
Can We Reliably Improve the Robustness to Image Acquisition of Remote Sensing of PV Systems?

Gabriel Kasmi^{1,2} Laurent Dubus^{2,3} Yves-Marie Saint Drenan¹ Philippe Blanc¹

¹MINES Paris, Université PSL Centre Observation Impacts Energie (O.I.E.)

²RTE France

³WEMC (World Energy & Meteorology Council, UK)

¹{firstname.lastname}@minesparis.psl.eu

²{firstname.lastname}@rte-france.com

Abstract

Photovoltaic (PV) energy is crucial for the decarbonization of energy systems. Due to the lack of centralized data, remote sensing of rooftop PV installations is the best option to monitor the evolution of the rooftop PV installed fleet at a regional scale. However, current techniques lack reliability and are notably sensitive to shifts in the acquisition conditions. To overcome this, we leverage the wavelet scale attribution method (WCAM) [21], which decomposes a model’s prediction in the space-scale domain. The WCAM enables us to assess on which scales the representation of a PV model rests and provides insights to derive methods that improve the robustness to acquisition conditions, thus increasing trust in deep learning systems to encourage their use for the safe integration of clean energy in electric systems.

1 Introduction

Photovoltaic (PV) energy grows rapidly and is crucial for the decarbonization of electric systems [11]. The rapid growth of rooftop PV makes the estimation of the global PV installed capacity challenging as centralized data is often lacking [22]. Remote sensing rooftop PV on orthoimagery with computer vision models is a promising solution for mapping rooftop PV installations. However, current approaches lack reliability and generalize poorly from one image provider to the other [45, 18]. Improving generalizability and reliability requires improving the robustness to acquisition conditions and proposing methods to assess the reliability of the decision process of the PV classifiers [5].

Deep learning-based pipelines became the standard method for remote sensing PV systems. DeepSolar [48] paved the way for country-wide mapping of PV systems using deep learning and overhead imagery. While several works discuss the poor generalizability of current methods [45, 7, 18], these works do not tackle the robustness to varying acquisition conditions, which can be assimilated to image corruptions [13].

In this work, we analyze the robustness to heterogeneous acquisition conditions of models for remote sensing of PV installations using the wavelet scale attribution method (WCAM, [21]). The WCAM assesses the reliability of a model’s decision by decomposing it into the scale-space domain. We analyze the model’s sensitivity to acquisition conditions using the WCAM and derive a principled method for improving the robustness to these acquisition conditions. Our work shows that the WCAM provides a finer understanding of what the model sees as a PV panel and guides us to improve the robustness to acquisition conditions. By improving the reliability and robustness of deep learning models for rooftop PV mapping, we aim to facilitate the mapping and, thus, the integration into the electric grid of rooftop PV.

2 Related works

Remote sensing of PV installations Many works leveraged overhead imagery and deep learning methods to map PV installations [31, 29, 9, 49]. The DeepSolar [48] method marked a significant milestone with mapping distributed and utility-scale installations over the continental United States using state-of-the-art deep learning models. Many works built on DeepSolar to map regions or countries, especially in Europe [24, 1, 22, 7, 32, 33]. However, current methods cannot be transposed from one region to another without incurring accuracy drops, thus limiting their practical usability [18] due to a lack of reliability of the generated data [5]. To address this gap, we propose to study and mitigate the impact of acquisition conditions, ubiquitous with overhead imagery, which prevents reusing trained models for registry updates.

Sensitivity to distribution shifts The sensitivity to distribution shifts [25] prevents from using pre-trained models without further training, whether temporally or spatially, i.e., it limits their ability to generalize [5, 30]. Some works empirically discussed this issue [45] and argued that the generalization ability depended on how hard to recognize the PV panels are. However, no work properly disentangled the effect of each source of variability identified by [42]: geographical conditions, varying acquisition conditions, and the ground sampling distance (GSD).

Frequency-centric explanations A line of works aimed at explaining the behavior of neural networks through the lenses of frequency analysis. Several works showed that convolutional neural networks (CNNs) are biased towards high frequencies [44, 47] and that robust methods tend to limit this bias [51, 2]. Other works highlighted a so-called spectral bias [37, 46, 20], showing that CNNs learn the input frequencies from the lowest to the highest. More recently, using wavelet transforms, [21] expanded attribution methods from the pixel to the space-scale (wavelet) domain. This work connects the fields of interpretability and robustness and enables understanding *what* models see on images. It has not yet been applied to orthoimagery, where scales are explicitly indexed.

3 Data and methods

3.1 Data

We consider the crowdsourced training dataset *Base de données d'apprentissage profond photovoltaïque* (BDAPPV, [23]). This dataset contains annotated images of 28,000 PV panels in France and neighboring countries. This dataset also proposes annotations of images that depict the same PV panel but from two different image providers: images coming from the Google Earth Engine (hereafter referred to as "Google") [10] and from the IGN, the French public operator for geographic information. We have double annotations for more than 8,000 PV systems. It allows us to assess the impact of the acquisition conditions as the only change factor between two images is the varying acquisition condition: the semantic content (the PV panel and its surroundings) remains unchanged. The native ground sampling distance (GSD) of Google images is 10 cm/pixel and 20 cm/pixel for IGN images. We define the acquisition conditions as the properties of the technical infrastructure (airborne or spaceborne, camera type, image quantization, and postprocessing) and the atmospheric and meteorological conditions the day the image was taken. Figure 3 in the appendix A presents images samples of the BDAPPV dataset.

3.2 Methods

3.2.1 Identifying where the sensitivity to distribution shifts comes from

Empirical framework BDAPPV features images of the same installations from two providers and records the crude location of the PV installations. Using this information, we can define three test cases to disentangle the distribution shifts that occur with remote sensing data: the resolution, the acquisition conditions, and the geographical variability. We train a ResNet-50 model [12] on Google images downsampled at 20cm/pixel of resolution and evaluate it on three datasets: a dataset with Google images at their native 10cm/pixel resolution ("Google 10 cm/pixel"), the IGN images with a native 20cm/pixel resolution ("IGN") and Google images downsampled at 20 cm/pixel located outside

of France ("Google OOD¹"). We add the test set to record the test accuracy without distribution shift ("Google baseline"). We only do random crops, rotations, and ImageNet normalizations during training. Figure 4 in appendix A presents examples of images seen during training and test.

3.2.2 Data augmentations for improving the robustness to acquisition conditions

Benchmarking current approaches The literature on robustness to image corruptions [13] proposed numerous data augmentation methods to improve the robustness of classification models to image corruptions [14, 15, 4, 3, 8]. We consider the well-established AugMix method [14] and the recently-proposed RandAugment [4] and AutoAugment [3] methods. These methods apply a random composition of perturbations to images during training to learn an invariance against these perturbations. We do not consider the case of training from multiple sources as our setting is that we wish to generalize to unseen images (either temporally or spatially, so we cannot incorporate knowledge about these images).

Lowering the reliance on high-frequency components Since we know that varying acquisition conditions mainly alter high-frequency components, we introduce two data augmentation techniques that aim at reducing the reliance on high-frequency components: Gaussian blurring ("Blurring") and Blurring + wavelet perturbation (WP). Blurring consists of a fixed image blur, while the Blurring + WP also perturbs the wavelet coefficients of the image to force the model to rely on several rather than one scale for prediction. We refer the reader to the appendix C.1 for more details on the data augmentation strategies and a review of the hyperparameters. Figure 9 in appendix C.2 illustrates the effect of the different data augmentation techniques. We compare this approach with a baseline without augmentations ("ERM") and existing data augmentation techniques.

3.2.3 Understanding the sensitivity to acquisition conditions with the Wavelet sScale Attribution Method (WCAM)

Attribution [40, 39, 35] indicates the important regions for prediction, i.e., decomposes the prediction in the pixel (spatial) domain. The WCAM [21] generalizes attribution to the wavelet (space-scale domain). The WCAM provides us with two pieces of information: where the model sees and what scale (i.e., frequency) it sees at this location. Therefore, we can see if a prediction relies on robust or fragile frequencies. Additionally, the decomposition of the prediction in terms of scales is interpretable, particularly in the case of orthoimagery. For example, on Google images, details at the 1-2 pixel scale correspond to physical objects with a size between 0.1 and 0.2 m on the ground. Thus, we know what the model sees as a panel; we can interpret it and assess whether it is sensitive to varying acquisition conditions. The decomposition brought by the WCAM enables the interpretation of the model's decision process. Appendix B provides additional background for reading WCAMs.

4 Results

4.1 Acquisition conditions mainly explain the poor generalization of PV mapping algorithms

Table 1: **F1 Score** and decomposition in true positives, true negatives, false positives, and false negatives rates of the disentanglement of the distribution shift between the GSD (Google 10 cm/px), the geographical variability (Google OOD) and the acquisition conditions (IGN).

	F1 Score ($\uparrow$)	True positives rate	True negatives rate	False positives rate	False negatives rate
Google baseline	0.98	0.99	0.98	0.02	0.01
Google 10cm/px	0.89	0.81	1.00	0.00	0.19
Google OOD	0.98	0.99	0.98	0.02	0.01
IGN	0.46	0.32	0.95	0.03	0.68

Results Table 1 shows the results of the decomposition of the effect of distribution shifts into three components: resolution, acquisition conditions, and geographical shift. We can see that the F1 score drops the most when the model faces new acquisition conditions. The second most significant impact

¹OOD: out-of-distribution.

comes from the change in the ground sampling distance, but the performance drop remains relatively small compared to the effect of the acquisition conditions. In our framework, there is no evidence of an effect of the geographical variability once we isolate the effects of the acquisition conditions and ground sampling distance. This effect is probably underestimated, as images of our dataset that are not in France are near France. However, the effect of the acquisition conditions is sizeable enough to seek methods for addressing it.

Mechanisms: when important factors disappear

Changing the provider (i.e., altering the acquisition conditions) alters the scales describing the image of PV panels. If the model relied on a scale no longer on the image, it could no longer recognize the PV panel. On Figure 1, we can see that on Google, the important factor was the factor **(b)** (on the leftmost image), which is no longer important on the IGN image (on the right). On the IGN image, the model instead relied on the factor **(a)** as the factor **(b)** is no longer visible. This change in the important factor (at the same scale in this example) seems to have driven the shift from predicting to not predicting the PV panel. Interestingly, we can see that the scales highlighted in **(c)** are visible on both images but not important for the prediction in the IGN image: the model no longer "sees" these details. We refer the reader to the appendix B.2 for further guidance on interpreting a WCAM.




Figure 1: Predictions on Google image (left, upper row) and IGN image (right, upper row) and associated WCAMs (bottom row, displayed with the same color scale). The brighter, the more important the highlighted region for the prediction

4.2 Lowering the reliance on high frequencies improves generalization

Blurring and wavelet perturbation improve accuracy Table 2 reports the results of the evaluated data augmentation techniques to mitigate the effect of acquisition conditions. Augmentations that explicitly discard small scales (high frequencies) information perform the best. However, the blurring method sacrifices the recall (which drops to 0.6) to improve the F1 score. On Table 2, this can be seen by the increase in false positives. Therefore, this method is unreliable for improving the robustness to acquisition conditions. On the other hand, adding wavelet perturbation yields improvements and outperforms existing approaches without sacrificing precision or recall.

Table 2: **F1 Score** and decomposition in true positives, true negatives, false positives, and false negatives for models trained on Google with different mitigation strategies. All models are evaluated on the same test set, so we report the raw values rather than the rates. Evaluation on IGN images. The oracle corresponds to a model trained on IGN images with standard augmentations. Best results are **bolded**.

	F1 Score ($\uparrow$)	True positives	True negatives	False positives	False negatives
Oracle	0.88	1818	1992	428	83
ERM [43]	0.44	566	2321	99	1335
AutoAugment [3]	0.46	598	2318	102	1303
AugMix [14]	0.48	624	2318	102	1277
RandAugment [4]	0.51	707	2280	140	1194
Blurring	0.74	1855	1196	1224	46
Blurring + WP	<u>0.58</u>	896	2114	306	1005

Relying on consistent scales Figure 2 compares the scales on which the best-performing methods rely. In our case, we want our models to rely on the largest scales (i.e., lowest frequencies) to entail robustness [50] against varying acquisition conditions. We can see that the blurring and wavelet perturbation enforces this property better than other data augmentation techniques. Indeed, the model




Figure 2: WCAMs on IGN of models trained on Google with different augmentation techniques.

relies on coarser scales (which are more robust) and on scales on which the ERM also relies. More generally, the WCAM lets us compare methods that perform quantitatively similarly.

On the choice of the training images Our results show that lowering the reliance on high-frequency content in the image improves generalization. This content is located on the 10-20cm scale and only appears on Google images. In Table 3, we flip our experiment to study how a model trained on IGN images generalizes to Google images. Results show that the model trained on IGN generalizes better to the downscaled Google images than the opposite. This result further supports the idea that higher GSD is not necessarily better for good robustness to acquisition conditions.

Table 3: **F1 Score** and true positives, true negatives, false positives, and false negatives. Evaluation computed on the Google dataset. ERM was trained on Google and Oracle on IGN images.

	F1 Score ($\uparrow$)	True positives	True negatives	False positives	False negatives
ERM [43]	0.98	1891	2355	36	39
Oracle (ERM trained on IGN)	0.91	1815	2127	264	115

5 Conclusions and future work

We set up an experiment to disentangle the effects of heterogeneous acquisition conditions, geographical variability, and ground sampling distance on the generalization of deep neural networks to unseen data. Our results show that the sensitivity to acquisition conditions is the leading cause of poor generalization. To explain why models are sensitive to acquisition conditions, we leverage the wavelet scale attribution method (WCAM, [21]). Acquisition conditions perturb the scales the model relied on to make a prediction. If these scales correspond to high frequencies, they are likely to be disrupted by the acquisition conditions. We show that models biased towards low frequencies are more robust to acquisition conditions. We design a data augmentation method that outperforms other methods to improve the robustness to varying acquisition conditions. More generally, models trained on images with a lower GSD generalize better.

Broader impact Currently, transmission system operators (TSOs) lack quality data regarding rooftop PV installations [22]. The lack of information leads to imprecise estimations and forecasts of the overall PV power generation, which in a context of sustained growth of the PV installed capacity could increase the uncertainty and threaten the grid’s stability [36]. On the other hand, current methods for mapping rooftop PV installations lack reliability, owing to their poor generalization abilities beyond their training dataset [5]. This work addresses this gap and thus demonstrates that remote sensing of PV installations is a reliable way for TSOs to improve their knowledge regarding small-scale PV installations.

Future works We wish to discuss further the conditions on the training images for good robustness to acquisition conditions. In particular, we plan to discuss the trade-off between the minimal GSD to *reliably* see PV panels [26] and a notion of image quality for the training data.

6 Acknowledgements

This work is funded by RTE France, the French transmission system operator, and benefited from CIFRE funding from the ANRT. The authors gratefully acknowledge the support of this project.

References

- [1] Edoardo Arnaudo, Giacomo Blanco, Antonino Monti, Gabriele Bianco, Cristina Monaco, Paolo Pasquali, and Fabrizio Dominici. A Comparative Evaluation of Deep Learning Techniques for Photovoltaic Panel Detection from Aerial Images. *IEEE Access*, pages 1–1, 2023.
- [2] Yiting Chen, Qibing Ren, and Junchi Yan. Rethinking and Improving Robustness of Convolutional Neural Networks: a Shapley Value-based Approach in Frequency Domain. October 2022.
- [3] Ekin D. Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V. Le. AutoAugment: Learning Augmentation Policies from Data, April 2019. arXiv:1805.09501 [cs, stat].
- [4] Ekin D. Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V. Le. RandAugment: Practical automated data augmentation with a reduced search space, November 2019. arXiv:1909.13719 [cs].
- [5] Tim De Jong, Stefano Bromuri, Xi Chang, Marc Debusschere, Natalie Rosenski, Clara Schartner, Katharina Strauch, Marion Boehmer, and Lyana Curier. Monitoring Spatial Sustainable Development: semi-automated analysis of Satellite and Aerial Images for Energy Transition and Sustainability Indicators. *arXiv preprint arXiv:2009.05738*, 2020.
- [6] Thomas Fel, Remi Cadene, Mathieu Chalvidal, Matthieu Cord, David Vigouroux, and Thomas Serre. Look at the Variance! Efficient Black-box Explanations with Sobol-based Sensitivity Analysis, November 2021. arXiv:2111.04138 [cs].
- [7] Âzeddine Frimane, Robert Johansson, Joakim Munkhammar, David Lingfors, and Johan Lindahl. Identifying small decentralized solar systems in aerial images using deep learning. *Solar Energy*, 262:111822, September 2023.
- [8] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. April 2023.
- [9] Vladimir Golovko, Alexander Kroshchanka, Sergei Bezobrazov, Anatoliy Sachenko, Myroslav Komar, and Oleksandr Novosad. Development of Solar Panels Detector. In *2018 International Scientific-Practical Conference Problems of Infocommunications. Science and Technology (PIC S&T)*, pages 761–764, Kharkiv, Ukraine, October 2018. IEEE.
- [10] Noel Gorelick, Matt Hancer, Mike Dixon, Simon Ilyushchenko, David Thau, and Rebecca Moore. Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote sensing of Environment*, 202:18–27, 2017. Publisher: Elsevier.
- [11] Nancy M Haegel, Robert Margolis, Tonio Buonassisi, David Feldman, Armin Froitzheim, Raffi Garabedian, Martin Green, Stefan Glunz, Hans-Martin Henning, Burkhard Holder, and others. Terawatt-scale photovoltaics: Trajectories and challenges. *Science*, 356(6334):141–143, 2017. Publisher: American Association for the Advancement of Science.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [13] Dan Hendrycks and Thomas Dietterich. Benchmarking Neural Network Robustness to Common Corruptions and Perturbations, March 2019. arXiv:1903.12261 [cs, stat].
- [14] Dan Hendrycks, Norman Mu, Ekin D. Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. AugMix: A Simple Data Processing Method to Improve Robustness and Uncertainty, February 2020. arXiv:1912.02781 [cs, stat].
- [15] Dan Hendrycks, Andy Zou, Mantas Mazeika, Leonard Tang, Bo Li, Dawn Song, and Jacob Steinhardt. PixMix: Dreamlike Pictures Comprehensively Improve Safety Measures, March 2022. arXiv:2112.05135 [cs].
- [16] Wassily Hoeffding. A Class of Statistics with Asymptotically Normal Distribution. In Samuel Kotz and Norman L. Johnson, editors, *Breakthroughs in Statistics: Foundations and Basic Theory*, Springer Series in Statistics, pages 308–334. Springer, New York, NY, 1992.
- [17] Toshimitsu Homma and Andrea Saltelli. Importance measures in global sensitivity analysis of nonlinear models. *Reliability Engineering & System Safety*, 52(1):1–17, April 1996.
- [18] Wei Hu, Kyle Bradbury, Jordan M. Malof, Boning Li, Bohao Huang, Artem Streltsov, K. Sydney Fujita, and Ben Hoen. What you get is not always what you see—pitfalls in solar array assessment using overhead imagery. *Applied Energy*, 327:120143, December 2022.

- [19] Michiel J. W. Jansen. Analysis of variance designs for model output. *Computer Physics Communications*, 117(1):35–43, March 1999.
- [20] Jason Jo and Yoshua Bengio. Measuring the tendency of CNNs to Learn Surface Statistical Regularities, November 2017. arXiv:1711.11561 [cs, stat].
- [21] Gabriel Kasmi, Laurent Dubus, Yves-Marie Saint Drenan, and Philippe Blanc. Assessment of the Reliability of a Model’s Decision by Generalizing Attribution to the Wavelet Domain, September 2023. arXiv:2305.14979 [cs, stat].
- [22] Gabriel Kasmi, Laurent Dubus, Yves-Marie Saint-Drenan, and Philippe Blanc. Towards Unsupervised Assessment with Open-Source Data of the Accuracy of Deep Learning-Based Distributed PV Mapping. In Thomas Corpetti, Dino Ienco, Roberto Interdonato, Minh-Tan Pham, and Sébastien Lefèvre, editors, *Proceedings of MACLEAN: MACHINE Learning for EArth Observation Workshop co-located with the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML/PKDD 2022), Grenoble, France, September 18-22, 2022*, volume 3343 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2022.
- [23] Gabriel Kasmi, Yves-Marie Saint-Drenan, David Trebosc, Raphaël Jolivet, Jonathan Leloux, Babacar Sarr, and Laurent Dubus. A crowdsourced dataset of aerial images with annotated solar photovoltaic arrays and installation metadata. *Scientific Data*, 10(1):59, January 2023.
- [24] Bala Bhavya Kausika, Diede Nijmeijer, Iris Reimerink, Peter Brouwer, and Vera Liem. GeoAI for detection of solar photovoltaic installations in the Netherlands. *Energy and AI*, 6:100111, December 2021.
- [25] Pang Wei Koh and Percy Liang. Understanding Black-box Predictions via Influence Functions, December 2020. arXiv:1703.04730 [cs, stat].
- [26] Peiran Li, Haoran Zhang, Zhiling Guo, Suxing Lyu, Jinyu Chen, Wenjing Li, Xuan Song, Ryosuke Shibasaki, and Jinyue Yan. Understanding rooftop PV panel semantic segmentation of satellite and aerial images for better using machine learning. *Advances in Applied Energy*, 4:100057, November 2021.
- [27] S.G. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):674–693, July 1989. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [28] Stéphane Mallat. *A wavelet tour of signal processing*. Elsevier, 1999.
- [29] Jordan M. Malof, Kyle Bradbury, Leslie M. Collins, and Richard G. Newell. Automatic detection of solar photovoltaic arrays in high resolution aerial imagery. *Applied Energy*, 183:229–240, December 2016.
- [30] Jordan M Malof, Boning Li, Bohao Huang, Kyle Bradbury, and Artem Stretslov. Mapping solar array location, size, and capacity using deep learning and overhead imagery. page 6, 2019.
- [31] Jordan M. Malof, Rui Hou, Leslie M. Collins, Kyle Bradbury, and Richard Newell. Automatic solar photovoltaic panel detection in satellite imagery. In *2015 International Conference on Renewable Energy Research and Applications (ICRERA)*, pages 1428–1431, Palermo, Italy, November 2015. IEEE.
- [32] Kevin Mayer, Benjamin Rausch, Marie-Louise Arlt, Gunther Gust, Zhecheng Wang, Dirk Neumann, and Ram Rajagopal. 3D-PV-Locator: Large-scale detection of rooftop-mounted photovoltaic systems in 3D. *Applied Energy*, 310:118469, March 2022.
- [33] Kevin Mayer, Zhecheng Wang, Marie-Louise Arlt, Dirk Neumann, and Ram Rajagopal. DeepSolar for Germany: A deep learning framework for PV system mapping from aerial imagery. In *2020 International Conference on Smart Energy Systems and Technologies (SEST)*, pages 1–6, Istanbul, Turkey, September 2020. IEEE.
- [34] William J. Morokoff and Russel E. Caflisch. Quasi-Monte Carlo Integration. *Journal of Computational Physics*, 122(2):218–230, December 1995.
- [35] Vitali Petsiuk, Abir Das, and Kate Saenko. RISE: Randomized Input Sampling for Explanation of Black-box Models, September 2018. arXiv:1806.07421 [cs].
- [36] Marco Pierro, Fabio Romano Liolli, Damiano Gentili, Marcello Petitta, Richard Perez, David Moser, and Cristina Cornaro. Impact of PV/Wind Forecast Accuracy and National Transmission Grid Reinforcement on the Italian Electric System. *Energies*, 15(23):9086, November 2022.
- [37] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred A. Hamprecht, Yoshua Bengio, and Aaron Courville. On the Spectral Bias of Neural Networks, May 2019. arXiv:1806.08734 [cs, stat].
- [38] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge, January 2015. arXiv:1409.0575 [cs].
- [39] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. *International Journal of Computer Vision*, 128(2):336–359, February 2020. arXiv:1610.02391 [cs].

- [40] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition, April 2015. arXiv:1409.1556 [cs].
- [41] Ilya Meerovich Sobol. On sensitivity estimation for nonlinear mathematical models. *Matematicheskoe modelirovanie*, 2(1):112–118, 1990. Publisher: Russian Academy of Sciences, Branch of Mathematical Sciences.
- [42] Devis Tuia, Claudio Persello, and Lorenzo Bruzzone. Domain adaptation for the classification of remote sensing data: An overview of recent advances. *IEEE geoscience and remote sensing magazine*, 4(2):41–57, 2016. Publisher: IEEE.
- [43] Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 1999.
- [44] Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P. Xing. High-Frequency Component Helps Explain the Generalization of Convolutional Neural Networks. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8681–8691, Seattle, WA, USA, June 2020. IEEE.
- [45] Rui Wang, Joseph Camilo, Leslie M. Collins, Kyle Bradbury, and Jordan M. Malof. The poor generalization of deep convolutional networks to aerial imagery from new geographic locations: an empirical study with solar array detection. In *2017 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*, pages 1–8, Washington, DC, October 2017. IEEE.
- [46] Zhi-Qin John Xu, Yaoyu Zhang, Tao Luo, Yanyang Xiao, and Zheng Ma. Frequency Principle: Fourier Analysis Sheds Light on Deep Neural Networks. *Communications in Computational Physics*, 28(5):1746–1767, June 2020. arXiv:1901.06523 [cs, stat].
- [47] Dong Yin, Raphael Gontijo Lopes, Jonathon Shlens, Ekin D. Cubuk, and Justin Gilmer. A Fourier Perspective on Model Robustness in Computer Vision, September 2020. arXiv:1906.08988 [cs, stat].
- [48] Jiafan Yu, Zhecheng Wang, Arun Majumdar, and Ram Rajagopal. DeepSolar: A Machine Learning Framework to Efficiently Construct a Solar Deployment Database in the United States. *Joule*, 2(12):2605–2617, December 2018.
- [49] Jiangye Yuan, Hsiu-Han Lexie Yang, Olufemi A. Omitaomu, and Budhendra L. Bhaduri. Large-scale solar panel mapping from aerial images using deep convolutional networks. In *2016 IEEE International Conference on Big Data (Big Data)*, pages 2703–2708, Washington DC, USA, December 2016. IEEE.
- [50] Jiajin Zhang, Hanqing Chao, Amit Dhurandhar, Pin-Yu Chen, Ali Tajer, Yangyang Xu, and Pingkun Yan. When Neural Networks Fail to Generalize? A Model Sensitivity Perspective, December 2022. arXiv:2212.00850 [cs].
- [51] Zhuang Zhang, Dejian Meng, Lijun Zhang, Wei Xiao, and Wei Tian. The range of harmful frequency for DNN corruption robustness. *Neurocomputing*, 481:294–309, April 2022.

A Additional figures

Examples of images from BDAPPV Figure 3 plots examples of images coming from BDAPPV. These images depict the same scene for two different providers.




Figure 3: Examples of test of the same PV panels but with different providers (Up Google, down: IGN).

Test sets Figure 4 plots examples of the different test images to disentangle the effects of distribution shifts. The baseline and IGN images represent the same panel at the same spatial resolution. The Google 10 cm/pixel depicts the same scene but with the native resolution of Google images. Finally, the OOD test set contains images located outside of France.




Figure 4: Test images on which a model trained on Google images (downsampled to 20 cm/px of GSD) is evaluated.

B Reading WCAMs

B.1 Identification of the important scales in the input image

The content of this section comes from section 2.2. of [21]. Our only addition is Figure 5 and its comment (between "Figure 5 presents" and "larger than 8 pixels on the image.")

Dyadic wavelet transform A wavelet is an integrable function $\psi \in L^2(\mathbb{R})$ with zero average, normalized and centered around 0. Unlike a sinewave, a wavelet is localized in space and in the Fourier domain. It implies that dilatations of this wavelet enable to scrutinize different frequencies (scales) while translations enable to scrutinize spatial location. To compute an image's (continuous) wavelet transform (CWT), one first defines a filter bank $\mathcal{D}$ from the original wavelet ψ with the scale factor s and the 2D translation in space u . We have

$$\mathcal{D} = \left\{ \psi_{s,u}(x) = \frac{1}{\sqrt{s}} \psi \left(\frac{x-u}{s} \right) \right\}_{u \in \mathbb{R}^2, s \geq 0}, \quad (1)$$

where $|\mathcal{D}| = J$, and J denotes the number of levels. The computation of the wavelet transform of a function $f \in L^2(\mathbb{R})$ at location x and scale s is given by

$$\mathcal{W}(f)(x, s) = \int_{-\infty}^{+\infty} f(u) \frac{1}{\sqrt{s}} \psi^* \left(\frac{x-u}{s} \right) du, \quad (2)$$

which can be rewritten as a convolution [28]. Computing the multilevel decomposition of f requires applying Equation 2 J times with all dilated and translated wavelets of $\mathcal{D}$. [27] showed that one could implement the multilevel dyadic decomposition of the discrete wavelet transform (DWT) by applying a high-pass filter H to the original signal f and subsampling by a factor of two to obtain the *detail* coefficients and applying a low-pass filter G and subsampling by a factor of two to obtain the *approximation* coefficients. Iterating on the approximation coefficients yields a multilevel transform where the j^{th} level extracts information at resolutions between 2^j and 2^{j+1} pixels. The detail coefficients can be decomposed into horizontal, vertical, and diagonal components when dealing with images.

Figure 5 presents an example of a two-level dyadic wavelet transform. The leftmost image is the input image. On the center-left image, the northwest, southwest, and southeast panels represent the *detail* coefficients (horizontal, diagonal, and vertical, respectively) at the smallest scale obtained from the high pass filtering of the input image. The northeast panel represents the *approximation* coefficients obtained from the low-pass filtering of the input image. On the center-right image, in the northeast panel, the localization of the detail and approximation coefficients is analogous to the decomposition from stage 1; instead, the reference is the approximation from stage 1, not the input image. On the rightmost image, the third decomposition level is obtained from the high pass filtering of the approximation coefficients from stage 2.

The detail coefficients at stage 1 show details at the 1-2 pixel scale. The detail coefficients at stage 2 are the 2-4 pixel scale details. At stage 3, we can see the detailed coefficients at the 4-8 pixel scale. The approximation coefficients contain everything larger than 8 pixels on the image.




Figure 5: Example of a multilevel dyadic wavelet transform

Sobol sensitivity analysis Let $(X_1, \dots, X_K)$ be independent random variables and $\mathcal{K} = \{1, \dots, K\}$ denote the set of indices. Let f be a model, X an input, and $f(X)$ the model's decision (e.g., the output probability). We denote $f_\kappa = f_\kappa(X_\kappa)$ the partial contributions of the variables $(X_k)_{k \in \kappa}$ to the score $f(X)$. The Sobol-Hoeffding decomposition [16] decomposes the decision score $f(X)$ into summands of increasing dimension

$$f(X) = f_\emptyset + \sum_{\kappa \in \mathcal{P}(\mathcal{K}) \setminus \{\emptyset\}} f_\kappa(X_\kappa), \quad (3)$$

Where $f_\emptyset$ denotes the prediction with no features. Under the orthogonality condition $\forall (u, v) \in \mathcal{K}^2$ such that $u \neq v$, $\mathbb{E}[f_u(X_u)f_v(X_v)] = 0$, we derive from Equation 3 the variance of the model's score

$$\text{Var}(f(X)) = \sum_{\kappa \in \mathcal{P}(\mathcal{K})} \text{Var}(f_\kappa(X_\kappa)), \quad (4)$$

Equation 4 enables us to describe the influence of a subset κ of features as the ratio between its own and total variance. This corresponds to the first order **Sobol index** given by

$$S_\kappa = \frac{Var(f_\kappa(X_\kappa))}{Var(f(X))}. \quad (5)$$

S_κ measures the proportion of the output variance $Var(f(X))$ explained by the subset of variables X_κ [41]. In particular, S_k only captures the *direct* contribution of the feature X_k to the model's decision. To capture the indirect effect, due to the effect of X_k on the other variables, **total Sobol indices** S_{T_k} [17] can be computed as

$$S_{T_k} = \sum_{\kappa \in \mathcal{P}(\mathcal{K}), k \in \kappa} S_\kappa. \quad (6)$$

Total Sobol indices (TSIs) measure the contribution of the k^{th} feature, taking into account both its *direct* effect and its *indirect* effect through its interactions with the other features.

Efficient estimation of Sobol indices As seen from Equation 5, estimating the impact of a feature k on the model's decision requires recording the partial contribution $f_k(X_k)$. This partial contribution corresponds to a *forward*. Estimating Sobol indices requires computing variances by drawing at least N samples and computing N forwards to estimate a first-order Sobol index S_k of a single feature k . As we are interested in the TSI of a feature k , we need to estimate the Sobol index of all sets of features $\kappa \in \mathcal{K}$ such that $k \in \kappa$. To minimize the computational cost of this computation, [6] introduced an efficient sampling strategy based on Quasi-Monte Carlo methods [34] to generate the N perturbations of dimension K applied to the input and used Jansen's estimator [19] to estimate the TSIs given the models' outputs and the quasi-random perturbations. Their approach requires $N(K+2)$ forwards [6].

To estimate the TSIs, they draw two matrices from a Quasi-Monte Carlo sequence of size $N \times K$ and convert them into perturbations, which they apply to X . The perturbed input yields two matrices, A and B . a_{jk} (resp. b_{kj}) is the element of A (resp. B) corresponding to the k^{th} feature and the j^{th} sample. For the k^{th} feature, they define $C^{(k)}$ in the same way as A , except that the column corresponding to feature k is replaced by the column of B . They then derive an empirical estimator for the Sobol index and TSI as

$$\hat{S}_k = \frac{\hat{V} - \frac{1}{2N} \sum_{j=1}^N [f(B_j) - f(C_j^{(k)})]^2}{\hat{V}}, \quad \hat{S}_{T_k} = \frac{\frac{1}{2N} \sum_{j=1}^N [f(A_j) - f(C_j^{(k)})]^2}{\hat{V}}, \quad (7)$$

where $f_\emptyset = \frac{1}{N} \sum_{j=1}^N f(A_j)$ and $\hat{V} = \frac{1}{N-1} \sum_{j=0}^N [f(A_j) - f_\emptyset]^2$. Further implementational details can be found in [6]. Figure 6 summarizes the Wavelet sCale Attribution Method of [21].




Figure 6: Flowchart of the wavelet scale attribution method (WCAM). [21] generate a set of perturbed images by perturbing the image’s discrete wavelet transform (DWT). Expanding [6], they apply masks generated from a Quasi Monte-Carlo sequence to the DWT of the image. They evaluate the model on the altered samples reconstructed from the perturbed DWT. They compute the Sobol index for each mask component using the predicted probabilities and the perturbation masks using the Jansen estimator [19]

B.2 Decomposing the scales of the PV panels

On orthoimagery, PV panels are located in space and in scale. Figure 7 shows an example. Depending on the scale of interest, the PV panel will show different characteristics. At the smallest scales (i.e., high frequencies), the PV panel corresponds to small details within the individual PV modules. On the opposite, if we consider the PV system as a whole, it has a size of about 10m, i.e., 100 pixels in this image.




Figure 7: Decomposition of the scales of a PV panel

The WCAM enables us to see the important scales for the prediction and their location. It lets us know what the model sees from the panel: the overall system, the cluster, or individual modules.

Besides, it also illustrates the impact of background noise on the detection. If small scales not located on the panel are important for the prediction, it indicates that the model relies on background noise for the prediction. For example, the image of Figure 7 could be the grass around the house.

B.3 Some examples

Figure 8 presents some examples of predictions made by a PV classifier. We can see that for our illustration image (leftmost image on Figure 8), the grass, but also the shape of the garden on the upper left of the image, play a small role in the prediction of the PV panel. On the other hand, in the two subsequent images, the prediction relies on factors at different scales located around the PV panel.

The WCAM lets us see that when a model focuses on a PV panel, its focus can be decomposed into different scales. Besides, these scales do not always have the same importance. For instance, the importance of the 1-2 pixel details is more important on the third image from the left than for the other images. It could be because the panel is frameless, so the model has to rely on other factors, as this one is missing for this image.




Figure 8: Decomposition in the space-scale domain of PV panel predictions

C Data augmentation strategies

C.1 Description of the data augmentations

AugMix [14] The data augmentation strategy "Augment-and-Mix" (AugMix) consists of producing a high diversity of augmented images from an input sample. A set of operations (perturbations) to be applied to the images are sampled, along with sampling weights. The image resulting x_{aug} is obtained through the composition $x_{aug} = \omega_1 op_1 \circ \dots \omega_n op_n(x)$ where x is the original image. Then, the augmented image is interpolated with the original image with a weight m that is also randomly sampled. We have $x_{augmix} = mx + (1 - m)x_{aug}$.

AutoAugment [3] This strategy aims at finding the best data augmentation for a given dataset. The authors determined the best augmentations strategy S as the outcome of a reinforcement learning problem: a controller predicts an augmentation policy from a search space. Then, the authors train a model, and the controller updates its sampling strategy S based on the train loss. The goal is that the controller generates better policies over time. The authors derive optimal augmentation strategies for various datasets, including ImageNet [38], and show that the optimal policy for ImageNet generalizes well to other datasets.

RandAugment [4] This strategy's primary goal is to remove the need for a computationally expansive policy search before model training. Instead of searching for transformations, random probabilities are assigned to the transformations. Then, each resulting policy (a weighted sequence of K transformations) is graded depending on its strength. The number of transformations and the strength are passed as input when calling the transformation.

Blurring We apply a nonrandom Gaussian blur to the image. The value is set by comparing visually Google and IGN images and trying to remove details from Google images that are not visible on IGN images. After a manual inspection, we set the blur level so that the details at 10-20cm scale are discarded from the image. It corresponds to a blurring value $\sigma = 2$. in the `ImageFilter.GaussianBlur` method of the PIL library.

Blurring + Wavelet perturbation (WP) We first blur the image. Then, for each color channel, we compute the dyadic wavelet transform of the image and randomly perturb the coefficients (we randomly set some coefficients to 0). The set of coefficients set to 0 is determined with uniform sampling. This results in a random perturbation that removes information for some precise scales and locations. We then reconstruct the image from its perturbed wavelet coefficients. For each call, 20% of the coefficients are canceled. This value balances between the loss of information and the input perturbation. We perturb each color channel independently.

C.2 Plots

Figure 9 plots examples of the different data augmentations implemented in this work. Along with these augmentations, we apply random rotations, symmetries, and normalization to the input during training. At test time, we only normalize the input images.




Figure 9: Visualization of the different data augmentation techniques implemented in this work.

D Complementary results

D.1 Training results on the Google test set

Training results Table 4 reports the training results of our methods on the source (Google) test set. We can see that our spectral method performs slightly less than the other methods on the source dataset.

Table 4: **F1 Score** and decomposition in true positives, true negatives, false positives, and false negatives for models trained on Google images with different strategies to mitigate the sensitivity to acquisition conditions. Evaluation computed on the Google (source) dataset.

	F1 Score ($\uparrow$)	True positives	True negatives	False positives	False negatives
ERM [43]	0.98	1891	2355	36	39
AutoAugment [3]	0.98	1906	2340	51	24
AugMix [14]	0.98	1894	2354	37	36
RandAugment [4]	0.98	1907	2342	49	23
Blurring	0.82	1636	1958	433	294
Blurring + WP	0.90	1798	2135	256	132

D.2 Probability shift

Figure 10 shows the change in predicted probabilities for positive Google images. We can see that when the classifier no longer recognizes the PV panel, the probability shift is large, suggesting that the important factor for prediction disappeared from the image. These results provide more systematic evidence of the qualitative pattern highlighted by Figure 1: if a critical component is no longer depicted due to the change in the acquisition condition, then the model no longer sees a PV panel (despite having other information about it).




Figure 10: Evolution of the predicted probabilities for images depicting a PV panel on the Google test set and the corresponding images on the IGN test set. The predicted probability completely flips over when the model no longer recognizes the PV panel.

D.3 Understanding false positives

On Figure 11, we plot some examples of false positives and their associated WCAM. We can see that the WCAM helps us understand why the model misleads an element of the image with a PV panel.

In the leftmost image, we can see that the shape of the roof mainly causes the false alarm, which is reminiscent of the delineation of a PV system on a roof. On the other hand, in the fourth image (from the left), multiple factors at different scales lead to a wrong prediction.




Figure 11: Examples of false positives on IGN and corresponding WCAM.