Global utility-scale photovoltaics are structurally displaced from electricity demand centres

Boxiong Su[1,2], Mingquan Wu[1,2,*], Jie Liu[1,4], Feng Yan[3], Zheng Niu [1,2], Fang Chen [1,2],
Meng Wang[4], Wang Li[1,2], Zongru Li[2,5], Hanguang Yu[6]

1. International Research Center of Big Data for Sustainable Development Goals, State Key Laboratory of Remote Sensing and Digital Earth, Aerospace Information Research Institute, Chinese Academy of Sciences, Beijing 100094, China;
2. University of Chinese Academy of Sciences, Beijing 100049, China;
3. Foreign Environmental Cooperation Center, Ministry of Ecology and Environment of China, Beijing 100812, China;
4. International Cooperation Office, Aerospace Information Research Institute, Chinese Academy of Sciences, Beijing 100094, China
5. State Key Laboratory of Geographic Information Science and Technology, Institute of Geographic Sciences and Natural Resources Research, Chinese Academy of Sciences, Beijing, 100101, China
6. School of Artificial Intelligence, China University of Geosciences (Beijing), Beijing 100083, China

Corresponding author: Mingquan Wu (wumq@aircas.ac.cn)

Supplementary Information

Supplementary Methods

S1. Global utility-scale PV mapping

S1.1 Sample construction and Sentinel-2 input features

Model inputs were derived from Sentinel-2 multispectral surface reflectance, which was used to capture spectral differences between PV installations and surrounding land surfaces across the visible, near-infrared and shortwave-infrared ranges. Ten bands were selected as input features: B2, B3, B4, B5, B6, B7, B8, B8A, B11 and B12.

Positive samples were obtained from an existing global PV distribution dataset^18. Negative samples were generated by random sampling outside a 1-km buffer surrounding known PV samples, so as to reduce interference from roads, bare land, industrial sites and other highly reflective backgrounds during model training. For all sample locations, multispectral reflectance values were extracted in Google Earth Engine at the corresponding pixel positions, and compiled into a binary labelled sample set.

To improve training stability and reduce the influence of scale differences among spectral bands, all input features were standardized using statistics calculated from the training set, and the same normalization parameters were then applied consistently to the validation and test sets. The samples were split by stratified sampling into training, validation and test subsets to preserve the class proportions of PV and non-PV samples. Detailed sample statistics and descriptions of the input bands are provided in Supplementary Table 2.

S1.2 Model training and global-scale inference

PV samples were extracted from the existing global PV distribution dataset, whereas non-PV samples were obtained by random sampling outside a 1-km buffer surrounding PV samples. All samples were processed in Google Earth Engine, where spectral values were extracted according to site location and exported to form a labelled dataset with a binary field, label_mode. To reduce scale differences among spectral bands and improve convergence stability, all input features were standardized to zero mean and unit variance. The normalization parameters were calculated from the training set and then applied to the validation and test sets. The full sample set was divided by stratified sampling into training (70%), validation (15%) and test (15%) subsets to ensure consistent PV/non-PV class proportions across all partitions.

The classifier was a lightweight three-layer fully connected neural network (Fully Connected Deep Neural Network, FC-DNN). The input layer takes a 10-dimensional feature vector. Two hidden layers, with 32 and 64 neurons and ReLU activation, are used for nonlinear mapping, followed by a sigmoid-activated output layer that produces a probability value between 0 and 1 (Table 1), representing the confidence that a given pixel corresponds to a PV installation. The total number of network parameters is approximately 2.5×10^3, allowing efficient inference over large areas.

The forward propagation of the network can be written as:

$$h_1 = \mathrm{ReLU}(xW_1 + b_1),$$

$$h_2 = \mathrm{ReLU}(h_1W_2 + b_2),$$

$$\hat{y} = \sigma(h_2W_3 + b_3),$$

where σdenotes the sigmoid function and $\hat{y} \in (0,1)$is the predicted probability that the pixel belongs to a PV installation.

Model training was performed in PyTorch using the binary cross-entropy loss function:

$$L = -\frac{1}{N}\sum_{i=1}^{N}[y_i\log(\hat{y}_i) + (1 - y_i)\log(1 - \hat{y}_i)].$$

The optimizer was stochastic gradient descent (SGD), with a learning rate of 0.001, a batch size of 32 and 100 training epochs. After each epoch, model performance on the validation set was evaluated using overall accuracy (OA), precision (P), recall (R) and F1-score to monitor convergence and generalization. All training was conducted in a GPU (CUDA) environment.

To evaluate classification performance comprehensively, we used the following metrics:

$$OA = \frac{TP + TN}{TP + TN + FP + FN},$$

$$P = \frac{TP}{TP + FP},$$

$$R = \frac{TP}{TP + FN},$$

$$F1 = \frac{2PR}{P + R},$$

where TP, FP, TN and FN denote the numbers of true positives, false positives, true negatives and false negatives, respectively.

S1.3 Post-processing and extraction of utility-scale PV assets

The raw model output is a 10-m raster of PV occurrence probability. To derive a high-quality and physically meaningful PV map suitable for subsequent estimation of generation capacity and matching with electricity demand, we further post-processed the probability surface through probability thresholding, binarization, removal of small patches, spatial aggregation and exclusion of anomalously large patches.

Elevation, slope, snow and glacier masks

To remove, at an early stage, areas that are unsuitable for PV development and prone to false detections, we applied pixel-wise topographic and cryospheric screening to the model-derived PV probability surface $P(x)$, yielding a pre-filtered probability map $P_{\mathrm{pre}}(x)$.

We first derived elevation $H(x)$ and slope $S(x)$ from a digital elevation model, and set thresholds of $H_{\max} = 4{,}500\mathrm{m}$ and $S_{\max} = 35°$. A topographic validity mask was then defined as:

$$M_T(x) = \begin{cases} 1, & H(x) < H_{\max} \text{ and } S(x) < S_{\max} \\ 0, & \text{otherwise} \end{cases}$$

where x denotes pixel location; $M_T(x) \in \{0,1\}$ is the topographic mask; $H(x)$ is pixel elevation (m); $S(x)$ is pixel slope (°); and $H_{\max}$ and $S_{\max}$ are the elevation and slope thresholds, respectively.

Permanent snow- and ice-covered areas were identified from land-cover class labels $L(x)$, and glacier extents were further identified from glacier boundary data. The snow–ice–glacier validity mask $M_C(x)$ was defined as:

$$M_C(x) = \begin{cases} 1, & L(x) \neq \text{Snow/Ice and } x \notin \text{Glacier} \\ 0, & \text{otherwise} \end{cases}$$

where $M_C(x) \in \{0,1\}$ is the snow–ice–glacier mask; $L(x)$ is the land-cover class label; Snow/Ice denotes the snow-and-ice category; and $x \notin$ Glacier indicates that the pixel does not fall within glacier boundaries.

The two masks were combined pixel by pixel and applied to update the probability surface, producing the preliminarily filtered PV probability map $P_{\mathrm{pre}}(x)$.

Extraction and binarization of high-confidence PV pixels

The pre-filtered probability map $P_{\mathrm{pre}}(x)$ was then thresholded, retaining only high-confidence pixels with probability greater than or equal to 0.90. This produced an initial 10-m binary PV mask:

$$M_{10}(x) = T \square \big(P_{\mathrm{pre}}(x)\big) = \begin{cases} 1, & P(x) \geq 0.90 \\ 0, & \text{otherwise} \end{cases}$$

This threshold was chosen to maximize precision, consistent with common remote-sensing practice for identifying medium- to large-scale artificial facilities.

Removal of small patches

Small patches in the binary image, often caused by noise, mixed pixels or local reflectance anomalies, were removed using 8-neighbour connected-component analysis. For each connected component C_i, the number of pixels $| C_i |$was compared with a minimum patch threshold $T_{\min}$(50–100 pixels, corresponding to approximately 0.005–0.01 km^2). The filtered mask was defined as:

$$M'_{10} = \bigcup_{i:|C_i|\geq T_{\min}} C_i$$

Removal of morphologically irregular patches

To further remove candidate PV patches with strongly jagged boundaries, narrow fragmented shapes or spurious connections created by noise, we applied a joint post-processing strategy consisting of morphological regularization, shape-consistency screening and hole filling to the mask M'_{10}. This strategy first uses morphological opening to suppress small protrusions and narrow connections, then evaluates shape regularity from the retention ratio before and after opening, and finally applies morphological closing to fill small holes and smooth boundaries.

(1) Morphological opening: boundary regularization and suppression of narrow connections

We first applied morphological opening to M'_{10} to weaken small protrusions, serrated edges and narrow connections. Let S be the structuring element (here a 3×3square kernel). Morphological opening is defined as erosion followed by dilation:

$$M_{\text{open}} = (M'_{10} \ominus S) \oplus S$$

where $\ominus$and $\oplus$denote binary erosion and binary dilation, respectively. To strengthen the regularization effect, the opening operation was iterated $n_o$times ($n_o = 5$), thereby more effectively suppressing fragmented and highly irregular boundaries.

(2) Shape-consistency screening: removal of irregular patches based on overlap ratio

We then performed 8-neighbour connected-component analysis on M'_{10}, yielding the set of connected patches $\{C_i\}_{i=1}^{N}$. For each patch C_i, we calculated its retention ratio in the morphologically opened result M_{open}:

$$r_i = \frac{| C_i \cap M_{\text{open}} |}{| C_i | + \varepsilon}$$

where $\varepsilon$is a small constant introduced to avoid division by zero ($\varepsilon = 10^{-6}$). This metric reflects the stability of each patch under morphological regularization: compact patches with relatively regular boundaries typically retain high values of r_i, whereas elongated, fragmented or strongly serrated patches lose substantial numbers of pixels during opening and thus yield low r_i.

Using a shape-consistency threshold of $\rho = 0.6$, we retained only patches with $r_i \geq \rho$, producing the morphologically filtered mask M^*_{10}:

$$M_{10}^{*} = \bigcup_{i:\, r_i \geq \rho} C_i$$

This overlap-based screening can be regarded as an explicit constraint on the geometric regularity of candidate patches, and it effectively removes false detections caused by noise connections, fragmented structures or highly irregular boundaries.

(3) Morphological closing: hole filling and boundary smoothing

After irregular patches had been removed, we applied morphological closing to $M_{10}^{*}$to fill small internal holes, bridge minor gaps and further smooth patch boundaries. Morphological closing, defined as dilation followed by erosion, is given by:

$$M_{\text{close}} = (M_{10}^{*} \oplus S) \ominus S$$

The closing operation was iterated $n_c$times ($n_c = 1$). Without substantially altering the main patch geometry, this step improves spatial continuity and provides a more stable binary segmentation result for subsequent area-based filtering of exceptionally large patches and for scale aggregation.

Removal of anomalously large patches

In some environments, such as bare land, salt lakes and other highly reflective surfaces, spatial aggregation may produce exceptionally large connected regions with high predicted PV probability. To ensure that the mapped assets remain physically plausible and meaningful from an engineering perspective, we applied a maximum patch-area threshold $A_{\max}$. Any connected patch $C_k$exceeding the size of the largest currently known PV plant worldwide (approximately 40–60 km ^{2}) was removed:

$$M'_{1\text{km}} = \bigcup_{k:\, A_k \leq A_{\max}} C_k$$

where $A_k$is the area of connected patch C_k.

Supplementary Figures

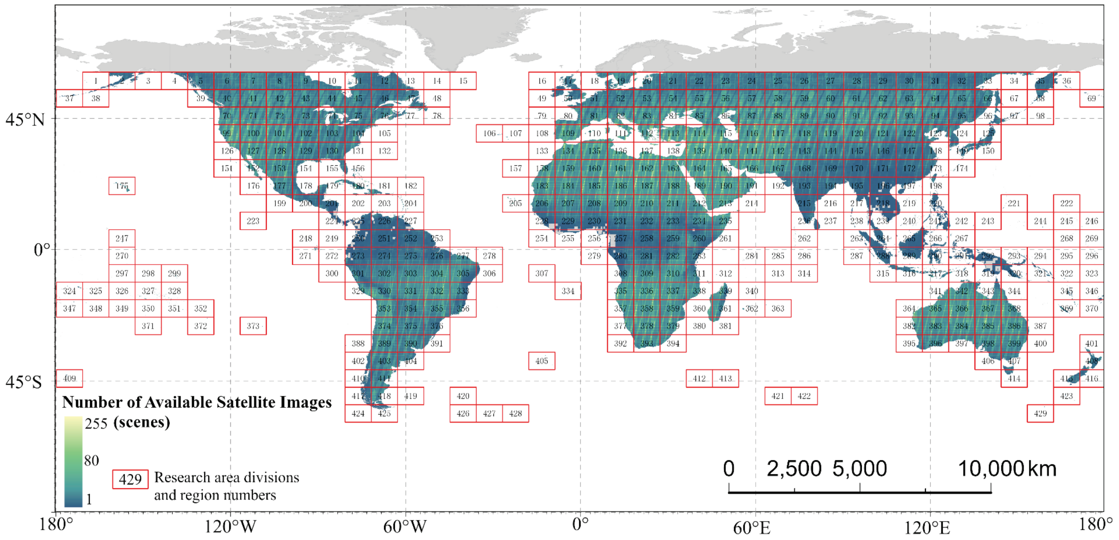

Supplementary Fig. 1 | Global partitioning of the 60°S-60°N study area and the 429 valid tiles used in this study.

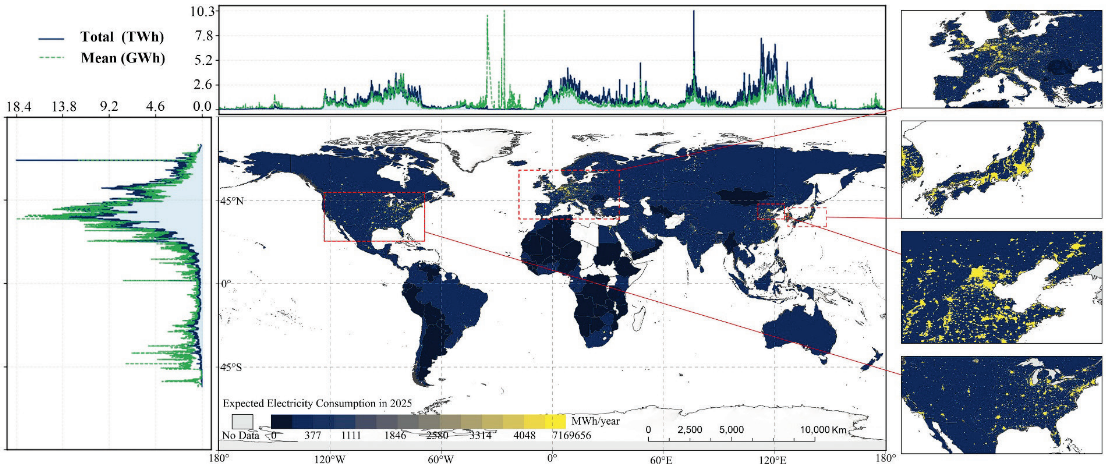

Supplementary Fig. 2 | Global pattern of projected electricity demand in 2025 used in the PV supply-demand coupling analysis.

Supplementary Fig. 3 | Representative comparisons of utility-scale PV asset delineation.(a) High-resolution Google Earth imagery for representative utility-scale PV sites.(b) PV extraction results produced by the model developed in this study.(c) Corresponding PV extraction results from the global PV inventory dataset released by Kruitwagen et al. in 2021.

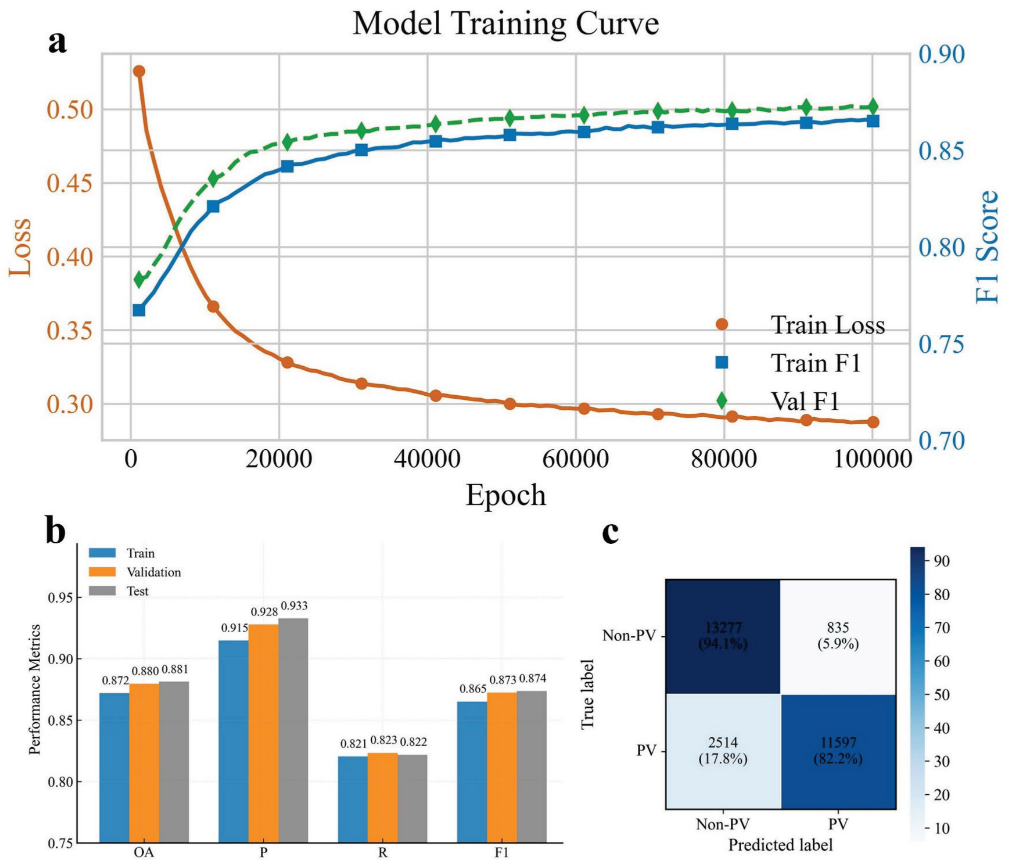

Supplementary Fig. 4 | Training process and classification performance of the deep neural network used for global mapping of utility-scale PV assets.

Supplementary Tables

Supplementary Table. 1 | Installed capacity density by latitude band, conservatively estimated based on the Land-use Control Indicators for Photovoltaic Power Station Engineering Projects issued by the Ministry of Natural Resources of China

Latitude band (°)	Installed capacity density (MW/km^2)）	Corresponding Q_i (Kwp/m^2)
0-18(<18)	97	0.097
18-20	93	0.093
20-25	84	0.084
25-30	74	0.074
30-35	63	0.063
35-40	51	0.051
40-45	39	0.039
45-50	25	0.025
50-60	12.5	0.0125

Supplementary Table. 2 | Architecture of the proposed FC-DNN model.

Layer	Type	Units	Activation
Input	Feature vector	10	-
FC1	Fully connected (FC)	32	ReLU
FC2	Fully connected (FC)	64	ReLU
Output	Fully Connected (Dense)	1	Sigmoid

Supplementary Table. 3 | Sentinel-2 Level-2A input bands used for global utility-scale PV mapping

Band	Spectral region	Central wavelength(nm)	Spatial resolution	Role in PV mapping
B2	Blue	490	10	Captures reflectance differences in the short visible wavelengths and helps distinguish PV panels from bare land, roads and some artificial surfaces.
B3	Green	560	10	Helps characterize differences in surface brightness and colour, improving the separability of PV facilities from surrounding land surfaces.
B4	Red	665	10	Used to identify differences in red-band reflectance between PV sites and vegetation, bare land and built-up areas.
B5	Red-edge1	705	20	Provides red-edge information and enhances discrimination between PV facilities and vegetated backgrounds.
B6	Red-edge2	740	20	Complements red-edge spectral variation and improves target-recognition stability under complex background conditions.
B7	Red-edge3	783	20	Further characterizes red-edge spectral shape and helps suppress confusion with vegetation and some highly reflective backgrounds.
B8	NIR	842	10	Provides near-infrared reflectance information and strengthens discrimination of PV facilities from vegetation, water bodies and bare land.
B8a	Narrow NIR	865	20	Provides narrow-band near-infrared information complementary to B8 and assists discrimination under complex land-surface conditions.
B11	SWIR 1	1610	20	Reflects differences in surface moisture and material properties, helping distinguish PV facilities from soil, construction materials and some bare-land backgrounds.
B12	SWIR 2	2190	20	Provides longer-wave shortwave-infrared information, enhances characterization of surface material differences, and helps reduce false positives from highly reflective non-PV surfaces.

Supplementary Table. 4 | Summary of public data sources used in this study

Dataset	Provider	Variables used	Application in this study
Sentinel-2 Level-2A MSI imagery	ESA/Copernicus	Surface reflectance	Utility-scale PV asset mapping
Global PV distribution location dataset	Published dataset[21]	Known PV locations	Positive-sample construction for PV model training
ESA WorldCover	ESA WorldCover project	Land-cover classes	Land-cover attribution; snow/ice screening; post-processing support
Digital elevation model	NASA JPL / USGS[68]	Elevation; slope	Terrain filtering in PV post-processing
Long-term PV generation climatology	Global Solar Atlas; ESMAP; Solargis; World Bank Group	PV generation per unit installed capacity	Realized PV supply estimation
Land-use Control Indicators for Photovoltaic Power Station Engineering Projects	Ministry of Natural	Latitude-dependent installed-capacity density parameters	Realized PV supply estimation
VIIRS nighttime-lights monthly composites (VCMCFG)	NOAA / VIIRS	Nighttime radiance	Electricity demand estimation
Global 1-km gridded electricity-consumption dataset	Published dataset[67]	Electricity consumption	Baseline demand surface for extrapolation to 2025
World Economic Outlook database (October 2025 edition)	IMF	GDP growth rate	Electricity demand extrapolation
GHS-POP R2023A	JRC / GHSL	Population distribution	Population-weighted interpretation of under-coverage
Economically developable PV potential dataset	Published dataset[43]	Developable PV potential	Scenario-based allocation; potential-centre extraction

201