



HAL
open science

Space-scale Exploration of the Poor Reliability of Deep Learning Models: the Case of the Remote Sensing of Rooftop Photovoltaic Systems

Gabriel Kasmi, Laurent Dubus, Yves-Marie Saint-Drenan, Philippe Blanc

► To cite this version:

Gabriel Kasmi, Laurent Dubus, Yves-Marie Saint-Drenan, Philippe Blanc. Space-scale Exploration of the Poor Reliability of Deep Learning Models: the Case of the Remote Sensing of Rooftop Photovoltaic Systems. 2024. <hal-04843911>

HAL Id: hal-04843911

<https://minesparis-psl.hal.science/hal-04843911v1>

Preprint submitted on 17 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

METHODS ARTICLE

Space-scale Exploration of the Poor Reliability of Deep Learning Models: the Case of the Remote Sensing of Rooftop Photovoltaic Systems

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

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

²RTE France, Paris, France

³WEMC (World Energy & Meteorology Council, UK)

Received xx xxx xxxx

Keywords: deep learning, distribution shifts, solar energy, remote sensing, wavelets

Abstract

Photovoltaic (PV) energy grows rapidly and is crucial for the decarbonization of electric systems. However, centralized registries recording the technical characteristics of rooftop PV systems are often missing, making it difficult to accurately monitor this growth. The lack of monitoring could threaten the integration of PV energy into the grid. To avoid this situation, the remote sensing of rooftop PV systems using deep learning emerged as a promising solution. However, existing techniques are not reliable enough to be used by public authorities or transmission system operators (TSOs) to construct up-to-date statistics on the rooftop PV fleet. The lack of reliability comes from the fact that deep learning models are sensitive to distribution shifts. This work proposes a comprehensive evaluation of the effects of distribution shifts on the classification accuracy of deep learning models trained to detect rooftop PV panels on overhead imagery. We construct a benchmark to isolate the sources of distribution shift and introduce a novel methodology that leverages explainable artificial intelligence (XAI) and decomposition of the input image and model's decision in terms of scales to understand how distribution shifts affect deep learning models. Finally, based on our analysis, we introduce a data augmentation technique meant to improve the robustness of deep learning classifiers to varying acquisition conditions. We show that our proposed approach outperforms competing methods. We discuss some practical recommendations for mapping PV systems using overhead imagery and deep learning models.

Impact Statement

This paper analyzes the effects of distribution shifts on deep learning models trained to detect rooftop photovoltaic (PV) systems on aerial imagery by combining explainable artificial intelligence methods. It then proposes practical solutions grounded in this analysis to improve the robustness of these models, thus improving their reliability and facilitating the use of remote sensing techniques to improve insertion of rooftop PV systems into the grid.

1. Introduction

Photovoltaic (PV) energy grows rapidly and is crucial for the decarbonization of electric systems (Haegel et al., 2017). The rapid growth of rooftop PV makes the estimation of the global PV installed

capacity challenging as centralized data is often lacking (Hu et al., 2022; Kasmi et al., 2022). Remote sensing of rooftop PV on orthoimagery with deep learning models is a blooming solution for mapping rooftop PV installations. Deep learning-based pipelines became the standard method for remote sensing PV systems as works such as DeepSolar (Yu et al., 2018) paved the way for country-wide mapping of PV systems using deep learning and overhead imagery. Recently, methods for mapping rooftop PV systems in many regions, especially in Europe, have been proposed (Frimane et al., 2023; Kasmi et al., 2022; Kausika et al., 2021; Lindahl et al., 2023; Mayer et al., 2020; Rausch et al., 2020; Zech & Ranalli, 2020). Some of these works (Kasmi et al., 2022; Mayer et al., 2022) introduced methods to estimate the technical characteristics of the PV systems (individual localization, orientation, PV installed capacity).

However, current approaches are sensitive to so-called distribution shifts (De Jong et al., 2020), i.e., differences between the training and testing data (Koh et al., 2021). This sensitivity limits their ability to generalize to new images without incurring significant accuracy drops, thus limiting their practical usability (Hu et al., 2022) as the generated data lacks reliability (De Jong et al., 2020). Steps towards improving the reliability of the deep learning-based registries of rooftop PV systems have been taken, with Hu et al., 2022 and Kasmi et al., 2022 discussing the practical evaluation of the mapping algorithms or Li et al., 2021 defining the minimum resolution to detect rooftop PV systems from overhead images. However, to date, only Wang et al., 2017 studied the poor generalizability of PV mapping algorithms, but their evaluation framework was limited to the case study of two cities and one image dataset. Therefore, a comprehensive evaluation of the causes and consequences of distribution shifts on PV panel detectors and potential remedies to improve their robustness is missing despite its importance for improving their reliability in practice.

This work sets up an empirical benchmark using the training dataset BDAPPV¹ (Kasmi et al., 2023b) to disentangle the effects of distribution shifts on rooftop PV detectors. We then combine explainable artificial intelligence (XAI) methods to understand how these shifts affect the deep learning models. Based on our findings, we propose simple solutions that can be implemented to improve the robustness to distribution shifts of deep learning models trained to detect PV panels on overhead imagery. 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. Code for replicating the results of this paper can be found at https://github.com/gabrielkasmi/robust_pv_mapping and model weights can be found at <https://zenodo.org/records/12179554>.

2. Related works

2.1. Remote sensing of rooftop photovoltaic installations

The remote sensing of rooftop PV systems is now a well-established field with early works dating back to Golovko et al., 2018; Malof et al., 2015, 2016; Yuan et al., 2016. The DeepSolar project (Yu et al., 2018) marked a significant milestone by 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, covering areas such as North-Rhine Westphalia (Mayer et al., 2020), Switzerland (Casanova et al., 2021), Oldenburg in Germany (Zech & Ranalli, 2020), parts of Sweden (Frimane et al., 2023; Lindahl et al., 2023), Northern Italy (Arnaudo et al., 2023), the Netherlands (Kausika et al., 2021) or the surroundings of Berkeley in California (Parhar et al., 2021). Several works even included GIS data to construct registries of PV installations (Kasmi et al., 2022; Kausika et al., 2021; Mayer et al., 2022; Rausch et al., 2020). In the current context of rapid rooftop PV growth (Haegel et al., 2017; RTE France, 2022), remote sensing of rooftop PV installations using deep learning and orthoimagery offers the potential to address the lack of systematic registration of small-scale PV installations (Kasmi et al., 2022; Kausika, 2022).

¹BDAPPV: Base de données d'apprentissage profond pour les installations photovoltaïques (Database for deep learning applied to PV systems).

However, current methods cannot be transposed from one region to another without incurring accuracy drops, thus limiting their practical usability (Hu et al., 2022). Unpredictable accuracy drops cast doubt on the reliability of the generated data (De Jong et al., 2020). Recently, Kasmi et al., 2022 introduced a method aiming at indirectly assessing the accuracy of the detections by automatically comparing the registry generated with deep learning algorithms to reference data, which is often aggregated at the scale of the city. While this work enabled the quantification of the drop in accuracy encountered during deployment, no cues as to why the accuracy varied during deployment were discussed. In this work, we propose to study and mitigate the impact of these distribution shifts.

2.2. Sensitivity to distribution shifts

Sensitivity to distribution shifts in remote sensing. Distribution shifts, i.e., the sensitivity to the fact that "the training distribution differs from the test distribution" (Koh et al., 2021) are ubiquitous in machine learning (Torralba & Efros, 2011). The sensitivity to distribution shifts causes unpredictable performance drops, which can have dire consequences as models are deployed in safety-critical settings such as autonomous driving (Sun et al., 2022b) or medical diagnoses (Pooch et al., 2020).

In remote sensing applications, Tuia et al., 2016 identified two primary sources of shifts in the input data models are sensitive to: variations in the geographical scenery, and the varying acquisition conditions. Following Murray et al., 2019, we can add the ground sampling distance (GSD). The acquisition conditions encompass the conversion of a scene into a digital image and include all sources of variability in the input images caused by different sensors, exposure, attitude and altitude during acquisition, and atmospherical conditions. The ground sampling distance is the upper bound to the image resolution. The lower the ground sampling distance, the more detailed the image. In practice, the resolution is limited by the GSD and the image quality (noise, optical transfer function, intrinsic geometric consistency, etc). So far, the only work that investigated the poor reliability of deep learning systems applied to the remote sensing of PV panels is Wang et al., 2017. The authors argued that the generalization ability from one city to another depending on how "hard" to recognize the PV panels are. However, no proper definition of the "hardness" to recognize PV panels or a proper disentanglement the effect of each source of variability was carried out, and there was no prescription regarding model training or data preprocessing.

Mitigating the sensitivity to distribution shifts. Numerous approaches have been introduced to mitigate the sensitivity to distribution shifts. We refer the reader to surveys such as Csurka, 2017; Csurka et al., 2021; Guan and Liu, 2022; Tuia et al., 2016; Zhou et al., 2023 for reviews of these methods in various settings. One simple yet effective approach is to leverage data augmentations during training to incentivize the model to learn a given invariance during training. The aim is that the model is no longer sensitive to a given set of perturbations of the input images. Popular data augmentation methods consist in defining a method to generate as many perturbed samples as possible while preserving the semantic content of the image. To this end, AugMix (Hendrycks et al., 2020) applies a random sequence with random weights of perturbations to the input image. Similarly, Hendrycks et al., 2022 augment an input image with fractal patterns, and Sun et al., 2022a perturb the Fourier spectrum of the input image.

However, improving the robustness against distribution shifts is a long-tailed problem, meaning that unseen situations eventually arise, and all situations cannot be accounted for (Recht et al., 2019; Torralba & Efros, 2011). Therefore, we should first understand how distribution shifts affect a model's performance before implementing a mitigation method. To this end, we propose to leverage explainable artificial intelligence (XAI) methods.

2.3. Explainable artificial intelligence (XAI)

Modern deep learning algorithms are often qualified as black boxes, meaning it is hard to fully grasp their inner workings. This black-box nature limits the applicability of machine learning in safety-critical settings (Achtibat et al., 2022). We can distinguish two main approaches for machine learning explainability: by-design interpretable models and post-hoc explainability (Parekh, 2023)². By-design interpretability aims at constructing models that are transparent and self-explanatory (Sudjianto & Zhang, 2021), e.g., the decision boundaries of a decision tree. On the other hand, post-hoc explainability seeks to explain a model’s decision by highlighting important features contributing to this decision, without explicitly stating how these features affected the model. Methods such as class activation maps (CAMs, Zhang et al., 2017), which plot a heatmap of the important image regions for the classification of this image, fall into this category.

XAI methods for model debugging. One of the main motivations for XAI is to inspect the decision of models to assess whether they relied on relevant factors to make predictions. Several works highlighted biases in the decision process, such as the reliance on spurious features. Lapuschkin et al., 2019 leveraged the GradCAM (Selvaraju et al., 2017) to show how classifiers could rely on watermarks rather than relevant areas of the input image for horses classification, thus highlighting a so-called "Clever Hans" (Pfungst, 1911) effect³. CAMs (Zhang et al., 2017) have also been used to understand the behavior of convolutional neural networks (CNNs) in medical imagery classification by Zhang et al., 2021. Another example of usage of XAI tools to understand and debug a model was proposed by Dardouillet et al., 2023, who leveraged SHapley Additive exPlanations (SHAP, Lundberg and Lee, 2017) to understand a model deployed for oil slick pollution detection on the sea surface. In this work, we go one step further and show how combining post-hoc and by-design interpretable XAI methods can help understand and mitigate the sensitivity to distribution shifts of CNNs deployed for mapping PV systems from overhead imagery.

3. Data

We consider the crowdsourced training dataset BDAPPV⁴ (Kasmi et al., 2023b). 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 panels, but from two different image providers: images coming from the Google Earth Engine (hereafter referred to as "Google," Gorelick et al., 2017) and from the IGN (IGN, 2024), the French public operator for geographic information. We have double annotations for 7,686 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 (almost) 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 1 plots examples of images coming from BDAPPV. These images depict the same scene for both providers (although potentially at different dates).

²Flora et al., 2022 note that there is no consensus yet in the literature regarding the use of the terms explainability and interpretability. Following Flora et al., 2022, we say that a model is interpretable if it is *inherently* or *by design* interpretable and that a model is explainable if we can compute an *post-hoc* explanation of its decision.

³Clever Hans was a horse that appeared to perform arithmetic and other intellectual tasks but was actually responding to subtle cues from his handler. In machine learning, it is used as an example of the reliance of deep learning models on spurious features (e.g., background) rather than causal features (e.g., shape of an object) for tasks such as classification. See Lapuschkin et al., 2019 for more details.

⁴The dataset is accessible here: <https://zenodo.org/records/7358126>.



Figure 1. Examples of images of the same PV panels but with different providers and acquisition dates (Up Google, down: IGN).

4. Methods

We aim to explain why convolutional neural networks (CNNs) applied to detect PV panels on orthoimages are sensitive to distribution shifts. We first construct a benchmark to isolate the effect of the three main sources of distribution shifts on overhead images highlighted by Tuia et al., 2016 and Murray et al., 2019. On our benchmark dataset, we then compare the behavior of a by-design interpretable model, the Scattering transform (Bruna & Mallat, 2013), with the predictions of our CNN to understand on which scales the CNN relies and how it is affected by a disruption of these scales. To verify that our mechanism is correct, we leverage the wavelet scale attribution method (WCAM, Kasmir et al., 2023a), which is a post-hoc explainability method, to isolate the important scales in the predictions of our black-box CNN model. We chose the Scattering transform and the WCAM because these methods rely on a decomposition of the input image in the space-scale (wavelet domain), which is particularly well suited in the case of remote sensing images since the scales, expressed in pixels on images, are indexed in meters and can thus point towards actual elements depicted on the images⁵. Finally, based on our findings, we propose a data augmentation method to improve the robustness of CNNs and draw some lessons regarding the choice of image data.

4.1. Disentangling the sources of distribution shifts on overhead images

BDAPPV features images of the same installations from two providers and records the approximate 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 (He et al., 2016) on Google images downsampled at 20 cm/pixel of resolution and evaluate it on three datasets: a dataset with Google images at their native 10 cm/pixel resolution ("Google 10 cm/pixel"), the IGN images with a native 20 cm/pixel resolution ("IGN") and Google images downsampled at 20 cm/pixel located outside of France ("Google OOD⁶"). 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 2 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

⁵In appendix A, we provide further evidence of the limitation of "traditional" feature attribution methods for explaining the false detection of deep learning models in our use case.

⁶OOD: out-of-distribution.

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 2. Test images on which a model trained on Google images (downsampled to 20 cm/px of GSD, "Google baseline") is evaluated. "Google 10 cm/pixel" corresponds to the source Google image, before downsampling and evaluates the effect of varying ground sampling distances. "Google OOD" corresponds to Google images taken outside of France. "IGN" corresponds to images depicting the same installations as Google baseline but with a different provider.

4.2. Space-scale decomposition of a model's decision process

4.2.1. Background: the wavelet transform of an image

Motivation and definition. We propose to analyze the decision process of an off-the-shelf CNN model through the lenses of the space-scale or wavelet decomposition. Wavelets are a natural tool to decompose an image into scales while maintaining local analysis in space: they provide a single space-scale decomposition. As scales are indexed in terms of actual distances on the ground, we can directly identify the important objects that contribute to a model's decision by studying the important scales. Figure 3 illustrates the objects that can be found at different scales of an orthoimage.




Figure 3. Decomposition of a PV panel into scales.

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 dilations of this wavelet enable to scrutinize different scales while translations enable to scrutinize spatial location. In other words, scales correspond to different spatial frequency ranges or spectral domains.

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 (Mallat, 1999). Computing the multi-level decomposition of f requires applying Equation 2 J times with all dilated and translated wavelets of \mathcal{D} . Mallat, 1989 showed that one could implement the multi-level 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 multi-level 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 various rotations (usually horizontal, vertical, and diagonal) when dealing with 2D signals (e.g., images).

Interpreting the wavelet transform of an image. Figure 4 illustrates how to interpret the (two-level) wavelet transform of an image. Reading is the same for any multi-level decomposition. The right image plots the two-level dyadic decomposition of the original image depicted on the left. Following this transform, the localization on the image highlighted by the red polygon can be decomposed into six detail components (marked yellow and blue) and one approximation component (marked pink). The yellow components correspond to the details at the 1-2 pixel scale, and the blue components to the details at the 2-4 pixel scale. For each location, the wavelet transform summarizes the information contained in the image at this scale and location.




Figure 4. Image and associated two-level dyadic wavelet transform with indications to interpret the wavelet transform of the image. "Horizontal," "diagonal," and "vertical" indicate the direction of the detail coefficients. The direction is the same at all levels.

4.2.2. By design interpretable XAI method: the Scattering transform

The Scattering transform (Bruna & Mallat, 2013) is a deterministic feature extractor. CNNs and the Scattering transform share the same multi-level architecture, where the previous layer's output is passed onto the next after a nonlinearity is applied. The nonlinearities in a CNN are generally rectified linear units (ReLU), whereas in the Scattering transform, it is a modulus operation. Unlike CNNs, whose kernel coefficients are learned during training, the coefficients of the Scattering transform are fixed. Bruna and Mallat, 2013 showed that the Scattering transform computes representations from an input image that share the same properties of translational invariance as the representations computed with a CNN. The advantage of the Scattering transform is that as filters are fixed, we can know precisely what information they extract from the input image. Figure 5 summarizes the feature extraction process of the Scattering transform.

The input image x is downsampled, and a wavelet filter ϕ is applied in J directions. The wavelet coefficients at that scale are retrieved (black arrows), and the image is passed onto the next layer (blue arrows). As the depth increases, the spatial extent covered by the filters decreases. At each spatial location, one takes the modulus of the wavelet transform to compute a scale invariant representation which indicates the amount of "energy" in the image at this scale and localization.

The Scattering transform is parameterized by the number m of layers and the number J of orientations. We have a total of $mJ + m^2J(J - 1)/2$ coefficients. At the end of the decomposition, the features, i.e., the scattering coefficients, are flattened into a single vector of size $mJ + m^2J(J - 1)/2$. We can identify to which scale, location, and orientation on the input image this feature corresponds.




Figure 5. A scattering propagator U_J applied to x computes each $U[\lambda_1]x = |x \star \psi_{\lambda_1}|$ and outputs $S_J[\emptyset]x = x \star \phi_{2^j}$ (black arrow). Applying U_J to each $U[\lambda_1]x$ computes all $U[\lambda_1, \lambda_2]x$ and outputs $S_J[\lambda_1] = U[\lambda_1] \star \phi_{2^j}$ (black arrows). Applying U_J iteratively to each $U[p]x$ outputs $S_J[p]x = U[p]x \star \phi_{2^j}$ (black arrows) and computes the next path layer. Figure borrowed from Bruna and Mallat, 2013. Note: On the image, the input x corresponds to f and $\lambda = 2^j r$ is a frequency variable corresponding to the j^{th} scale with r rotations.

We implement three variants of the Scattering transform with depths m varying from one to three levels. Bruna and Mallat, 2013 stated that first-order coefficients were insufficient to discriminate between two very different images but that coefficients of order $m = 2$ could. We consider $J = 8$ orientations.

We stack the scattering coefficients into a vector of dimension $mJ + m^2J(J - 1)/2$, akin to the penultimate layer of a CNN. We train a linear classifier on this feature vector. Our implementation of the Scattering transform is based on the Python library Kymatio (Andreux et al., 2020).

4.2.3. Post-hoc XAI method: the Wavelet scale attribution method (WCAM)

Traditional feature attribution methods (Petsiuk et al., 2018; Selvaraju et al., 2020; Simonyan & Zisserman, 2015) highlight the important areas for the prediction of a classifier in the pixel (spatial) domain. The WCAM (Kasmi et al., 2023a) generalizes attribution to the wavelet (space-scale domain). The WCAM provides us with two pieces of information: where the model sees and what scale it sees at this location. The decomposition of the prediction in terms of scales points towards actual elements on the input image since on orthoimagery scales are indexed in meters. 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. We refer the reader to appendix B or to Kasmi et al., 2023a for more details on the computation of the WCAM.

Reading a WCAM. Figure 6 presents an example of an explanation computed using the WCAM. On the right panel, we can see the important areas in the model prediction highlighted in the wavelet domain. On the left panel, we can see the spatial localization of the important components. We can see two main spatial locations: the center of the image, which depicts the PV panel, and the bottom left, which depicts a pool. Disentangling the scales, we can see that the PV panel’s importance spreads across three scales (orange arrows), while the pool is only important at the 4-8 pixel scale. This underlines that the model focuses on the PV panel because it sees details ranging from small details in the PV modules to the cluster of modules.

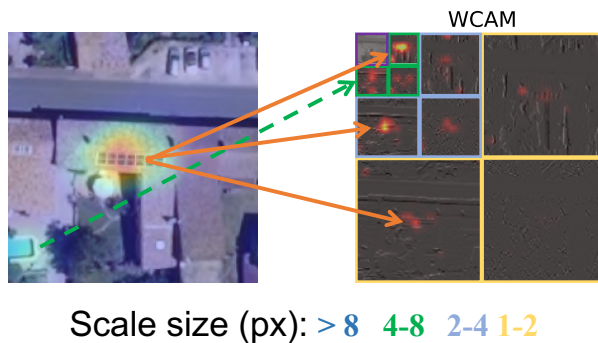


Figure 6. Decomposition in the wavelet domain of the important regions for a model’s prediction with the WCAM.

4.3. Improving the robustness through data augmentations

Improving the robustness to noise and scale distortions. Since we know that varying acquisition conditions induce perturbations (Lone & Siddiqui, 2018), which primarily affects high-frequency components (i.e., the finest scales), 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 at specific scales to force the model to learn potential distortions incurring at these scales. We refer the reader to the appendix D.1 for more details on the data augmentation

strategies and a review of the hyperparameters. Figure 10 in appendix D.2 illustrates the effect of the different data augmentation techniques. We compare this approach with a baseline without augmentations (labeled "ERM" for empirical risk minimization, Vapnik, 1999) and existing data augmentation techniques. We also include an "Oracle", i.e. an ERM trained on IGN images. The Oracle is a proxy for the best possible performance on the domain of interest.

Comparisons with existing approaches. The literature on robustness to image corruptions (Hendrycks & Dietterich, 2019) proposed numerous data augmentation methods to improve the robustness of classification models to image corruptions (Cubuk et al., 2019, 2020; Geirhos et al., 2019; Hendrycks et al., 2020, 2022). We consider the AugMix method (Hendrycks et al., 2020) and the recently-proposed RandAugment (Cubuk et al., 2020) and AutoAugment (Cubuk et al., 2019) 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). For comparison, we include an Oracle, a model trained on the target dataset (i.e., a model trained on IGN images). The Oracle is a proxy for the best performance that can be achieved on this dataset.

5. Results

5.1. Deep models are mostly sensitive to varying acquisition conditions, leading to an increase in the number of false negatives

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 comes from the change in the ground sampling distance, but the performance drop remains relatively small compared to the effect of the acquisition conditions (which can also be assimilated to variations in the image quality). 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.

Table 1. *F1 Score and decomposition in true positives, true negatives, false positives, and false negatives rates of the classification accuracy of a CNN model trained on Google images (Google baseline) and tested on the three instances of distributions shifts: GSD (Google 10 cm/px), the geographical variability (Google OOD) and the acquisition conditions (IGN). Values highlighted in red indicate the worst performance and values in orange the second-to-last worst performance.*

Model	F1 Score (↑)	True positive rate (↑)	True negative rate (↑)	False positive rate (↓)	False negative rate (↓)
Google baseline	0.98	0.99	0.98	0.02	0.01
Google 10 cm/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

5.2. The Scattering transform shows that clean, fine-scale features are transferable but poorly discriminative

Discriminative and transferable features. In the following, we distinguish between two kinds of features: the *discriminative* and the *transferable* features. Discriminative features enable the model to discriminate well between PV and non-PV images. Relying on discriminative features ensures a low

number of false positives. On the other hand, transferable features correspond to features that generalize well across domains. If a model relies on transferable features, its performance should remain even across domains. Ideally, we would like a model to rely on discriminative and transferable features. Analysis of the errors of the CNN using the Scattering transform highlights a potential trade-off between transferable and discriminative features.

Accuracy of the Scattering transform. Table 2 presents the accuracy results of the Scattering transform and compares it with a random classifier and the ERM (which is the same model as the one evaluated in Table 1). We can see that the performance on the source domain lags behind the performance of the CNN, but the Scattering transform generalizes better to IGN than the CNN. However, this comes at the cost of a high false positive rate. Table 5 in appendix E presents similar accuracy results for variants of the Scattering transform model in the depth and number of features.

Table 2. F1 Score and decomposition in true positives, true negatives, false positives and false negative rate of the classification accuracy of the Scattering Transform model trained on Google images and deployed on IGN images. Best results are **bolded, values in red highlight problematic cases.**

Model	Dataset	F1 Score (\uparrow)	True positive rate (\uparrow)	True negative rate (\uparrow)	False positive rate (\downarrow)	False negative rate (\downarrow)
Scattering transform	<i>Google baseline</i>	0.57	0,89	0,10	0,56	0,48
	IGN	0,59	0,54	0,31	0,62	0,54
CNN (ERM)	<i>Google baseline</i>	0,98	0,99	0,98	0,02	0,01
	IGN	0,46	0,32	0,95	0,03	0,68
Random classifier	<i>Google baseline</i>	0,47	0,5	0,50	0,55	0,45
	IGN	0,47	0,50	0,50	0,56	0,44

Implications for the CNN. We know which features the Scattering transform relies on. It leverages information at the two-pixel scale after downsampling the input image. In other words, the Scattering transform makes predictions based on *clean* features at the two-pixel scale. Therefore, we can deduce that these features are *transferable*, as the performance remains even across datasets, but not very *discriminative* as the false positives rate is high (across both domains).

On the other hand, the CNN should rely on *discriminative* features, which are located at coarser scales than 8 pixels, and on noisy features. In section 5.3, we investigate how the distortion of the input image’s coarse scales impacts the CNN’s decision process and the shift in its predicted probability. In section 5.4.1, we discuss how noise in input images affects the generalization ability of the CNN.

5.3. CNNs are sensitive to the distortion of coarse-scale discriminative features

Predicted probability shifts. The CNN outputs a predicted probability of a PV panel on the input image. When evaluating the CNN on the same scene coming from two providers, we compute *predicted probability shift* $\Delta p = |p_{ign} - p_{google}|$ when the model trained on Google is evaluated on IGN images. By construction, $\Delta p \in [0, 1]$. If $\Delta p = 0$, the predicted probability did not change when changing the provider. On the other hand, if $\Delta p \rightarrow 1$, then it means that the model made a different prediction solely because of the new acquisition condition.

Correlations between the probability shift and low-scale similarity of the images. For all images in our test set ($n = 4321$), we compute the similarity between the approximation coefficients at the 3rd level (corresponding to details above 8 pixels) and the predicted probability shift of the model. We expect that the greater the dissimilarity, the higher the shift in probability. We evaluate the dissimilarity between the low-scale details using the Structural similarity index measure (SSIM, Wang et al., 2004) and the

Euclidean distance between the images. The SSIM takes values between -1 and 1, where 1 indicates perfect similarity, 0 indicates no similarity, and -1 indicates perfect anti-correlation. On the other hand, the Euclidean distance takes positive values; the greater the distance, the greater the dissimilarity between the images.

As expected, we obtain a negative Pearson correlation coefficient equal to -0.41 (with a p -value $< 10^{-5}$) between the input images' SSIMs and the predicted probability shift. Using the Euclidean distance, we obtain a correlation coefficient of 0.250 ($p < 10^{-5}$). These results back the idea that the CNN is sensitive to low-scale perturbations of the input image, which results in a shift in the predicted probability.

Visualization of the model's response with the WCAM. The WCAM disentangles the important scales in a model's prediction. It enables us to see which scales were disrupted. On Figure 7, we present an example of an image that was initially identified as a PV panel but turned out to be no longer recognized.

We can see that in both cases, the approximation details are important in the model's prediction. The model responds to distortions at this scale by no longer focusing on a single area. Indeed, the model weights more components located at the 2-4 and 4-8 pixel scale (orange circles), which were not as important at the beginning. At the level of the perturbed scales, we can also witness that the model is disrupted by factors lying next to the PV panel (green circle). We supply more examples of such cases in the appendix F and discuss quantitative analysis of this result in appendix C.




Figure 7. Analysis with the WCAM of the CNNs prediction on an image no longer recognized as a PV panel.

5.4. Pathways towards improving the robustness to acquisition conditions

5.4.1. Blurring and wavelet perturbation improve accuracy

Table 3 reports the results of our data augmentation techniques and compares them with existing methods. We can see that 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 3, 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 (WP) improves the accuracy of the classification model without sacrificing the precision or the recall. While the drop in accuracy is still sizeable compared to the Oracle, the gain is consistent compared to other data augmentation techniques. Compared to RandAugment, the best-benchmarked method, our Blurring + WP is closer to the targets regarding true positives and true negatives and makes lower false negatives. This experiment shows that it is possible to consistently and reliably improve the robustness of acquisition conditions using a data augmentation technique, which does not leverage any information on the IGN dataset.

Table 3. F1 Score and decomposition in true positives, true negatives, false positives, and false negatives rate for models trained on Google with different mitigation strategies. Evaluation on IGN images. The oracle corresponds to a model trained on IGN images with standard augmentations. The best results are **bolded** and second best underlined. Values in red highlight problematic cases.

Model	F1 Score (↑)	True positive rate (↑)	True negative rate (↑)	False positive rate (↓)	False negative rate (↓)
Oracle	0.88	0.96	0.82	0.18	0.04
ERM	0.44	0.30	0.96	0.04	0.70
AutoAugment	0.46	0.31	0.96	0.04	0.69
AugMix	0.48	0.33	0.96	0.04	0.67
RandAugment	0.51	0.37	0.94	0.06	0.63
Blurring	0.74	0.98	0.49	0.51	0.02
Blurring + WP	<u>0.58</u>	0.47	0.87	0.13	0.53

5.4.2. On the role of the input data: towards some practical recommendations regarding the training data

Our results show that lowering the reliance on high-frequency content in the image improves generalization. This content is located on the 0.1-0.2 m scale and only appears on Google images. In Table 4, 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 4. F1 Score and true positives, true negatives, false positives, and false negatives rates. Evaluation computed on the Google dataset. ERM was trained on Google and Oracle on IGN images.

Model	F1 Score (↑)	True positive rate (↑)	True negative rate (↑)	False positive rate (↓)	False negative rate (↓)
ERM (Vapnik, 1999)	0.98	0.98	0.98	0.02	0.02
Oracle (ERM trained on IGN)	0.91	0.94	0.89	0.11	0.06

6. Discussion

6.1. Conclusion

This work is a comprehensive evaluation of the effects of distribution shifts on the classification accuracy of deep learning models trained to detect rooftop PV panels on overhead imagery. We first set up an experiment to disentangle the effects of the three primary shifts incurring in remote sensing (Murray et al., 2019; Tuia et al., 2016), namely geographical variability, varying acquisition conditions, and varying ground sampling distance. We show that the varying acquisition conditions contribute significantly to the observed performance drop. To explain why this drop occurs, we leverage space-scale analysis to disentangle the different scales from the input images. We combine two types of explainable AI methods grounded in the wavelet decomposition of the input images to show that the CNN relies on noisy features (at the finest scales) and features that are not very well transferable across domains (at the coarsest scales). We then introduce a data augmentation technique to improve the model’s robustness to distortions of the coarse-scale features and remove noise from the fine-scale features. We compare this method against various popular data augmentation techniques and show that our approach outperforms these baselines. We then discuss the practical takeaways of this work for the training or the choice of the training data for the initial training of the deep learning model.

Broader impact. Currently, transmission system operators (TSOs) lack quality data regarding rooftop PV installations (Kasmi et al., 2022). 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 (Pierro et al., 2022). On the other hand, current methods for mapping rooftop PV installations lack reliability, owing to their poor generalization abilities beyond their training dataset (De Jong et al., 2020). 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 and provides guidance on how to carry out the initial training of the model and how to deal with registry updates on newly available images.

6.2. Limitations and future works

Further discussion of the geographical variability. Our training data was limited to a narrow area around France. Therefore, we suspect the effect of the geographical variability to be underestimated. For instance, Freitas et al., 2023 showed that fine-tuning a model with data that is *not far* from the target area (e.g., France when the goal is to map PV systems in Portugal) enables accuracy gains compared to directly transferring a model trained over the United States. It could be interesting to study how the performance varies with the distance between the training data and the target mapping area once all other factors (acquisition conditions, ground sampling distance) are accounted for.

Extensions to other models. Over the last couple of years, foundation models (Bommasani et al., 2022) have been redefining the standards in deep learning. These very large models, trained on large data corpora, have shown remarkable performance for many challenging tasks, especially for text (Brown et al., 2020) and image (Rombach et al., 2022) generation. These models are used for more conventional and specialized tasks such as image segmentation (Kirillov et al., 2023) and achieve superior performance to conventional approaches while only requiring a few samples to learn their new task. Extending this benchmark and evaluating the performance of foundation models fine-tuned for segmenting PV panels such as Yang et al., 2024 under distribution shifts could be interesting.

A. Limitations of the GradCAM and related feature attribution methods for our use case

Figure 8 presents the explanations obtained using the GradCAM (Selvaraju et al., 2020). We can see two different prediction patterns depending on whether the model predicts a positive (true or false) or

a negative (true or false). In the case of a true positive prediction, the model will focus on a specific, narrow region of the image, which indeed corresponds to a PV panel. However, for false positives, the model also focuses on a narrow image region. Inspecting the samples of Figure 8 reveals that this region of the image depicts items that *resemble* PV panels. On the image on the first row (second column) of Figure 8, we can see that the model confuses a shadehouse that shares the same color and overall shape of a PV panel with an actual panel. In the image on the second row, the verandas with groves fool the model.

On the other hand, when the model does not see a PV panel, it does not focus on a specific image region. This remains true for the false negatives, where we can see that the model does not see the panels on any of the images.

However, we can also see that as the GradCAM only assesses where the model is looking, it is challenging to understand why it focused on a given area that resembles a PV panel on false positives and why it did not identify the PV panel on the false negatives. Achtab et al., 2022 underlines the necessity for reliable model evaluation to assess where models are looking at and *what* they are looking at on input images. The choice of the WCAM as an attribution method and, more broadly, the space-scale decomposition is an attempt to address this question by assessing the scales the models consider when making their predictions.

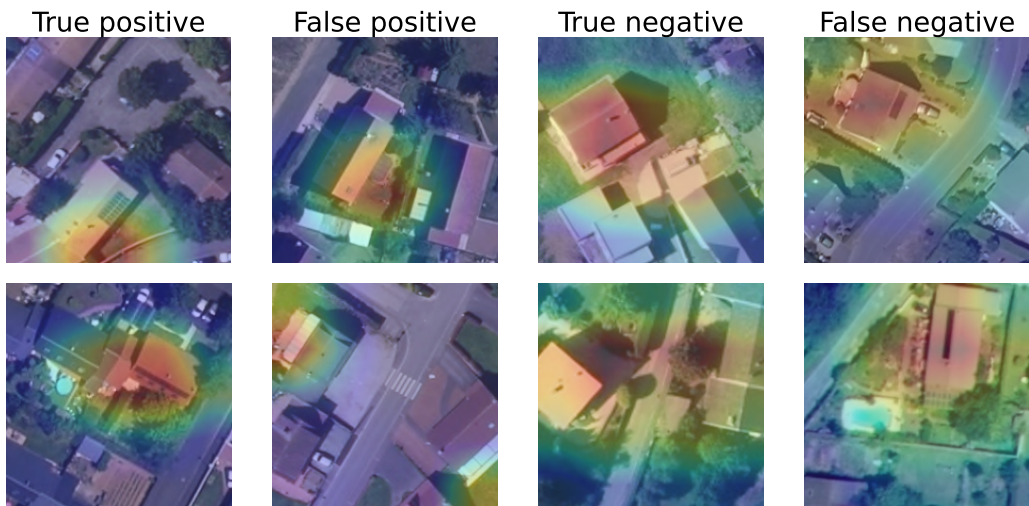


Figure 8. Model explanations using the GradCAM (Selvaraju et al., 2020) for some true positives, false positives, true negatives and false negatives. The redder, the higher the contribution of an image region to the predicted class (1 for true and false positives, 0 for true and false negatives).

B. Computation of the WCAM (from Kasmi et al., 2023a)

Figure 9 depicts the principle of the WCAM. The importance of the regions of the wavelet transform of the input image is estimated by (1) generating masks from a Quasi-Monte Carlo sequence, (2) evaluating the model on perturbed images. We obtain these images by computing the discrete wavelet transform (DWT) of the original image, applying the masks on the DWT to obtain perturbed DWT,⁷ and inverting the perturbed DWT to generate perturbed images. We generate $N(K+2)$ perturbed images

⁷On an RGB image, we apply the DWT channel-wise and apply the same perturbation to each channel.

for a single image. **(3)** We estimate the total Sobol indices of the perturbed regions of the wavelet transform using the masks and the model’s outputs using Jansen’s estimator (Jansen, 1999). Fel et al., 2021 introduced this approach to estimate the importance of image regions in the pixel space. We generalize it to the wavelet domain.

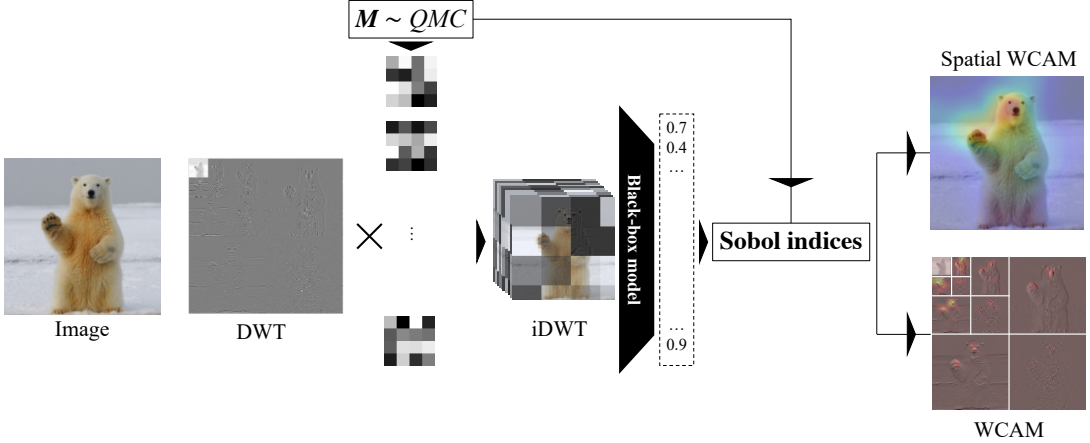


Figure 9. Flowchart of the wavelet scale attribution method (WCAM). Source: Kasmi et al., 2023a.

C. Quantitative relationship between the WCAM’s scale embeddings and the model’s response to distribution shifts

Definition. A *scale embedding* is a vector $z = (z_1, \dots, z_L) \in \mathbb{R}^L$ where each component z_s encodes the importance of the l^{th} scale component in the prediction.

Scale embeddings compute the importance of each scale and each direction and summarize it into a vector $z \in \mathbb{R}^L$ where L indicates the number of levels. In our case, we have ten levels (1 corresponding to the approximation coefficients and $L = 3 \times 3$ corresponding to the three scales of details coefficients and their three respective orientations). Scale embeddings summarize the importance of each scale, irrespective of the spatial localization of importance

Results. We computed the distance (measured by the Euclidean distance) between two images’ scale embeddings and computed the correlation between this distance and the predicted probability shift. As a baseline, we also computed the distance between the two WCAMs.

We obtained correlation coefficients of 0.18 ($p = 0.19$) for the scale embedding and 0.17 ($p = 0.19$) for the raw WCAM. Although weaker than the correlation between the distortion and the predicted probability shift, this result highlights that the WCAM consistently captures the change in behavior of the model resulting from the shift in acquisition conditions.

D. Data augmentation strategies

D.1. Description of the data augmentations

AugMix (Hendrycks et al., 2020). 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 $x_{aug} = \omega_1 op_1 \circ \dots \circ \omega_n op_n(x)$ 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 (Cubuk et al., 2019). 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 for the controller to generate better policies over time. The authors derive optimal augmentation strategies for various datasets, including ImageNet (Russakovsky et al., 2015), and show that the optimal policy for ImageNet generalizes well to other datasets.

RandAugment (Cubuk et al., 2020). 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 to discard the details at 0.1-0.2 m scale from the image. It corresponds to a blurring value $\sigma = 2$. in the `ImageFilter.GaussianBlur` method of the Python Imaging Library (PIL).

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) at targeted scales. 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. The wavelet perturbation aims to disrupt information at specific scales, as it can happen with varying acquisition conditions.

D.2. Plots

Figure 10 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 10. Visualization of the different data augmentation techniques implemented in this work.

E. Accuracy results for variants of the Scattering transform

Table 5 presents the accuracy of the Scattering transform for two depth variants (labeled $m = 1$ and $m = 2$). We can see that the performance of the Scattering transform remains relatively poor regardless of the depth of the scattering coefficients. Contrary to the claims of Bruna and Mallat, 2013, including second-order coefficients does not seem enough to discriminate between images, as the number of false positives remains high. This could be caused by the fact that our task, namely the detection of small objects on overhead imagery, is more challenging than digit classifications.

Table 5. F1 Score and decomposition in true positives, true negatives, false positives and false negative rate of the classification accuracy of the Scattering Transform model trained on Google images and deployed on IGN images.

Depth	Dataset	F1 Score (\uparrow)	True positive rate (\uparrow)	True negative rate (\uparrow)	False positive rate (\downarrow)	False negative rate (\downarrow)
$m = 1$	Google baseline	0.57	0,84	0,09	0,57	0,57
	IGN	0.57	0,71	0,40	0,52	0,36
$m = 2$	Google baseline	0.57	0,89	0,10	0,56	0,48
	IGN	0.59	0,54	0,31	0,62	0,54
ERM	Google baseline	0.98	0,99	0,98	0,02	0,01
	IGN	0.46	0,32	0,95	0,03	0,68
Random classifier	Google baseline	0.47	0,5	0,50	0,55	0,45
	IGN	0.47	0,50	0,50	0,56	0,44

F. Additional figures

Figure 11 to Figure 13 present additional examples of qualitative assessment of the effects of distribution shifts on the model’s prediction. On Figure 11, we can see that the model initially primarily relied on the gridded pattern, particularly visible at the 4-8 pixel scale. The acquisition conditions discarded this factor, thus explaining why the model could no longer recognize the PV panel. A similar phenomenon occurs on Figure 12. Figure 13 presents an example of a prediction not affected by the acquisition conditions. We can see that the important scales (especially at the 4-8 pixel scale) remain the same.

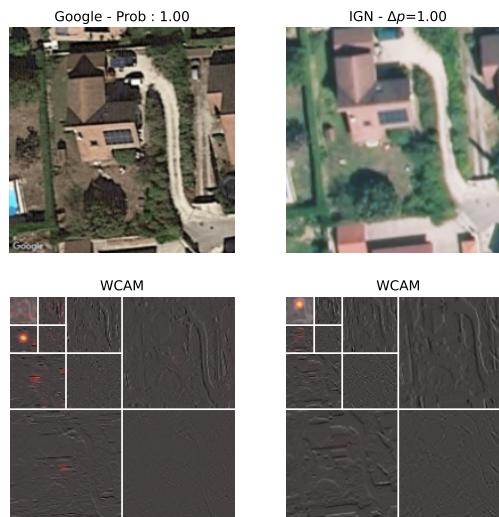


Figure 11. Analysis with the WCAM of the CNNs prediction on an image no longer recognized as a PV panel.




Figure 12. Analysis with the WCAM of the CNNs prediction on an image no longer recognized as a PV panel.




Figure 13. Analysis with the WCAM of the CNNs prediction on an image that is insensitive to varying acquisition conditions.

Funding Statement. This research was supported by a grant from the ANRT (CIFRE funding 2020/0685) and was funded by the French transmission system operator RTE.

Competing Interests. The authors declare no conflicts of interest.

Data Availability Statement. Code for replicating the results of this paper can be found at: https://github.com/gabrielkasmi/robust_pv_mapping. Model weights can be found at: <https://zenodo.org/records/12179554>.

Ethical Standards. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

Author Contributions. Conceptualization, Gabriel Kasmi; Formal analysis, Gabriel Kasmi; Funding acquisition, Laurent Dubus; Investigation, Gabriel Kasmi; Methodology, Gabriel Kasmi; Project administration, Laurent Dubus; Software, Gabriel Kasmi; Supervision, Philippe Blanc, Yves-Marie Saint-Drenan and Laurent Dubus; Validation, Gabriel Kasmi; Writing – original draft, Gabriel Kasmi; Writing – review & editing, Gabriel Kasmi, Philippe Blanc, Yves-Marie Saint-Drenan Laurent Dubus. All authors approved the final submitted draft

References

- Achtibat, R., Dreyer, M., Eisenbraun, I., Bosse, S., Wiegand, T., Samek, W., & Lapuschkin, S. (2022, June). From "Where" to "What": Towards Human-Understandable Explanations through Concept Relevance Propagation [arXiv:2206.03208 [cs]]. <https://doi.org/10.48550/arXiv.2206.03208>
- Andreux, M., Angles, T., Exarchakis, G., Leonarduzzi, R., Rochette, G., Thiry, L., Zarka, J., Mallat, S., Adén, J., Belilovsky, E., Bruna, J., Lostanlen, V., Chaudhary, M., Hirn, M. J., Oyallon, E., Zhang, S., Cella, C., & Eickenberg, M. (2020). Kymatio: Scattering Transforms in Python. *Journal of Machine Learning Research*, 21(60), 1–6. <http://jmlr.org/papers/v21/19-047.html>
- Arnaudo, E., Blanco, G., Monti, A., Bianco, G., Monaco, C., Pasquali, P., & Dominici, F. (2023). A Comparative Evaluation of Deep Learning Techniques for Photovoltaic Panel Detection from Aerial Images. *IEEE Access*, 1–1. <https://doi.org/10.1109/ACCESS.2023.3275435>
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., ... Liang, P. (2022, July). On the Opportunities and Risks of Foundation Models [arXiv:2108.07258 [cs]]. Retrieved December 16, 2023, from <http://arxiv.org/abs/2108.07258>
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., ... Amodei, D. (2020, July). Language Models are Few-Shot Learners [arXiv:2005.14165 [cs]]. <https://doi.org/10.48550/arXiv.2005.14165>
- Bruna, J., & Mallat, S. (2013). Invariant Scattering Convolution Networks [Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence]. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1872–1886. <https://doi.org/10.1109/TPAMI.2012.230>
- Casanova, A., Careil, M., Verbeek, J., Drozdal, M., & Romero Soriano, A. (2021). Instance-Conditioned GAN. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, & J. W. Vaughan (Eds.), *Advances in Neural Information Processing Systems* (pp. 27517–27529, Vol. 34). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2021/file/e7ac288b0f2d41445904d071ba37aaff-Paper.pdf
- Csurka, G. (2017). A Comprehensive Survey on Domain Adaptation for Visual Applications. In G. Csurka (Ed.), *Domain Adaptation in Computer Vision Applications* (pp. 1–35). Springer International Publishing. https://doi.org/10.1007/978-3-319-58347-1_1
- Csurka, G., Volpi, R., & Chidlovskii, B. (2021, December). Unsupervised Domain Adaptation for Semantic Image Segmentation: A Comprehensive Survey [arXiv:2112.03241 [cs]]. <https://doi.org/10.48550/arXiv.2112.03241>
- Cubuk, E. D., Zoph, B., Mane, D., Vasudevan, V., & Le, Q. V. (2019). AutoAugment: Learning Augmentation Strategies From Data. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Cubuk, E. D., Zoph, B., Shlens, J., & Le, Q. V. (2020). RandAugment: Practical automated data augmentation with a reduced search space [ISSN: 2160-7516]. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 3008–3017. <https://doi.org/10.1109/CVPRW50498.2020.00359>
- Dardouillet, P., Benoit, A., Amri, E., Bolon, P., Dubucq, D., & Credo, A. (2023). Explainability of Image Semantic Segmentation Through SHAP Values. In J.-J. Rousseau & B. Kapralos (Eds.), *Pattern Recognition, Computer Vision, and Image Processing. ICPR 2022 International Workshops and Challenges* (pp. 188–202). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-37731-0_19
- De Jong, T., Bromuri, S., Chang, X., Debusschere, M., Rosenski, N., Schartner, C., Strauch, K., Boehmer, M., & Curier, L. (2020). Monitoring Spatial Sustainable Development: Semi-automated analysis of Satellite and Aerial Images for Energy Transition and Sustainability Indicators. <https://doi.org/https://doi.org/10.48550/arXiv.2009.05738>

- Fel, T., Cadene, R., Chalvidal, M., Cord, M., Vigouroux, D., & Serre, T. (2021). Look at the Variance! Efficient Black-box Explanations with Sobol-based Sensitivity Analysis. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, & J. W. Vaughan (Eds.), *Advances in Neural Information Processing Systems* (pp. 26005–26014, Vol. 34). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2021/file/da94cbeff56cfd450785df477941308b-Paper.pdf
- Flora, M., Potvin, C., McGovern, A., & Handler, S. (2022, November). Comparing Explanation Methods for Traditional Machine Learning Models Part 1: An Overview of Current Methods and Quantifying Their Disagreement [arXiv:2211.08943 [physics, stat]]. Retrieved March 16, 2024, from <http://arxiv.org/abs/2211.08943>
- Freitas, S., Silva, M., Silva, E., Marceddu, A., Miccoli, M., Gnatyuk, P., Marangoni, L., & Amicone, A. (2023). An Artificial Intelligence-Based Framework to Accelerate Data-Driven Policies to Promote Solar Photovoltaics in Lisbon [_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/solr.202300597>]. *Solar RRL*, *n/a*(*n/a*), 2300597. <https://doi.org/10.1002/solr.202300597>
- Frimane, Å., Johansson, R., Munkhammar, J., Lingfors, D., & Lindahl, J. (2023). Identifying small decentralized solar systems in aerial images using deep learning. *Solar Energy*, *262*, 111822. <https://doi.org/10.1016/j.solener.2023.111822>
- Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F. A., & Brendel, W. (2019). ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. *International Conference on Learning Representations*. <https://openreview.net/forum?id=Bygh9j09KX>
- Golovko, V., Kroshchanka, A., Bezobrazov, S., Sachenko, A., Komar, M., & Novosad, O. (2018). Development of Solar Panels Detector. *2018 International Scientific-Practical Conference Problems of Infocommunications. Science and Technology (PIC S&T)*, 761–764. <https://doi.org/10.1109/INFOCOMMST.2018.8632132>
- Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., & Moore, R. (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone [Publisher: Elsevier]. *Remote sensing of Environment*, *202*, 18–27.
- Guan, H., & Liu, M. (2022). Domain Adaptation for Medical Image Analysis: A Survey [Conference Name: IEEE Transactions on Biomedical Engineering]. *IEEE Transactions on Biomedical Engineering*, *69*(3), 1173–1185. <https://doi.org/10.1109/TBME.2021.3117407>
- Haegel, N. M., Margolis, R., Buonassisi, T., Feldman, D., Froitzheim, A., Garabedian, R., Green, M., Glunz, S., Henning, H.-M., Holder, B., et al. (2017). Terawatt-scale photovoltaics: Trajectories and challenges [Publisher: American Association for the Advancement of Science]. *Science*, *356*(6334), 141–143.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hendrycks, D., & Dietterich, T. (2019). Benchmarking Neural Network Robustness to Common Corruptions and Perturbations. *International Conference on Learning Representations*. <https://openreview.net/forum?id=HJz6tiCqYm>
- Hendrycks, D., Mu, N., Cubuk, E. D., Zoph, B., Gilmer, J., & Lakshminarayanan, B. (2020). AugMix: A Simple Data Processing Method to Improve Robustness and Uncertainty. *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. <https://openreview.net/forum?id=S1gmrXHFvB>
- Hendrycks, D., Zou, A., Mazeika, M., Tang, L., Li, B., Song, D., & Steinhardt, J. (2022). PixMix: Dreamlike Pictures Comprehensively Improve Safety Measures [arXiv:2112.05135 [cs]]. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. <https://doi.org/10.48550/arXiv.2112.05135>
- Hu, W., Bradbury, K., Malof, J. M., Li, B., Huang, B., Streltsov, A., Sydney Fujita, K., & Hoen, B. (2022). What you get is not always what you see—pitfalls in solar array assessment using overhead imagery. *Applied Energy*, *327*, 120143. <https://doi.org/10.1016/j.apenergy.2022.120143>
- IGN. (2024). BD ORTHO® | Géoservices. Retrieved January 18, 2024, from <https://geoservices.ign.fr/bdortho>
- Jansen, M. J. W. (1999). Analysis of variance designs for model output. *Computer Physics Communications*, *117*(1), 35–43. [https://doi.org/10.1016/S0010-4655\(98\)00154-4](https://doi.org/10.1016/S0010-4655(98)00154-4)
- Kasmi, G., Dubus, L., Saint-Drenan, Y.-M., & Blanc, P. (2022). Towards Unsupervised Assessment with Open-Source Data of the Accuracy of Deep Learning-Based Distributed PV Mapping. In T. Corpetti, D. Ienco, R. Interdonato, M.-T. Pham, & S. Lefèvre (Eds.), *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 (Vol. 3343). CEUR-WS.org. <https://ceur-ws.org/Vol-3343/paper6.pdf>
- Kasmi, G., Dubus, L., Saint-Drenan, Y.-M., & Blanc, P. (2023a). Assessment of the Reliability of a Model’s Decision by Generalizing Attribution to the Wavelet Domain [arXiv:2305.14979 [cs, stat]]. *XAI in Action: Past, Present, and Future Applications workshop at NeurIPS 2023*. <https://doi.org/10.48550/arXiv.2305.14979>
- Kasmi, G., Saint-Drenan, Y.-M., Trebosch, D., Jolivet, R., Leloux, J., Sarr, B., & Dubus, L. (2023b). A crowdsourced dataset of aerial images with annotated solar photovoltaic arrays and installation metadata. *Scientific Data*, *10*(1), 59. <https://doi.org/10.1038/s41597-023-01951-4>
- Kausika, B. B. (2022, May). *GIS4PV: A technological impact assessment of the application of GIS for Photovoltaic Solar Energy*. [Doctoral dissertation, Utrecht University]. <https://doi.org/10.33540/1371>
- Kausika, B. B., Nijmeijer, D., Reimerink, I., Brouwer, P., & Liem, V. (2021). GeoAI for detection of solar photovoltaic installations in the Netherlands. *Energy and AI*, *6*, 100111. <https://doi.org/10.1016/j.egyai.2021.100111>
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., Dollár, P., & Girshick, R. (2023, April). Segment Anything [arXiv:2304.02643 [cs]]. <https://doi.org/10.48550/arXiv.2304.02643>

- Koh, P. W., Sagawa, S., Marklund, H., Xie, S. M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R. L., Gao, I., Lee, T., David, E., Stavness, I., Guo, W., Earnshaw, B., Haque, I., Beery, S. M., Leskovec, J., Kundaje, A., ... Liang, P. (2021). WILDS: A Benchmark of in-the-Wild Distribution Shifts [ISSN: 2640-3498]. *Proceedings of the 38th International Conference on Machine Learning*, 5637–5664. Retrieved May 11, 2024, from <https://proceedings.mlr.press/v139/koh21a.html>
- Lapuschkin, S., Wäldchen, S., Binder, A., Montavon, G., Samek, W., & Müller, K.-R. (2019). Unmasking clever hans predictors and assessing what machines really learn [Publisher: Nature Publishing Group]. *Nature communications*, 10(1), 1–8.
- Li, P., Zhang, H., Guo, Z., Lyu, S., Chen, J., Li, W., Song, X., Shibasaki, R., & Yan, J. (2021). Understanding rooftop PV panel semantic segmentation of satellite and aerial images for better using machine learning. *Advances in Applied Energy*, 4, 100057. <https://doi.org/10.1016/j.adapen.2021.100057>
- Lindahl, J., Johansson, R., & Lingfors, D. (2023). Mapping of decentralised photovoltaic and solar thermal systems by remote sensing aerial imagery and deep machine learning for statistic generation. *Energy and AI*, 100300. <https://doi.org/10.1016/j.egyai.2023.100300>
- Lone, A. H., & Siddiqui, A. N. (2018). Noise models in digital image processing [Publisher: Al-falah School of Engineering and Technology]. *Global Sci-Tech*, 10(2), 63–66.
- Lundberg, S. M., & Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems* (Vol. 30). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf
- Mallat, S. (1989). A theory for multiresolution signal decomposition: The wavelet representation [Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence]. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7), 674–693. <https://doi.org/10.1109/34.192463>
- Mallat, S. (1999). *A wavelet tour of signal processing*. Elsevier.
- Malof, J. M., Bradbury, K., Collins, L. M., & Newell, R. G. (2016). Automatic detection of solar photovoltaic arrays in high resolution aerial imagery. *Applied Energy*, 183, 229–240. <https://doi.org/10.1016/j.apenergy.2016.08.191>
- Malof, J. M., Rui Hou, Collins, L. M., Bradbury, K., & Newell, R. (2015). Automatic solar photovoltaic panel detection in satellite imagery. *2015 International Conference on Renewable Energy Research and Applications (ICRERA)*, 1428–1431. <https://doi.org/10.1109/ICRERA.2015.7418643>
- Mayer, K., Rausch, B., Arlt, M.-L., Gust, G., Wang, Z., Neumann, D., & Rajagopal, R. (2022). 3D-PV-Locator: Large-scale detection of rooftop-mounted photovoltaic systems in 3D. *Applied Energy*, 310, 118469. <https://doi.org/10.1016/j.apenergy.2021.118469>
- Mayer, K., Wang, Z., Arlt, M.-L., Neumann, D., & Rajagopal, R. (2020). DeepSolar for Germany: A deep learning framework for PV system mapping from aerial imagery. *2020 International Conference on Smart Energy Systems and Technologies (SEST)*, 1–6. <https://doi.org/10.1109/SEST48500.2020.9203258>
- Murray, J., Marcos, D., & Tuia, D. (2019). Zoom In, Zoom Out: Injecting Scale Invariance into Landuse Classification CNNs. *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, 5240–5243.
- Parekh, J. (2023, July). *Un cadre flexible pour l'apprentissage automatique interprétable : Application à la classification d'images et d'audio* [Theses]. Institut Polytechnique de Paris [Issue: 2023IPPAT032]. <https://theses.hal.science/tel-04214919>
- Parhar, P., Sawasaki, R., Todeschini, A., Reed, C., Vahabi, H., Nusaputra, N., & Vergara, F. (2021). HyperionSolarNet: Solar Panel Detection from Aerial Images. *NeurIPS 2021 Workshop on Tackling Climate Change with Machine Learning*. <https://www.climatechange.ai/papers/neurips2021/41>
- Petsiuk, V., Das, A., & Saenko, K. (2018). RISE: Randomized Input Sampling for Explanation of Black-box Models [arXiv:1806.07421 [cs]]. <https://doi.org/10.48550/arXiv.1806.07421>
- Pfungst, O. (1911). *Clever Hans:(the horse of Mr. Von Osten.) a contribution to experimental animal and human psychology*. Holt, Rinehart; Winston.
- Pierro, M., Liolli, F. R., Gentili, D., Petitta, M., Perez, R., Moser, D., & Cornaro, C. (2022). Impact of PV/Wind Forecast Accuracy and National Transmission Grid Reinforcement on the Italian Electric System. *Energies*, 15(23), 9086. <https://doi.org/10.3390/en15239086>
- Pooch, E. H. P., Ballester, P., & Barros, R. C. (2020). Can We Trust Deep Learning Based Diagnosis? The Impact of Domain Shift in Chest Radiograph Classification. In J. Petersen, R. San José Estépar, A. Schmidt-Richberg, S. Gerard, B. Lassen-Schmidt, C. Jacobs, R. Beichel, & K. Mori (Eds.), *Thoracic Image Analysis* (pp. 74–83). Springer International Publishing. https://doi.org/10.1007/978-3-030-62469-9_7
- Rausch, B., Mayer, K., Arlt, M.-L., Gust, G., Staudt, P., Weinhardt, C., Neumann, D., & Rajagopal, R. (2020). An Enriched Automated PV Registry: Combining Image Recognition and 3D Building Data. *NeurIPS 2020 Workshop on Tackling Climate Change with Machine Learning*. <https://www.climatechange.ai/papers/neurips2020/46>
- Recht, B., Roelofs, R., Schmidt, L., & Shankar, V. (2019). Do ImageNet Classifiers Generalize to ImageNet? [ISSN: 2640-3498]. *Proceedings of the 36th International Conference on Machine Learning*, 5389–5400. Retrieved May 13, 2024, from <https://proceedings.mlr.press/v97/recht19a.html>

- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-Resolution Image Synthesis With Latent Diffusion Models. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–10695.
- RTE France. (2022). *Energy Pathways to 2050* (tech. rep.). RTE France. Retrieved September 7, 2023, from <https://rte-futursenergetiques2050.com/>
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., & Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3), 211–252. <https://doi.org/10.1007/s11263-015-0816-y>
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization [ISSN: 2380-7504]. *2017 IEEE International Conference on Computer Vision (ICCV)*, 618–626. <https://doi.org/10.1109/ICCV.2017.74>
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2020). Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *International Journal of Computer Vision*, 128(2), 336–359. <https://doi.org/10.1007/s11263-019-01228-7>
- Simonyan, K., & Zisserman, A. (2015). Very Deep Convolutional Networks for Large-Scale Image Recognition. In Y. Bengio & Y. LeCun (Eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. <http://arxiv.org/abs/1409.1556>
- Sudjianto, A., & Zhang, A. (2021, November). Designing Inherently Interpretable Machine Learning Models [arXiv:2111.01743 [cs, stat]]. Retrieved June 17, 2024, from <http://arxiv.org/abs/2111.01743>
- Sun, J., Mehra, A., Kailkhura, B., Chen, P.-Y., Hendrycks, D., Hamm, J., & Mao, Z. M. (2022a). A Spectral View of Randomized Smoothing Under Common Corruptions: Benchmarking and Improving Certified Robustness [Series Title: Lecture Notes in Computer Science]. In S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, & T. Hassner (Eds.), *Computer Vision – ECCV 2022* (pp. 654–671, Vol. 13664). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-19772-7_38
- Sun, T., Segu, M., Postels, J., Wang, Y., Van Gool, L., Schiele, B., Tombari, F., & Yu, F. (2022b). SHIFT: A Synthetic Driving Dataset for Continuous Multi-Task Domain Adaptation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 21371–21382. Retrieved November 6, 2023, from https://openaccess.thecvf.com/content/CVPR2022/html/Sun_SHIFT_A_Synthetic_Driving_Dataset_for_Continuous_Multi-Task_Domain_Adaptation_CVPR_2022_paper.html
- Torralba, A., & Efros, A. A. (2011). Unbiased look at dataset bias [ISSN: 1063-6919]. *CVPR 2011*, 1521–1528. <https://doi.org/10.1109/CVPR.2011.5995347>
- Tuia, D., Persello, C., & Bruzzone, L. (2016). Domain Adaptation for the Classification of Remote Sensing Data: An Overview of Recent Advances [Conference Name: IEEE Geoscience and Remote Sensing Magazine]. *IEEE Geoscience and Remote Sensing Magazine*, 4(2), 41–57. <https://doi.org/10.1109/MGRS.2016.2548504>
- Vapnik, V. (1999). *The nature of statistical learning theory*. Springer science & business media.
- Wang, R., Camilo, J., Collins, L. M., Bradbury, K., & Malof, J. M. (2017). The poor generalization of deep convolutional networks to aerial imagery from new geographic locations: An empirical study with solar array detection. *2017 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*, 1–8. <https://doi.org/10.1109/AIPR.2017.8457965>
- Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: From error visibility to structural similarity [Publisher: IEEE]. *IEEE transactions on image processing*, 13(4), 600–612.
- Yang, R., He, G., Yin, R., Wang, G., Zhang, Z., Long, T., & Peng, Y. (2024). Weakly-semi supervised extraction of rooftop photovoltaics from high-resolution images based on segment anything model and class activation map. *Applied Energy*, 361, 122964. <https://doi.org/10.1016/j.apenergy.2024.122964>
- Yu, J., Wang, Z., Majumdar, A., & Rajagopal, R. (2018). DeepSolar: A Machine Learning Framework to Efficiently Construct a Solar Deployment Database in the United States. *Joule*, 2(12), 2605–2617. <https://doi.org/10.1016/j.joule.2018.11.021>
- Yuan, J., Yang, H.-H. L., Omataomu, O. A., & Bhaduri, B. L. (2016). Large-scale solar panel mapping from aerial images using deep convolutional networks. *2016 IEEE International Conference on Big Data (Big Data)*, 2703–2708. <https://doi.org/10.1109/BigData.2016.7840915>
- Zech, M., & Ranalli, J. (2020). Predicting PV Areas in Aerial Images with Deep Learning. *2020 47th IEEE Photovoltaic Specialists Conference (PVSC)*, 0767–0774. <https://doi.org/10.1109/PVSC45281.2020.9300636>
- Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyals, O. (2017). Understanding deep learning requires rethinking generalization. *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. <https://openreview.net/forum?id=Sy8gdB9xx>
- Zhang, Y., Hong, D., McClement, D., Oladosu, O., Pridham, G., & Slaney, G. (2021). Grad-CAM helps interpret the deep learning models trained to classify multiple sclerosis types using clinical brain magnetic resonance imaging. *Journal of Neuroscience Methods*, 353, 109098. <https://doi.org/10.1016/j.jneumeth.2021.109098>
- Zhou, K., Liu, Z., Qiao, Y., Xiang, T., & Loy, C. C. (2023). Domain Generalization: A Survey [Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence]. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4), 4396–4415. <https://doi.org/10.1109/TPAMI.2022.3195549>