SATELLITE-BASED SOLAR NOWCASTING FOR ENERGY COMMUNITIES: AN AI-SUPPORTED FRAMEWORK FOR PV PANEL DETECTION AND PRODUCTION PREDICTION

Sead Mustafic[1], Irene Schicker[2], Lukas Prenner[3], Nikta Madjdi[2], Pascal Gfäller[2], Matthias Schlögl[2], Matthias Göbel[2], Petrina Papazek[2], Jasmina Hadzimustafic[2], Doris Oberleiter[3], Roland Perko[1]

[1] JOANNEUM RESEARCH, Graz, Austria, {firstname.lastname}@joanneum.at
[2] GeoSphere Austria, Vienna, Austria, {firstname.lastname}@geosphere.at
[3] energyfamily, Amstetten, Austria, office@energyfamily.at

Abstract: Energy communities require accurate photovoltaic (PV) forecasts to improve self-consumption and maintain grid stability under variable solar generation. This paper presents an AI-based framework that combines automated PV detection from high-resolution aerial imagery with satellite-driven solar nowcasting for community-level energy management. PV systems are identified using segmentation-based deep learning refined through region-specific labeling and active learning. Solar irradiance and PV production are forecast at 15-minute resolution with lead times up to 12 hours using an ensemble of satellite-based, numerical weather prediction, and data-driven models. Results demonstrate reliable detection of medium and large PV installations and complementary nowcasting strengths across seasonal conditions, enabling actionable, high-resolution forecasts for operational use in energy communities.

Keywords: energy communities, solar nowcasting, PV panel detection, satellite data

1 Introduction

Energy communities enable collective renewable energy production and sharing but face challenges from photovoltaic (PV) production variability and grid balancing. The research project entitled *Satellite-based solar nowcasting for energy communities* (PV4C[1]) develops an AI-supported framework combining satellite-based PV panel detection with high-resolution solar nowcasting for community-level energy management. **energyfamily** is the application-oriented project partner. As online market-place for distributed, renewable energy, energyfamily offers a web platform for the management & billing of energy communities. Currently > 300 communities and > 25,000 users are using the services of energyfamily. A core motivation of energyfamily is to provide energy communities with accessible, automated forecasts that improve self-consumption and serve as a foundation for operational decision-making and strategic flexibility deployment.

Within this work, the following technical objectives are presented: (i) automated PV panel detection using deep learning on aerial imagery (20 cm resolution), (ii) ensemble nowcasting

[1] https://www.joanneum.at/digital/projekte/pv4c/ (accessed 31.01.2026).

delivering 15-minute forecasts up to 12 hours ahead at 1 km resolution, downscaled to rooftop level (10-50 m resolution), and (iii) incorporation of all information into a web-based dashboard (see Figure 1). As primary test sites the greater Vienna area and Amstetten were chosen.

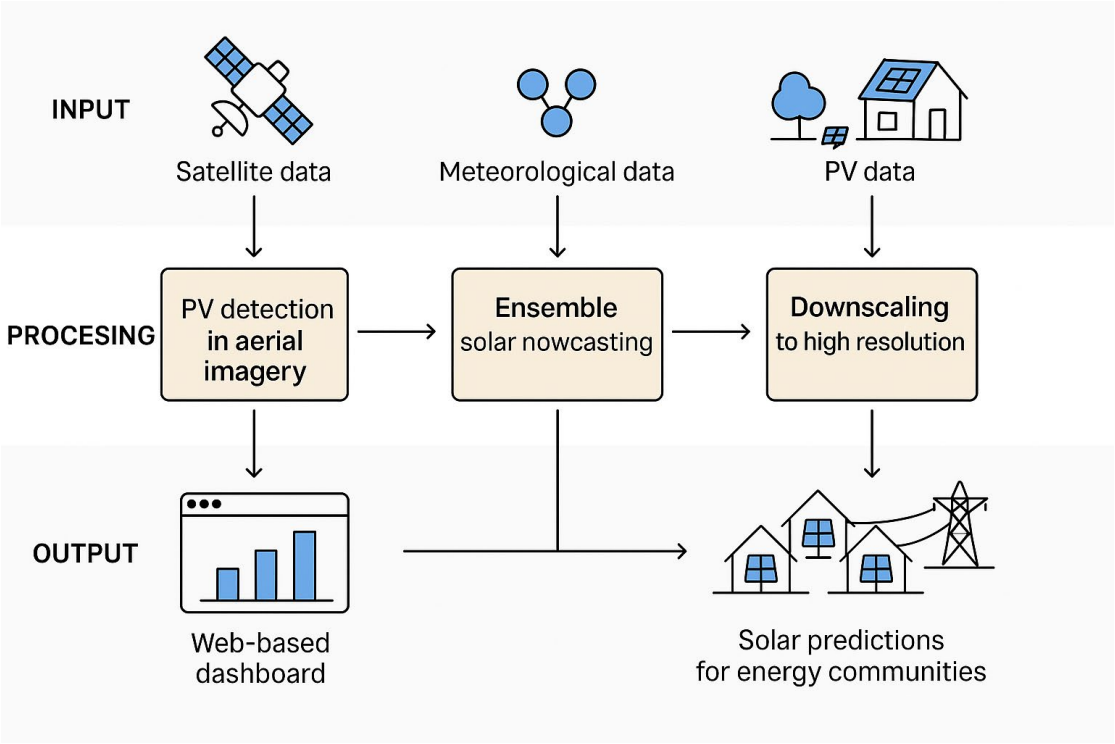

Figure 1: Workflow of solar nowcasting for energy communities.

2 Methodology

2.1 Data Integration

The framework integrates aerial imagery from the Federal Office of Metrology and Surveying in Austria (Bundesamt für Eich- und Vermessungswesen (BEV)) [3], European labeled PV datasets (BDAPPV [8], SolarDE [1], SolarDK [9]), 270+ TAWES meteorological stations, CAMS radiation data, and real energy community production and consumption data from over 200 communities. These data sources can be grouped into two main categories: (i) datasets used for PV system detection from aerial imagery, and (ii) datasets supporting nowcasting and energy-related modeling tasks.

Austrian BEV aerial imagery constitutes the primary data source for PV panel detection and represent the target domain for model inference. The BEV provides ortho-rectified airborne imagery acquired with UltraCam [10] sensors, offering four spectral bands (red, green, blue, and near-infrared) at a ground sampling distance (GSD) of approximately 20cm. These high-resolution images capture roof structures and urban layouts typical of Austrian cities and municipalities, enabling precise localization of rooftop PV panels.

To support the development of robust detection models, several publicly available labeled PV datasets were incorporated, such as BDAPPV, SolarDE, and SolarDK. These datasets provide manually annotated PV installations across diverse geographic regions, architectural styles, and acquisition conditions, thereby increasing variability in roof typologies and PV system appearances. Within the scope of this project, selected BEV images are further manually

labeled to reflect region-specific characteristics of Austrian urban and suburban architecture. This ensures that typical roof materials, building geometries, and PV installation patterns are explicitly represented in the training dataset.

Global radiation and sunshine duration observations from TAWES monitoring network, Cloud Type Level 2 product from the Meteosat Third Generation (MTG) Flexible Combined Imager (FSI), and high-resolution numerical weather prediction models AROME (Application of Research to Operations at Mesoscale, [14]) and C-LAEF (Convection permitting Limited Area Ensemble Forecast model, [15]) have been used as a data basis for the meteorological analysis and nowcasting models. Strong coupling to cloud analysis and cloud-motion nowcasting enables high spatial detail and very short-range accuracy of radiation forecasts, while also accounting for the effects of Alpine topography [11]. For satellite-based nowcasting, the LSA SAF (Satellite Application Facility on Land Surface Analysis) Meteosat Second Generation (MSG) products Downward Surface Shortwave Flux (DSSF) and diffuse fraction (FRACT) global irradiation and fraction of diffuse radiation forms the base. Additionally, the topographical information and cyclical encoded datetimes were added. In the next step, a migration to the MTG data is anticipated. For PV forecasting, Copernicus Atmospheric Monitoring Service (CAMS) data was used additionally.

2.2 PV Panel Detection

Photovoltaic (PV) panel detection from aerial imagery constitutes a critical component of the proposed framework, as it provides the spatial basis for estimating installed capacities and enabling rooftop-level production modeling within energy communities. Given the heterogeneity of roof structures, installation patterns, and sensor characteristics, PV detection remains a challenging computer vision task, particularly when models are transferred across geographic regions or imaging systems.

2.2.1 Detection Paradigms and Model Selection

Three principal detection paradigms were evaluated within this study: (i) binary classification at image patch level (presence/absence of PV systems), (ii) object detection using bounding boxes, and (iii) pixel-wise semantic segmentation. While classification approaches are computationally efficient, they do not provide spatially explicit information on PV system extent and are therefore unsuitable for downstream tasks such as capacity estimation or rooftop-level aggregation. Bounding-box detectors (YOLO [12] and DEIM [7]) were evaluated as baselines; however, while they provide coarse localization, their box-shaped outputs do not capture precise PV contours. This limits both reliable detection in complex rooftop scenes and the accurate estimation of PV surface area required for downstream aggregation.

Consequently, this work focuses primarily on segmentation-based deep learning architectures, which enable precise delineation of PV system contours. Three widely used segmentation architectures were evaluated: U-Net [13], U-Net++ [18], and SegFormer [17]. These architectures differ in their encoder–decoder design, multi-scale feature aggregation, and attention mechanisms, offering complementary strengths in capturing both local texture and global context. In addition, a pretrained Mask R-CNN [6] model from the GeoAI library [16] was included as a benchmark, providing instance-level segmentation and bounding boxes. The model was used to investigate cross-domain generalization behavior rather than to serve as a performance reference.

2.2.2 Training Data and Domain Shift Challenges

The primary target domain for inference is Austrian aerial imagery provided by BEV. However, substantial domain differences exist between BEV imagery and commonly used public PV datasets. These differences include, besides variations in ground sampling distance, sensor-specific characteristics (spectral response, radiometric calibration), acquisition parameters (flight altitude, viewing geometry), seasonal effects, and regional architectural styles.

Experiments with models trained exclusively on publicly available datasets, such as SolarDE, SolarDK, BDAPPV, and other European and non-European PV benchmarks, demonstrated limited generalization to BEV data. In particular, labels derived from different sensors and acquisition geometries were often not transferable due to geometric inconsistencies. A notable challenge arises from the fact that aerial imagery used for PV detection is often not provided as a true orthophoto. For elevated structures, including rooftops, this can result in horizontal displacements of several meters relative to ground-referenced objects, leading to substantial misalignment between imagery and externally sourced labels.

Temporal mismatch between imagery and annotations further exacerbates this issue. PV installations have increased significantly in Austria over the past few years, such that labels generated from older imagery or external databases frequently fail to reflect the current state of rooftop installations. These factors collectively necessitated the creation of BEV-specific training labels to ensure accurate and reliable PV detection.

2.2.3 BEV-Specific Label Generation

To address the limitations of label transferability, selected BEV image patches were manually annotated to generate high-quality, region-specific PV labels. Labeling focused on accurately delineating visible PV installations on rooftops, accounting for Austrian roof typologies, common installation layouts, and material characteristics. Particular attention was paid to ambiguous cases, such as dark-colored roofs, glass-covered terraces, and rooftop structures that visually resemble PV panels.

The labeling process was iterative and tightly integrated with model training. The greater Vienna area was selected as the primary labeling region, as it encompasses a wide range of dense urban, suburban, and peri-urban roof structures representative of typical Austrian settlement patterns. The availability of BEV imagery for multiple acquisition epochs enabled an efficient transfer and adaptation of manually created labels across different time periods, requiring comparatively limited additional annotation effort. In addition to the Vienna region, several smaller and structurally distinct areas were included to further increase model robustness and reduce potential regional bias.

2.2.4 Active Learning Strategy

Given the high cost of manual annotation and the lack of publicly available labeled BEV data, an active learning (AL) approach was adopted to iteratively improve model performance with minimal labeling effort. Active learning leverages model uncertainty to identify the most informative samples for annotation, typically those located near the decision boundary or exhibiting ambiguous visual features.

Based on preliminary experiments, two segmentation architectures (U-Net++ and SegFormer) were selected for the active learning experiments, while U-Net was excluded due to its lower baseline performance and to reduce computational overhead. The initial active learning

iteration started from weights pretrained on publicly available PV datasets, while all subsequent cycles followed two complementary training strategies in parallel: continued fine-tuning from the previous iteration and repeated reinitialization from the same pretrained weights used initially. This parallel training setup enabled a direct comparison between cumulative fine-tuning and repeated reinitialization strategies and helped identify potential model drift effects during active learning.

A total of seven active learning cycles were conducted. In each cycle, the current model ensemble was applied to a large pool of unlabeled BEV patches, and uncertainty measures were used to select between 50 and 250 new image patches for manual labeling. To mitigate model drift and error reinforcement, targeted manual interventions were introduced after the first and fourth active learning cycles by explicitly adding systematically misclassified or uncertain samples (e.g., roof windows mistaken for PV or false negatives on dark roofs).

2.3 Analysis and Nowcasting

IrradPhyDNet, the satellite-based machine learning nowcasting model, combines PhyCells for PDE-approximation with an ConvLSTM encoder-decoder architecture, trained on the LSA SAF DSSF satellite product using timestep-dropout for missing frames. INCA (Integrated Nowcasting through Comprehensive Analysis) [5] integrates sunshine duration measurements and MTG FCI cloud type product (2 km, 10-minute updates) with limited area NWP models AROME (2.5 km) and C-LAEF (1 km, hourly). The cloudiness nowcasts are based on the cloud motion vectors derived from consecutive cloudiness analyses, with the extrapolation of the cloud fields performed for a forecast range of up to +6 h. NWP total cloud cover is interpolated onto the INCA grid and combined with the extrapolated field in the nowcast period from +2 h to +6 h. For global irradiance, the NWP is used from the analysis time step on (first guess field), whereas for the nowcasts, persistence forecast and NWP are combined using time weight function, which gives the NWP model full weight from the +6 h on. As a result, global-, direct- and diffuse- irradiance analyses and forecasts are produced on a 1 km × 1 km grid at 15-minute frequency and time step, and with lead times up to 24 h. Derivation of diffuse and direct radiation components was implemented using an adapted version of the algorithm in [4]. The original Gassel method describes a physically consistent partitioning of horizontal global radiation into its beam and diffuse components based on solar geometry and atmospheric transmissivity. For the PV production nowcasts, different methods are combined into a prediction ensemble, namely a direct-model-output conversion to PV production, an LSTM and Random Forest model, and a Graph Neural Network model to account for inter-site dependencies for community predictions.

2.4 Dashboard

The PV4C dashboard serves as the central interface for energy communities and related stakeholders. It visualizes forecasted PV production at both rooftop and community level and supports decision-making through key features such as self-sufficiency and mismatch analytics, consumption and load forecasts, uncertainty visualization (e.g., ensemble spread), GDPR-compliant web-based tool designed to enable actionable insights for flexibility use, storage, and grid planning.

3 Results

3.1 PV Panel Detection

Our PV panel detection analysis shows that bounding-box–based detectors, despite their substantially larger architectures and higher training complexity, offer no clear advantage over lighter segmentation models in the exploratory phase of this study. Object detectors such as YOLO and DEIM were evaluated in preliminary experiments but did not yield sufficiently reliable results on BEV imagery, particularly with respect to the accurate delineation of PV installation geometry. Based on these observations, segmentation-based models, including U-Net, U-Net++, and SegFormer, were selected for further development, as they produced more spatially coherent PV masks and were better suited to represent rooftop installations for subsequent aggregation tasks.

The pretrained Mask R-CNN model from the GeoAI library showed reasonable performance when applied to imagery similar to its original training domain. However, when transferred to Austrian regions and BEV aerial imagery, its predictions exhibited noticeable degradation, particularly in terms of false positives and incomplete segmentation of PV systems. This observation highlights the limited generalizability of models trained on geographically or sensor-specific datasets.

Figure 2 illustrates representative segmentation results obtained from four models (U-Net++ and SegFormer), each trained using either continued fine-tuning or repeated reinitialization within the active learning framework. Large PV installations with clear contrast to the surrounding roof structures are detected relatively well by all models. In these cases, the predicted masks closely follow the visible panel outlines, indicating that dominant PV features are captured reliably.

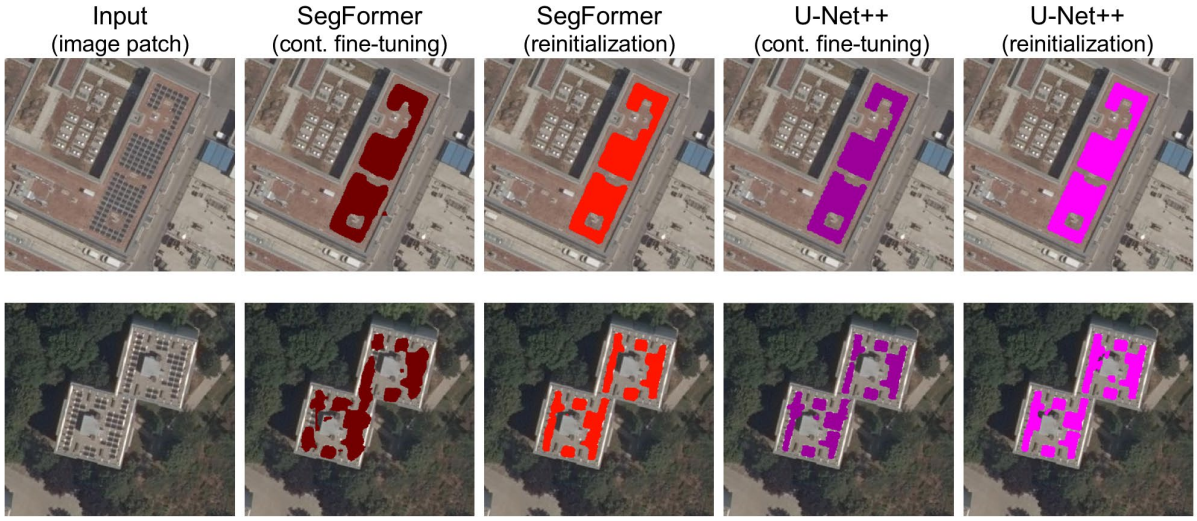

Figure 2: Examples of PV panel detection and segmentation results obtained from four different deep learning models. Each row shows a BEV aerial image patch and the corresponding segmentation outputs from U-Net++ and SegFormer. Large PV installations are generally detected reliably, whereas smaller systems, low-contrast panels on dark roofs, and visually similar rooftop structures remain challenging. Differences between model outputs highlight varying sensitivity to roof texture, illumination, and geometry.

In more challenging scenarios, however, systematic differences between training strategies become apparent. Small-scale installations and PV systems on dark or highly textured roofs are prone to being missed or only partially segmented across all models. In addition, visually similar rooftop elements, such as skylights or glass-covered terraces, are occasionally misclassified as PV panels. Notably, in several cases the continuously fine-tuned models exhibit less optimal segmentation for specific PV shapes, whereas models reinitialized from pretrained weights produce cleaner and more coherent masks. This effect is particularly pronounced for SegFormer, as illustrated by the second image in the lower-left position of Figure 2. This behavior suggests that repeated fine-tuning across active learning cycles may, in some cases, reduce flexibility in adapting to specific PV geometries, particularly after several AL iterations, whereas reinitialization helps preserve more generalizable feature representations.

Across the examined examples, U-Net++ tends to produce slightly more consistent and spatially accurate PV segment boundaries than SegFormer, particularly for elongated or regularly structured installations. While this difference is not uniform across all scenarios, it was observed repeatedly during visual inspection of the segmentation results. At the same time, the qualitative comparison shows that no single model consistently outperforms the others across all rooftop configurations. Instead, detection results vary depending on roof material, illumination conditions, and local texture, underlining the inherent uncertainty of PV detection from high-resolution aerial imagery. To better capture this variability, segmentation outputs from all four models were combined into a single color-coded agreement mask, as shown in Figure 3. The mask encodes the level of agreement between models on a per-pixel basis, ranging from detection by a single model to consensus across all four models.

Figure 3: Combined PV detection results based on multi-model agreement. Segmentation masks from four different models are merged into a single color-coded agreement map, indicating whether a pixel is detected as PV by one, two, three, or all four models. Regions with high agreement correspond to robust detections, while low-agreement areas reveal model uncertainty and ambiguous rooftop structures. This representation supports uncertainty analysis, consistency checks, and targeted refinement within the active learning process.

This combined representation provides an intuitive means of visualizing detection confidence and uncertainty. Regions with high model agreement typically correspond to clearly identifiable PV panels, whereas areas with low agreement often coincide with ambiguous rooftop

structures or low-contrast panels. Beyond visualization, this approach supports consistency checks across models and proved useful for guiding targeted sample selection and manual refinement within the active learning process.

Overall, the results demonstrate that segmentation-based deep learning models, when trained on publicly available PV datasets and subsequently refined on BEV imagery through active learning, can reliably detect large and medium-sized PV installations in Austrian aerial imagery. The detection of small or visually challenging systems remains limited at a ground sampling distance of 20 cm. These findings emphasize both the practical strengths and current limitations of aerial-image–based PV detection and motivate future work toward uncertainty-aware post-processing and multi-sensor extensions.

3.2 Nowcasting Performance

IrradPhyDNet achieves a lower MAE and a higher equitable threat score (ETS) than INCA (with MSG satellite data) during summer under dynamic cloud conditions (cf. Figure 5). However, INCA excels in winter fog conditions. Note that INCA includes the most recent sunshine duration and global irradiance observations whereas IrradPhyDNet does not. The MTG integration into INCA does improve spatial detail in cloud cover and radiation fields with better verification scores than the legacy MSG-driven INCA forecasts as shown in Figure 4. Furthermore, diffuse- and direct irradiance are added as new parameters to the INCA radiation module (Figure 6) and can be used for the PV production predictions.

Initial PV production predictions, using INCA forecasts with MSG data, with the baseline model reveal high accuracy scores for our test site Wien Hohe Warte (cf. Figure 7). One refinement for snow and freezing rain on the PV panels needs to be still implemented.

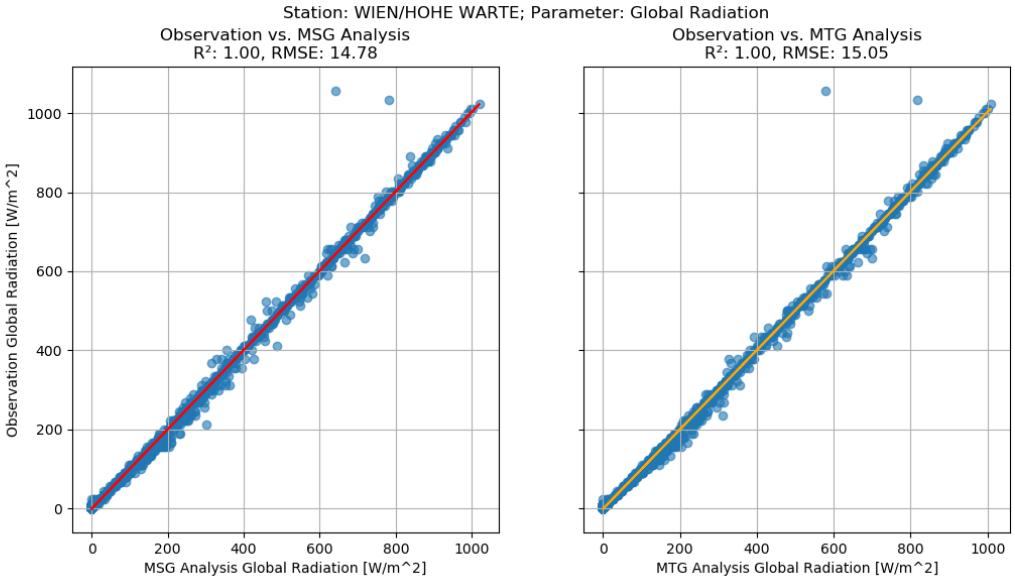

Figure 4: Comparison of 15 min INCA global irradiance analyses with observations at TAWES Hohe Warte for a selected test period May to June 2025. The difference in correlation comes from different satellite input, MSG- (left) and MTG FCI cloud types (right).

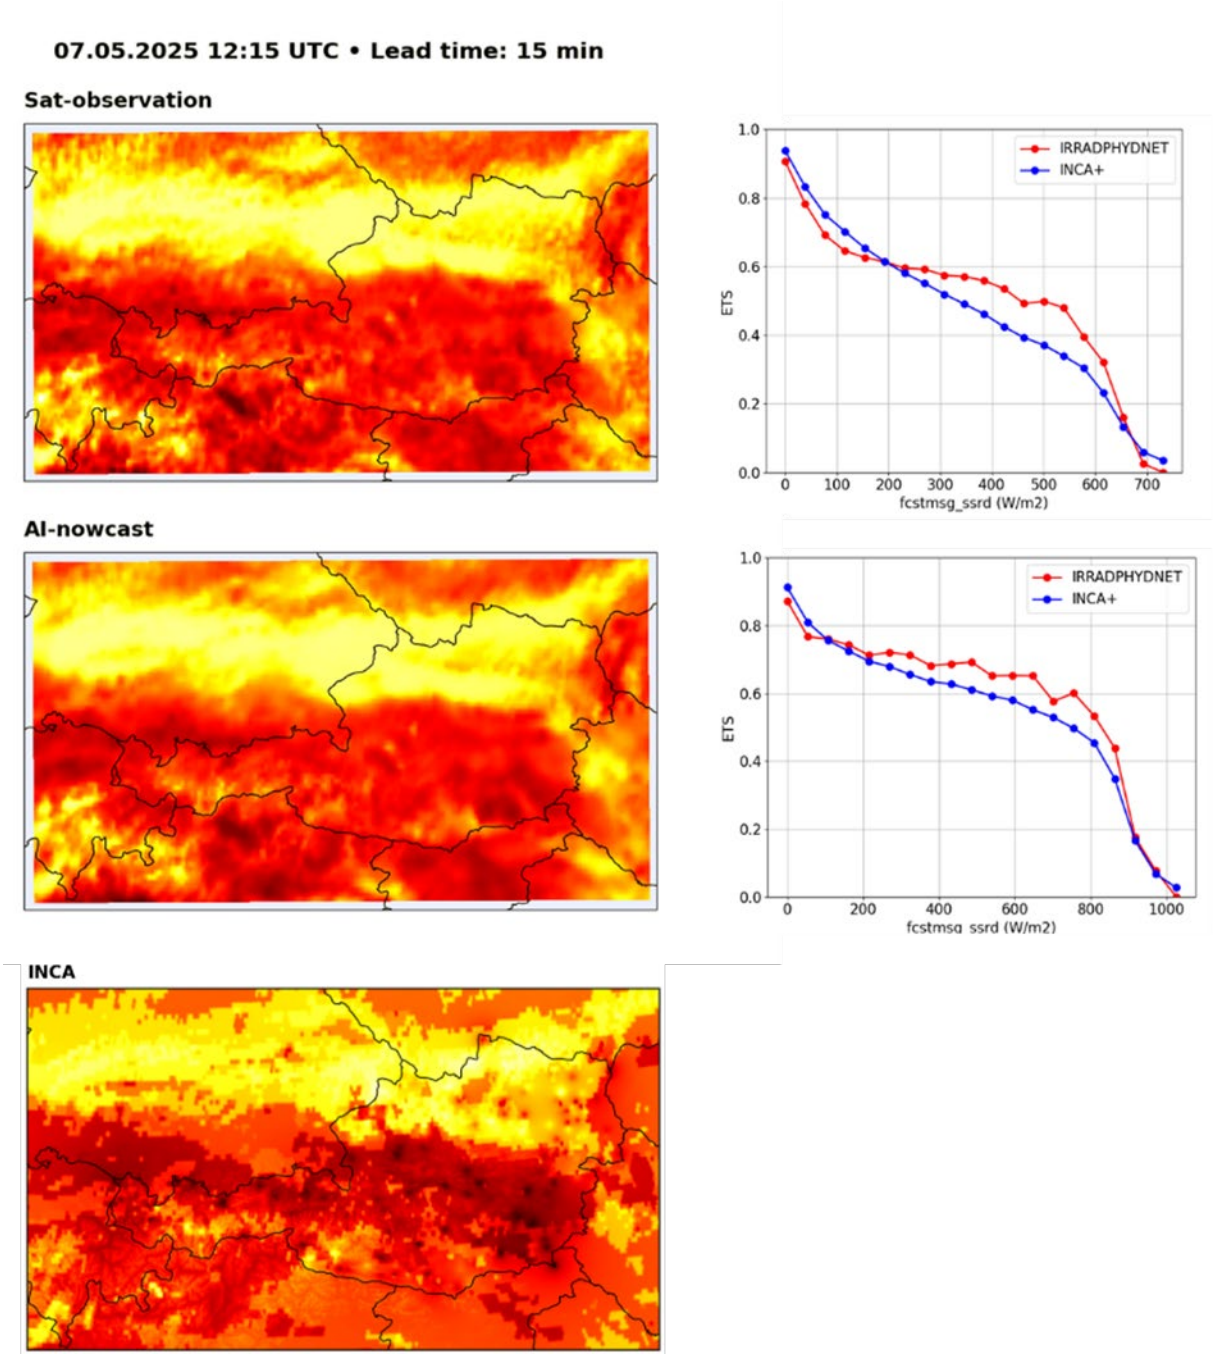

Figure 5: (top-left) LSA SAF observed global irradiance for a selected use case, the corresponding IrraPhyDnet nowcast (middle-left) and INCA nowcast (bottom-left) at lead time +15 minutes ahead. (top-right) ETS score for a winter period, and (bottom-right) for a summer period.

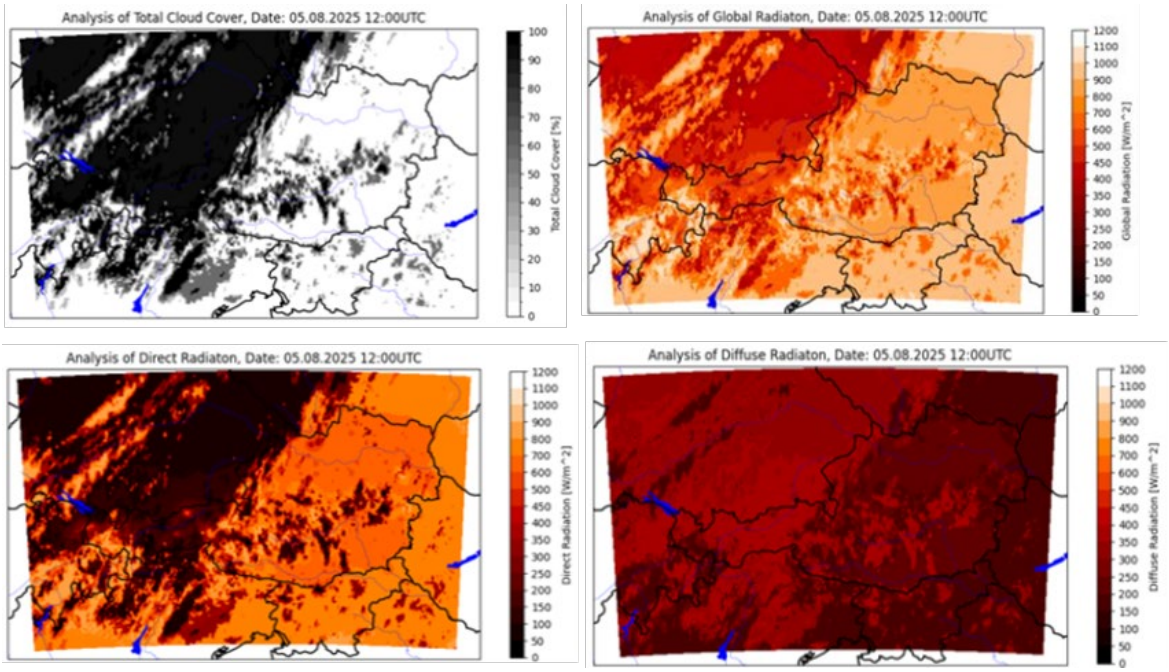

Figure 6: Analysis of cloudiness (top-left) and three irradiance components in INCA for 5.8.2025, 12 UTC calculated with MTG FCI cloud types and C-LAEF 1 km x 1 km NWP Model.

Figure 7: Meteorological and PV production (bottom-right) predictions for the test site Vienna Hohe Warte for January 2026, after the freezing rain event.

3.3 Energy Community Analysis

Real data reveals production-consumption temporal mismatch (midday PV peak vs. morning/evening consumption peaks), self-sufficiency ratios varying from 5:1 to zero, and critical ramp events concentrated mid-morning/midday. Member consumption patterns show high diversity, offering demand response opportunities.

4 Conclusions and Outlook

This work demonstrates how satellite remote sensing and AI address practical energy transition challenges. By automatically detecting PV installations from high-resolution aerial imagery using segmentation-based deep learning models and an active learning strategy tailored to Austrian BEV data, and predicting production at high spatial-temporal resolution, energy communities become active participants in flexible, decarbonized energy systems. The PV detection results show that region-specific training data are essential to overcome domain shifts caused by sensor characteristics, acquisition geometry, temporal label mismatch, and regional architectural differences. In particular, the combination of BEV-specific labeling and uncertainty-driven active learning enables robust and scalable PV mapping while substantially reducing manual annotation effort. Although limitations remain for small or low-contrast installations at a ground sampling distance of 20 cm, the generated PV masks provide a consistent spatial basis for rooftop-level aggregation and subsequent energy modeling.

The evaluation of the solar nowcasting framework shows that the satellite-based IrradPhyDNet model outperforms INCA during summer conditions with rapidly evolving cloud fields when compared against satellite observations, while INCA exhibits superior skill during winter fog situations, mainly due to the assimilation of recent surface observations. The integration of MTG satellite data into INCA further improves the spatial structure and verification scores of radiation fields compared to legacy MSG-based configurations. Initial PV production nowcasts derived from INCA radiation forecasts show high agreement with observed production at the Wien Hohe Warte test site, including during dynamically evolving winter conditions. Remaining forecast errors are primarily linked to surface effects such as snow and freezing rain on PV panels, highlighting clear pathways for future model refinement.

The multi-data integration verified against real operations provides robust forecasts for operational decision-making crucial for managing distributed generation while maintaining grid stability. The combination of scientific innovation and practical applicability ensures that energy communities and infrastructure operators alike benefit from more resilient and data-driven energy systems.

In future, the PV4C project focuses on PV panel detection in super-resolved Sentinel-2 imagery (based on the LDSR-S2 [2] method) using spectral analysis, full ensemble implementation with uncertainty quantification, rooftop-level downscaling using digital surface models, GNN PV production validation, web platform integration with APIs, and comprehensive verification using probabilistic metrics.

Expected impacts include (i) scientific contributions in small-object detection methodologies for medium-resolution imagery and hybrid ML-physics nowcasting, (ii) user-oriented, operational forecasting platform enabling Austrian energy communities to optimize self-consumption, increase resilience, and actively contribute to local grid stability, (iii) policy and planning support through reproducible, high-resolution evidence on distributed PV variability

and ramp events, supporting infrastructure and flexibility investment decisions, and (iv) commercial scalability to grid operators and energy service providers and regional flexibility markets across Austria and beyond.

For energyfamily, the project enables new platform features including improved PV production forecasts, real-time mismatch detection, and dynamic self-sufficiency tracking. These tools support data-driven growth of energy communities and help unlock the value of flexibility assets such as EV chargers, heat pumps, and battery storage – laying the foundation for more resilient and efficient local energy systems.

5 Acknowledgments

This research was partly funded by the Austrian Space Applications Programme (ASAP) through the project PV4C (FFG project number 911917).

6 References

[1] C.N. Clark, and F. Pacifici (2023). A solar panel dataset of very high resolution satellite imagery to support the Sustainable Development Goals, in Sci Data 10, 636.

[2] S. Donike, C. Aybar, L. Gómez-Chova, and F. Kalaitzis (2025). Trustworthy super-resolution of multispectral Sentinel-2 imagery with latent diffusion, in IEEE J. Sel. Topics Appl. Earth Obs. Remote Sens. 18, 6940–6952.

[3] Federal Office of Metrology and Surveying of Austria (2025). [online] https://www.bev.gv.at/

[4] A. Gassel (1999). Beiträge zur Berechnung solarthermischer und energieeffizienter Energiesysteme (Doctoral dissertation, Fraunhofer-IRB-Verlag).

[5] T. Haiden, A. Kann, C. Wittmann, G. Pistotnik, B. Bica, and C. Gruber (2011). The Integrated Nowcasting through Comprehensive Analysis (INCA) system and its validation over the Eastern Alpine region. Weather and Forecasting, 26(2), 166-183.

[6] K. He, G. Gkioxari, P. Dollár, and R. Girshick (2017). Mask R-CNN, in Proc. IEEE Int. Conf. Comput. Vis., pp. 2961–2969.

[7] S. Huang, X. Wang, and K. Chen (2024). DEIM: DETR with Improved Matching for Fast Convergence, arXiv:2412.04234.

[8] G. Kasmi, YM. Saint-Drenan, D. Trebosc, R. Jolivet, J. Leloux, B. Sarr, and L. Dubus (2023). A crowdsourced dataset of aerial images with annotated solar photovoltaic arrays and installation metadata, in Sci Data 10, 59.

[9] M. Khomiakov, J. H. Radzikowski, C. A. Schmidt, M. B. Sørensen, M. Andersen, M. R. Andersen, and J. Frellsen (2022). SolarDK: A High-Resolution Urban Solar Panel Image Classification and Localization Dataset, arXiv preprint arXiv:2212.01260.

[10] F. Leberl, M. Gruber, M. Ponticelli, S. Bernögger, and R. Perko (2003). The UltraCam large format aerial digital camera system, in American Society For Photogrammetry & Remote Sensing, pp. 5–9.

[11] M. Olefs, D. J. Baumgartner, F. Obleitner, C. Bichler, U. Foelsche, H. Pietsch, H. E. Rieder, P. Weihs, F. Geyer, T. Haiden, and W. Schöner (2016). The Austrian radiation monitoring network ARAD – best practice and added value. Atmospheric Measurement Techniques, 9(4), 1513-1531.

[12] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi (2016). You Only Look Once: Unified, Real-Time Object Detection, in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. pp. 779–788.

[13] O. Ronneberger, P. Fischer, and T. Brox (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation, in Medical Image Computing and Computer-Assisted Intervention, vol. 9351, pp. 234–241.

[14] Y. Seity, P. Brousseau, S. Malardel, G. Hello, P. Bénard, F. Bouttier, C. Lac, and V. Masson (2011). The AROME-France Convective-Scale Operational Model. Mon. Wea. Rev., 139, 976–991.

[15] C. Wastl, Y. Wang, A. Atencia, F. Weidle, C. Wittmann, C. Zingerle, and E. Keresturi (2021). C-LAEF: Convection-permitting Limited-Area Ensemble Forecasting system. Quarterly Journal of the Royal Meteorological Society. 147. 10.1002/qj.3986.

[16] Q. Wu (2024). GeoAI: Artificial Intelligence for Geospatial Data, GeoAI Python Package. [online]. https://geoai.gishub.org

[17] E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Álvarez, and P. Luo (2021). SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers, in Adv. Neural Inf. Process. Syst., vol. 34, pp. 12077–12090.

[18] Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh, and J. Liang (2020). UNet++: Redesigning Skip Connections to Exploit Multiscale Features in Image Segmentation, IEEE Trans. Med. Imaging, vol. 39, no. 6, pp. 1856–1867.