a change-detection step was not used, so the contribution of LULC, in particular of bare or sparsely vegetated land, may partly reflect the plants themselves.

Third, several factors that developers consider could not be characterized at 30 m: the hosting capacity of individual substations, land value and ownership, cadastral parcel boundaries, and the routing of new connection lines. The infrastructure layers come from OSM; their completeness in the three provinces could not be verified independently, about 14% of the mapped line length has no voltage tag, and the ETL feature may include distribution lines. The distance ranges reported in Section 3.3.2 depend on these layers.

Fourth, generalization is limited. Accuracy on unseen plants varied between provinces, and the leave-one-province-out test indicated weak transfer, suggesting feature effects are specific to this region and its regulatory context. It lacks an external validation area or a formal out-of-distribution analysis, and the hold-out is grouped by polygon and not by site. Held-out polygons close to a training polygon were predicted more completely than the others (Section 3.4), so performance on isolated new sites is likely to be lower than the hold-out values suggest [75,76]. The scores are not calibrated, and the five classes are relative to the study area.

Fifth, the thresholds of Table 3 are applied as a hard mask before sampling, and their uncertainty was not propagated to the map since this would require rebuilding the candidate surface and the sample for each threshold set. In addition, the framework is empirical and contains no mechanistic model of how a site would behave, so shortlisted parcels require site-level engineering assessment of structural and geotechnical response, for which physics-informed methods such as the energy approach of Samaniego et al. [81] may be used.

Finally, the model is a binary cell classifier that ranks land by its similarity to developed sites; it says nothing about the capacity, yield, or cost of a plant. A capacity assessment would require an energy-yield model, a levelized cost estimate including the connection line, and a spatial allocation procedure constrained by the minimum plant size.

5. Conclusions

This study developed an interpretable ML framework for solar PV site selection and applied it to Adıyaman, Elazığ, and Malatya in eastern Türkiye. Cells from 166 ground-mounted plants and an equal number of undeveloped cells were sampled under one restriction mask, plants were held out as whole units, and fourteen conditioning features on a 30 m grid were reduced to eleven by association analysis, mutual information, and tree-based backward selection.

CatBoost, XGBoost, and LightGBM, tuned under one protocol, reached an accuracy of about 0.72 in repeated plant-grouped cross-validation and 0.76–0.82 on a single hold-

out of 30 plants. Cell-level tests suggested clear differences between them, but a plant-cluster bootstrap and repeated plant-grouped cross-validation showed no consistent difference; CatBoost was carried forward for its accuracy on the held-out plants and its stability across random seeds. The methodological lesson is that treating the cells of one installation as independent overstates the differences between such models. Performance was uneven between provinces, transfer to an unseen province was weak, and the scores were not calibrated, so the output is a relative suitability score valid within the sampled region.

SHAP ranked distance to substations first in every resampled fold and distance to airports, a proxy for proximity to the provincial centers, second, while the order of the remaining features was not stable. GHI had a modest share; this finding is conditional on the narrow irradiance range of the candidate surface, the 200 m resolution of the GHI layer, and the fact that most GHI variation lies between provinces, and it does not imply that the solar resource is unimportant for PV generation. The attributions are associations, not causal effects.

The suitability map assigns 14.0% of the 10,971 km^2 candidate surface to the high and very high classes, which contain 76.1% of the cells of held-out plants (73.7% with plants weighted equally), against 27.3% for a BWM-based map built from the same features. Within the study region, plants have been built close to substations, main roads, and settlements, and the learned score reproduces this pattern on installations it has not seen.

The framework applies wherever an inventory of existing plants and geospatial data are available. It describes how development has been placed rather than prescribing where it should go and is best used alongside the planning process, followed by site-scale engineering assessment of shortlisted parcels.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/app162311768/s1, Text S1: formulations of the gradient boosting algorithms; Figure S1–S16: feature screening, hyperparameter sensitivity, statistical comparison, calibration, SHAP results of XGBoost and LightGBM, GHI discriminability, and ground validation by province; Table S1–S25: inventory and sample composition, BWM inputs, feature selection, robustness checks, statistical tests, SHAP stability, GHI statistics, suitability classes, ground validation, and proximity of held-out to training plants.

Author Contributions: Conceptualization, B.K. and A.B.; methodology, B.K. and A.B.; software, A.B.; validation, B.K. and A.B.; formal analysis, B.K. and A.B.; investigation, M.G.; resources, B.K.; data curation, Y.F.U.; writing—original draft preparation, B.K. and A.B.; writing—review and editing, B.K. and A.B.; visualization, A.B. and Y.F.U.; supervision, B.K.; funding acquisition, B.K. and M.G. All authors have read and agreed to the published version of the manuscript.

Funding: This study was supported by the Scientific Research Projects Coordination Unit of Firat University (Project Number: MF.26.72).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data and code supporting the findings of this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AHP	Analytic hierarchy process
ANN	Artificial neural network
AUC	Area under the receiver operating characteristic curve
BWM	Best–Worst Method
CatBoost	Categorical boosting
COP	Conference of the Parties
CV	Cross-validation
ECE	Expected calibration error
ETL	Energy transmission lines
FN	False negative
FP	False positive
GAM	Generalized additive model
GHI	Global horizontal irradiance
GIS	Geographic information system
GVIF	Generalized variance inflation factor
LightGBM	Light gradient boosting machine
LULC	Land use and land cover
MCDM	Multi criteria decision making
ML	Machine learning
MLP	Multilayer perceptron
OSM	OpenStreetMap
PV	Photovoltaic
RF	Random forest
SHAP	Shapley Additive exPlanations
SVM	Support vector machine
TN	True negative
TP	True positive
TUCBS	Turkish National Geographic Information System
VIF	Variance inflation factor
XGBoost	Extreme gradient boosting

References

1. UNFCCC. Outcome of the First Global Stocktake. In *Proceedings of the Decision 1/CMA.5, Conference of the Parties Serving as the Meeting of the Parties to the Paris Agreement, Fifth Session, Dubai, United Arab Emirates, 30 November–13 December 2023*; UNFCCC: Bonn, Germany, 2024. Available online: https://unfccc.int/documents (accessed on 10 September 2026). [Google Scholar]
2. UNFCCC. *The Road to Antalya: Türkiye to Host COP31 in 2026*; UNFCCC: Bonn, Germany, 2026. Available online: https://unfccc.int/cop31/the-road-to-antalya (accessed on 10 September 2026). [Google Scholar]
3. Republic of Türkiye. *The Second Nationally Determined Contribution (NDC 3.0)*; UNFCCC: Bonn, Germany, 2025. [Google Scholar]
4. Republic of Türkiye Ministry of Energy and Natural Resources. *Türkiye National Energy Plan*; Republic of Türkiye Ministry of Energy and Natural Resources: Ankara, Türkiye, 2022. Available online: https://enerji.gov.tr/announcements-detail?id=20331 (accessed on 5 September 2026).
5. Republic of Türkiye Ministry of Energy and Natural Resources. Elektrik [Electricity]. 2026. Available online: https://enerji.gov.tr/bilgi-merkezi-enerji-elektrik (accessed on 15 September 2026). [Google Scholar]
6. Jiang, F.; Zhu, Z.; Li, S.; Zhang, J. Navigating sustainable energy transitions and environmental constraints in PV siting. *Sustainability* **2026**, *18*, 3915. [Google Scholar] [CrossRef]

7. Wiggins, G.; Cary, M.; Chen, S. Optimal parcel-level photovoltaic and agrivoltaic site selection under physical and socioeconomic constraints. *Renew. Energy* **2026**, *257*, 124822. [Google Scholar] [CrossRef]
8. Arán Carrión, J.; Espín Estrella, A.; Aznar Dols, F.; Zamorano Toro, M.; Rodríguez, M.; Ramos Ridao, A. Environmental decision-support systems for evaluating the carrying capacity of land areas: Optimal site selection for grid-connected photovoltaic power plants. *Renew. Sustain. Energy Rev.* **2008**, *12*, 2358–2380. [Google Scholar] [CrossRef]
9. Sánchez-Lozano, J.M.; García-Cascales, M.S.; Lamata, M.T. Comparative TOPSIS-ELECTRE TRI methods for optimal sites for photovoltaic solar farms. Case study in Spain. *J. Clean. Prod.* **2016**, *127*, 387–398. [Google Scholar] [CrossRef]
10. Hooshangi, N.; Mahdizadeh Gharakhanlou, N.; Ghaffari Razin, S.R. Evaluation of potential sites in Iran to localize solar farms using a GIS-based Fermatean Fuzzy TOPSIS. *J. Clean. Prod.* **2023**, *384*, 135481. [Google Scholar] [CrossRef]
11. Gökler, S.H. Enhanced site selection for solar power plants utilizing the geographic information system and Pythagorean fuzzy analytical hierarchy process method. *Sustain. Energy Grids Netw.* **2025**, *44*, 101929. [Google Scholar] [CrossRef]
12. Ilhan, E.; Samadi, F.; Karimi, H. Solar farm suitability mapping for the contiguous United States using the fuzzy analytic hierarchy process and geospatial analysis. *Renew. Energy* **2026**, *260*, 125229. [Google Scholar] [CrossRef]
13. Deb, S.S.; Das, B.; Sharma, S. A spatial decision-support framework for solar photovoltaic planning in humid and forest-dominated regions: A case study of Tripura, India. *Renew. Energy* **2026**, *271*, 126016. [Google Scholar] [CrossRef]
14. Uyan, M. Optimal site selection for solar power plants using multi-criteria evaluation: A case study from the Ayranci region in Karaman, Turkey. *Clean Technol. Environ. Policy* **2017**, *19*, 2231–2244. [Google Scholar] [CrossRef]
15. Tercan, E.; Eymen, A.; Urfalı, T.; Saracoglu, B.O. A sustainable framework for spatial planning of photovoltaic solar farms using GIS and multi-criteria assessment approach in Central Anatolia, Turkey. *Land Use Policy* **2021**, *102*, 105272. [Google Scholar] [CrossRef]
16. Günen, M.A. A comprehensive framework based on GIS-AHP for the installation of solar PV farms in Kahramanmaraş, Turkey. *Renew. Energy* **2021**, *178*, 212–225. [Google Scholar] [CrossRef]
17. Türk, S.; Koç, A.; Şahin, G. Multi-criteria of PV solar site selection problem using GIS-intuitionistic fuzzy based approach in Erzurum province/Turkey. *Sci. Rep.* **2021**, *11*, 5034. [Google Scholar] [CrossRef]
18. Demir, A.; Dinçer, A.E.; Yılmaz, K. A novel method for the site selection of large-scale PV farms by using AHP and GIS: A case study in İzmir, Türkiye. *Sol. Energy* **2023**, *259*, 235–245. [Google Scholar] [CrossRef]
19. Akay, S.S. Investigating land suitability for PV farm and existing sites using a multi-criteria decision approach in Gaziantep, Türkiye. *Appl. Sci.* **2025**, *15*, 2441. [Google Scholar] [CrossRef]
20. Ranjgar, B.; Niccolai, A.; Leva, S. Solar PV site selection using GIS-MCDM: A comprehensive approach with spatially explicit collinearity and sensitivity analyses. *Renew. Energy* **2026**, *270*, 125873. [Google Scholar] [CrossRef]
21. Sun, Y.; Zhu, D.; Li, Y.; Wang, R.; Ma, R. Spatial modelling the location choice of large-scale solar photovoltaic power plants: Application of interpretable machine learning techniques and the national inventory. *Energy Convers. Manag.* **2023**, *289*, 117198. [Google Scholar] [CrossRef]
22. Sachit, M.S.; Shafri, H.Z.M.; Abdullah, A.F.; Rafie, A.S.M.; Gibril, M.B.A. Global spatial suitability mapping of wind and solar systems using an explainable AI-based approach. *ISPRS Int. J. Geo-Inf.* **2022**, *11*, 422. [Google Scholar] [CrossRef]
23. Cattani, G. Combining data envelopment analysis and Random Forest for selecting optimal locations of solar PV plants. *Energy AI* **2023**, *11*, 100222. [Google Scholar] [CrossRef]
24. Bostancı, B.; Kaynak, T.; Çapkurt, Z. Location selection for solar power plants by using support vector machines: Adıyaman province, Turkey. *Renew. Energy* **2024**, *229*, 120689. [Google Scholar] [CrossRef]
25. Ashraf, H.A.; Li, J.; Li, Z.; Sohail, A.; Ahmed, R.; Butt, M.H.; Ullah, H. Geographic information system and machine learning approach for solar photovoltaic site selection: A case study in Pakistan. *Processes* **2025**, *13*, 981. [Google Scholar] [CrossRef]
26. Liao, Y.; Miao, S.; Fan, W.; Liu, X. A novel hybrid fuzzy comprehensive evaluation and machine learning framework for solar PV suitability mapping in China. *Remote Sens.* **2025**, *17*, 2070. [Google Scholar] [CrossRef]
27. Sun, Y.; Li, Y.; Wang, R.; Ma, R. Dynamic geospatial modeling of solar energy expansion potential in China: Implications for national-level optimization of solar photovoltaic plant layouts. *Energy* **2025**, *333*, 137433. [Google Scholar] [CrossRef]
28. Tao, L.; Hayashi, K.; Gyeltshen, S.; Shimoyama, Y. Spatial assessment of utility-scale solar photovoltaic siting potential using machine learning approaches: A case study in Aichi prefecture, Japan. *Appl. Energy* **2025**, *383*, 125329. [Google Scholar] [CrossRef]
29. Mohlehli, M.G.; Nhamo, G. Explainable geo-informatics for spatial solar suitability analysis in Gauteng Province, South Africa. *Appl. Geomat.* **2026**, *18*, 81. [Google Scholar] [CrossRef]

30. Shivgan, Y.; Bhirud, C.; Jamgade, A.; Kulkarni, A.; Patil, R.; Kaur, G.; Bassir, D.; Tejani, G.G. SaurKshetra: A curated dataset and ML-based classification for solar energy site selection. *Data Brief.* **2026**, *67*, 112906. [Google Scholar] [CrossRef]
31. Zhou, J.; Xiao, J.; Ren, Y.; Chen, L.; Liu, Y.; Lai, J.; He, P.; He, Z.; Yin, D.; Yang, H. Assessment of suitability for photovoltaic plant siting and planning of power transmission corridors in China. *J. Clean. Prod.* **2026**, *552*, 147888. [Google Scholar] [CrossRef]
32. Khazael, S.M.; Abdul Maulud, K.N.; Gibril, M.B.A.; Sachit, M.S.; Karim, O.A. Enhancing solar PV suitability mapping in the Middle East using an optimized deep learning framework. *Alex. Eng. J.* **2025**, *129*, 553–571. [Google Scholar] [CrossRef]
33. Asadi, M.; Pourhossein, K. Neural network-based modelling of wind/solar farm siting: A case study of East-Azerbaijan. *Int. J. Sustain. Energy* **2021**, *40*, 616–637. [Google Scholar] [CrossRef]
34. Nadizadeh Shorabeh, S.; Neysani Samany, N.; Minaei, F.; Karimi Firozjaei, H.; Homaee, M.; Darvishi Boloorani, A. A decision model based on decision tree and particle swarm optimization algorithms to identify optimal locations for solar power plants construction in Iran. *Renew. Energy* **2022**, *187*, 56–67. [Google Scholar] [CrossRef]
35. Okeke, C.J.; Ebuara, D.E.; Edet, J.U.; Blanchard, R.E.; Wilson, J. A scalable machine learning framework for hybrid solar–wind site suitability screening. *Energy AI* **2026**, *25*, 100782. [Google Scholar] [CrossRef]
36. Chen, T.; Guestrin, C. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*; ACM: New York, NY, USA, 2016; pp. 785–794. [Google Scholar]
37. Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; Liu, T.-Y. LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems 30*; Curran Associates: Red Hook, NY, USA, 2017. [Google Scholar]
38. Prokhorenkova, L.; Gusev, G.; Vorobev, A.; Dorogush, A.V.; Gulin, A. CatBoost: Unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems 31*; Curran Associates: Red Hook, NY, USA, 2018. [Google Scholar]
39. Dietterich, T.G. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Comput.* **1998**, *10*, 1895–1923. [Google Scholar] [CrossRef]
40. McNemar, Q. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika* **1947**, *12*, 153–157. [Google Scholar] [CrossRef]
41. Cochran, W.G. The comparison of percentages in matched samples. *Biometrika* **1950**, *37*, 256–266. [Google Scholar] [CrossRef]
42. Lundberg, S.M.; Lee, S.-I. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30*; Curran Associates: Red Hook, NY, USA, 2017. [Google Scholar]
43. Rezaei, J. Best-worst multi-criteria decision-making method. *Omega* **2015**, *53*, 49–57. [Google Scholar] [CrossRef]
44. Rezaei, J. Best-worst multi-criteria decision-making method: Some properties and a linear model. *Omega* **2016**, *64*, 126–130. [Google Scholar] [CrossRef]
45. Turkish State Meteorological Service (MGM). Resmî İklim İstatistikleri: İllere Ait Mevsim Normalleri (1991–2020) [Official Climate Statistics: Seasonal Normals of Provinces (1991–2020)]. Available online: https://www.mgm.gov.tr/veridegerlendirme/il-ve-ilceler-istatistik.aspx (accessed on 2 September 2026).
46. Republic of Türkiye Ministry of Energy and Natural Resources. Güneş Enerjisi Potansiyel Atlası (GEPA) [Solar Energy Potential Atlas]. Available online: https://enerji.gov.tr/bilgi-merkezi-enerji-gunes (accessed on 2 September 2026).
47. Solargis. *Global Solar Atlas 2.13*; World Bank Group, Energy Sector Management Assistance Program (ESMAP): Washington, DC, USA, 2026. Available online: https://globalsolaratlas.info (accessed on 2 September 2026). [Google Scholar]
48. Enerji Atlası. Güneş Enerji Santralleri: Elazığ, Adıyaman ve Malatya [Solar Power Plants: Elazığ, Adıyaman and Malatya]. Available online: https://www.enerjiatlasi.com (accessed on 10 September 2026).
49. Shen, J.; Zheng, F.; Chen, T.; Ma, Y.; Deng, W.; Bellotti, A.G.; Tesema, F.B.; Prasad, D. Why theoretical photovoltaic suitability does not fully predict deployment: Interpretable deep learning of urban-scale siting across Chinese cities. *Appl. Energy* **2026**, *421*, 128228. [Google Scholar] [CrossRef]
50. Toprak Koruma ve Arazi Kullanımı Kanunu [Soil Protection and Land Use Law]. Law No. 5403. *Official Gazette of the Republic of Türkiye, No. 25880*, 19 July 2005. Available online: https://www.mevzuat.gov.tr/mevzuat?MevzuatNo=5403&MevzuatTur=1&MevzuatTertip=5 (accessed on 10 September 2026).
51. Dunnett, S.; Sorichetta, A.; Taylor, G.; Eigenbrod, F. Harmonised global datasets of wind and solar farm locations and power. *Sci. Data* **2020**, *7*, 130. [Google Scholar] [CrossRef]
52. Kruitwagen, L.; Story, K.T.; Friedrich, J.; Byers, L.; Skillman, S.; Hepburn, C. A global inventory of photovoltaic solar energy generating units. *Nature* **2021**, *598*, 604–610. [Google Scholar] [CrossRef]

53. Zanaga, D.; Van De Kerchove, R.; Daems, D.; De Keersmaecker, W.; Brockmann, C.; Kirches, G.; Wevers, J.; Cartus, O.; Santoro, M.; Fritz, S.; et al. *ESA WorldCover 10 m 2021 v200*; Zenodo: Geneva, Switzerland, 2022. [Google Scholar] [CrossRef]
54. General Directorate of Geographic Information Systems, Ministry of Environment, Urbanisation and Climate Change of the Republic of Türkiye. Türkiye Ulusal Coğrafi Bilgi Sistemi (TUCBS) [Turkish National Geographic Information System]. Available online: https://tucbs.gov.tr (accessed on 2 November 2026).
55. Copernicus Land Monitoring Service. CORINE Land Cover 2018 (Vector/Raster 100 m), Europe, 6-Yearly. Available online: https://data.europa.eu/data/datasets/-0b4b2fac-adf1-43a1-a829-70f02bf0c0e5-?locale=en (accessed on 2 November 2026).
56. Overpass Turbo: A Web-Based Data Mining Tool for OpenStreetMap. Available online: https://overpass-turbo.eu/ (accessed on 1 August 2026).
57. Emre, Ö.; Duman, T.Y.; Özalp, S.; Elmacı, H.; Olgun, Ş.; Şaroğlu, F. *Açıklamalı Türkiye Diri Fay Haritası, Ölçek 1:1,250,000 [Active Fault Map of Turkey with Explanatory Text, Scale 1:1,250,000]*; Special Publication Series 30; General Directorate of Mineral Research and Exploration (MTA): Ankara, Türkiye, 2013. [Google Scholar]
58. Zink, M.; Moreira, A.; Bachmann, M.; Rizzoli, P.; Fritz, T.; Hajnsek, I.; Krieger, G.; Wessel, B. The global TanDEM-X DEM—A unique data set. In Proceedings of the 2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS), Fort Worth, TX, USA, 23–28 July 2017; pp. 906–909. [Google Scholar]
59. Tarım Arazilerinin Korunması ve Kullanılması Hakkında Yönetmelik [Regulation on the Protection and Use of Agricultural Land]. *Official Gazette of the Republic of Türkiye*, No. 33214; 4 April 2026.
60. Milli Parklar Kanunu [National Parks Law]. Law No. 2873. *Official Gazette of the Republic of Türkiye*, No. 18132; 11 August 1983.
61. Kültür ve Tabiat Varlıklarını Koruma Kanunu [Law on the Conservation of Cultural and Natural Property]. Law No. 2863. *Official Gazette of the Republic of Türkiye*, No. 18113; 23 July 1983.
62. Uguz, Y.F.; Bilgili, A.; Uzar, M. Toward improved siting of wind–solar hybrid farms: A novel framework integrating multi-criteria decision making, machine learning–driven feature selection, and spatial clustering. *Appl. Energy* **2026**, *411*, 127609. [Google Scholar] [CrossRef]
63. Ryberg, D.S.; Robinius, M.; Stolten, D. Evaluating land eligibility constraints of renewable energy sources in Europe. *Energies* **2018**, *11*, 1246. [Google Scholar] [CrossRef]
64. Rehbein, J.A.; Watson, J.E.M.; Lane, J.L.; Sonter, L.J.; Venter, O.; Atkinson, S.C.; Allan, J.R. Renewable energy development threatens many globally important biodiversity areas. *Glob. Change Biol.* **2020**, *26*, 3040–3051. [Google Scholar] [CrossRef]
65. Swan, S.; Hadjian, A.H. *The 1986 North Palm Springs Earthquake: Effects on Power Facilities*; EPRI NP-5607, Research Project 2848; Electric Power Research Institute: Palo Alto, CA, USA, 1988. [Google Scholar]
66. Al Garni, H.Z.; Awasthi, A. Solar PV power plant site selection using a GIS-AHP based approach with application in Saudi Arabia. *Appl. Energy* **2017**, *206*, 1225–1240. [Google Scholar] [CrossRef]
67. Zoghi, M.; Ehsani, A.H.; Sadat, M.; Amiri, M.J.; Karimi, S. Optimization solar site selection by fuzzy logic model and weighted linear combination method in arid and semi-arid region: A case study Isfahan-IRAN. *Renew. Sustain. Energy Rev.* **2017**, *68*, 986–996. [Google Scholar] [CrossRef]
68. Dehra, H. An investigation on energy performance assessment of a photovoltaic solar wall under buoyancy-induced and fan-assisted ventilation system. *Appl. Energy* **2017**, *191*, 55–74. [Google Scholar] [CrossRef]
69. Mondol, J.D.; Yohanis, Y.G.; Norton, B. The impact of array inclination and orientation on the performance of a grid-connected photovoltaic system. *Renew. Energy* **2007**, *32*, 118–140. [Google Scholar] [CrossRef]
70. Fox, J.; Monette, G. Generalized collinearity diagnostics. *J. Am. Stat. Assoc.* **1992**, *87*, 178–183. [Google Scholar] [CrossRef]
71. Bilgili, A.; Arda, T.; Kilic, B. Explainability in wind farm planning: A machine learning framework for automatic site selection of wind farms. *Energy Convers. Manag.* **2024**, *309*, 118441. [Google Scholar] [CrossRef]
72. Geurts, P.; Ernst, D.; Wehenkel, L. Extremely randomized trees. *Mach. Learn.* **2006**, *63*, 3–42. [Google Scholar] [CrossRef]
73. Friedman, J.H. Greedy function approximation: A gradient boosting machine. *Ann. Stat.* **2001**, *29*, 1189–1232. [Google Scholar] [CrossRef]
74. Dorogush, A.V.; Ershov, V.; Gulin, A. CatBoost: Gradient boosting with categorical features support. *arXiv* **2018**, arXiv:1810.11363. [Google Scholar]
75. Roberts, D.R.; Bahn, V.; Ciuti, S.; Boyce, M.S.; Elith, J.; Guillera-Arroita, G.; Hauenstein, S.; Lahoz-Monfort, J.J.; Schröder, B.; Thuiller, W.; et al. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography* **2017**, *40*, 913–929. [Google Scholar] [CrossRef]

76. Wadoux, A.M.J.-C.; Heuvelink, G.B.M.; de Bruin, S.; Brus, D.J. Spatial cross-validation is not the right way to evaluate map accuracy. *Ecol. Model.* **2021**, *457*, 109692. [Google Scholar] [CrossRef]
77. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830. [Google Scholar]
78. Nadeau, C.; Bengio, Y. Inference for the generalization error. *Mach. Learn.* **2003**, *52*, 239–281. [Google Scholar] [CrossRef]
79. Shapley, L.S. A value for n-person games. In *Contributions to the Theory of Games II*; Kuhn, H.W., Tucker, A.W., Eds.; Princeton University Press: Princeton, NJ, USA, 1953; pp. 307–317. [Google Scholar]
80. Aydin, H.E.; Iban, M.C. Predicting and analyzing flood susceptibility using boosting-based ensemble machine learning algorithms with SHapley Additive exPlanations. *Nat. Hazards* **2023**, *116*, 2957–2991. [Google Scholar] [CrossRef]
81. Samaniego, E.; Anitescu, C.; Goswami, S.; Nguyen-Thanh, V.M.; Guo, H.; Hamdia, K.; Zhuang, X.; Rabczuk, T. An energy approach to the solution of partial differential equations in computational mechanics via machine learning: Concepts, implementation and applications. *Comput. Methods Appl. Mech. Eng.* **2020**, *362*, 112790. [Google Scholar] [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.