

Distribution Shifts at Scale: Out-of-distribution Detection in Earth Observation

Burak Ekim*

University of the Bundeswehr Munich

Girmaw Abebe Tadesse†

Microsoft AI for Good Research Lab

Caleb Robinson†

Microsoft AI for Good Research Lab

Gilles Hacheme

Microsoft AI for Good Research Lab

Michael Schmitt

University of the Bundeswehr Munich

Rahul Dodhia

Microsoft AI for Good Research Lab

Juan M. Lavista Ferres

Microsoft AI for Good Research Lab

Abstract

Training robust deep learning models is crucial in Earth Observation, where globally deployed models often face distribution shifts that degrade performance, especially in low-data regions. Out-of-distribution (OOD) detection addresses this by identifying inputs that deviate from in-distribution (ID) data. However, existing methods either assume access to OOD data or compromise primary task performance, limiting real-world use. We introduce TARDIS, a post-hoc OOD detection method designed for scalable geospatial deployment. Our core innovation lies in generating surrogate distribution labels by leveraging ID data within the feature space. TARDIS takes a pre-trained model, ID data, and data from an unknown distribution (WILD), separates WILD into surrogate ID and OOD labels based on internal activations, and trains a binary classifier to detect distribution shifts. We validate on EuroSAT and xBD across 17 setups covering covariate and semantic shifts, showing near-upper-bound surrogate labeling performance in 13 cases and matching the performance of top post-hoc activation- and scoring-based methods. Finally, deploying TARDIS on Fields of the World reveals actionable insights into pre-trained model behavior at scale. The code is available at <https://github.com/microsoft/geospatial-ood-detection>

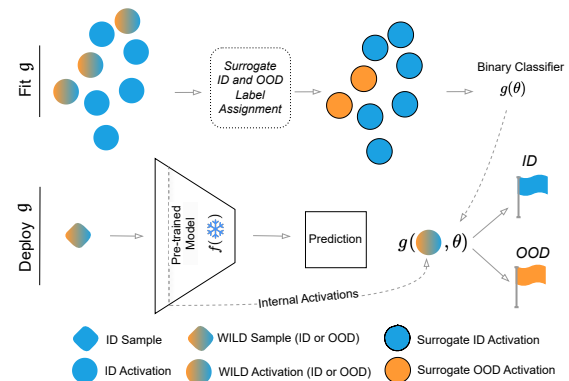


Figure 1. Overview of the proposed OOD detection method. Given a pre-trained model, ID samples, and WILD samples (from unknown distributions), TARDIS assigns surrogate ID/OOD labels to WILD samples using the ID set and fits a binary classifier g (top row). During deployment, g uses internal activations of unseen samples to predict whether they are ID or OOD (bottom row).

1. Introduction

Deep learning models have demonstrated remarkable capabilities across various domains but often exhibit overconfidence in their predictions, even when confronted with data that diverges from their training distribution [13, 24, 31]. This overconfidence arises from the assumption that inference data will follow the same independent and identically distributed (*i.i.d.*) properties as the training data. However, in real-world applications, this closed-world assumption [11, 18] is frequently violated by test-time distribu-

*Work done while a resident at the Microsoft AI for Good Research Lab. Contact: burak.ekim@unibw.de

†Residency supervisor.

tion shifts (e.g., lighting or angle variations, different devices or sensors) that can significantly degrade model performance and harm generalization. To address this, it is critical for predictive models to detect when new observations fall outside the training distribution, a task known as out-of-distribution (OOD) detection. Differentiating OOD from in-distribution (ID) samples is particularly challenging, partly due to the poor calibration of neural networks [9, 23]. This degradation poses a critical challenge, especially in applications where incorrect predictions can have severe consequences.

In Earth Observation (EO), specifically with satellite imagery, distribution shifts between training and inference commonly occur in the form of covariate and semantic shifts. Covariate shift refers to changes in the input distribution while keeping the task or label space fixed, whereas semantic shift involves changes in the label distribution or the meaning of inputs. A specific case of covariate shift is known as subpopulation shift [17], where the overall data distribution changes due to different proportions of underlying subgroups, even though the subgroups themselves remain the same. In EO, this arises when certain seasons, regions, or acquisition conditions are underrepresented during training but appear more prominently at test time. The magnitude of distribution shifts can further be characterized along a spectrum from near-distribution to far-distribution shifts. Near-distribution shifts may occur within the same satellite source due to temporal, geographic, or environmental variation—such as seasonal changes, sensor degradation, or atmospheric effects like cloud cover [16]. Far-distribution shifts involve more substantial differences, such as changes in satellite platforms, sensor types, or entirely different imagery domains, resulting in greater divergence from the training distribution. This variability necessitates tailored approaches for satellite-based machine learning models [26].

Despite the vital role of robust models in EO, few studies explore OOD detection in this domain [2, 3, 6, 7, 32]. These works often presume access to test-time distributions, adhere to closed-set assumptions, and/or degrade model performance on in-distribution tasks. This underscores the need for specialized OOD detection methods suited for large-scale EO deployment. A further challenge arises when applying geospatial models globally. Global geospatial models often suffer significant performance drops in low-data regions [1], such as sub-Saharan Africa [15]. These issues emphasize the need for global models equipped to detect OOD samples during inference.

In this paper, we build on well-established finding that OOD samples trigger internal activation patterns that diverge from ID samples [25, 30]. Building on this insight and recognizing the challenges of deploying geospatial models at scale, we introduce the following key contributions:

- We propose TARDIS (Test-time Addressing of Distribution Shifts at Scale), a lightweight post-hoc OOD detection method that preserves model performance and operates without access to labeled data from unknown test-time shifts.
- We evaluate TARDIS on EuroSAT (patch classification) and xBD (semantic segmentation) under near-distribution settings, analyzing covariate (geographical, temporal, environmental) and semantic (withheld class) shifts across 17 experimental setups. We also compare its performance against existing post-hoc OOD detection methods.
- We show that TARDIS scales effectively, providing actionable insights into the robustness and trustworthiness of geospatial models in real-world scenarios.

2. Related Works

Our method lies at the intersection of two notable categories of OOD detection methods: scoring functions and activation manipulation. Scoring functions assign a numerical score to each input, reflecting its alignment with ID samples based on the model’s output. Maximum Softmax Probability (MSP) [13] detects OOD samples by assessing the softmax confidence score, assuming that low confidence indicates OOD samples. The Out-of-Distribution detector for Neural networks (ODIN) [20] enhances MSP by applying input perturbations and temperature scaling to improve ID-OOD separation. The Mahalanobis score [19] calculates the distance between input features and class means in feature space, flagging inputs far from these means as OOD. The energy score [22] evaluates the model’s energy function to assess the likelihood of an input belonging to the ID distribution. Another category of OOD detection methods manipulates the internal activations of a pre-trained model to improve detection performance. ReAct [30] identifies differences in activation patterns of the penultimate layer between ID and OOD samples, enhancing separation by clipping activations at an upper limit. DICE [29] applies weight sparsification on a specific layer to further distinguish ID from OOD data and, when combined with ReAct, can enhance detection performance. Activation Shaping [4] prunes a portion of an input sample’s activations and slightly adjusts the remaining activations; when paired with the energy score, this approach has been shown to outperform contemporary OOD methods. Neuron Activation Patterns (NAP) [25] extracts, downsamples, and binarizes activation patterns from convolutional layers, computing the smallest Hamming distance between binarized test and training patterns. This distance provides a measure of model uncertainty, aiding in OOD detection. Both scoring functions and activation-based strategies face notable challenges in dynamic, real-world environments. Scoring functions, while maintaining ID accuracy, struggle with OOD detection due to their reliance on static data distributions,

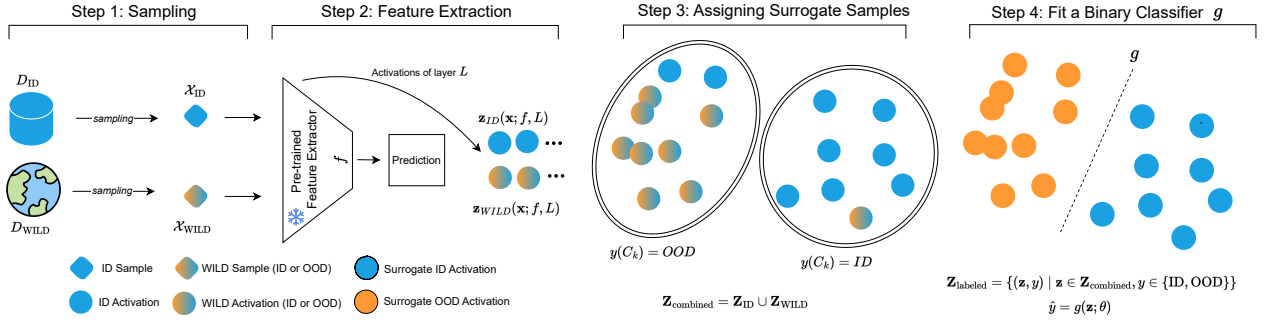


Figure 2. The proposed framework consists of four key steps: (1) Sampling in-distribution (ID) and WILD samples; (2) Extracting internal activations from a pre-trained model f for both ID and WILD samples; (3) Clustering the combined feature space and labeling WILD samples as surrogate-ID or surrogate-OOD; (4) Fitting a binary classifier g on the labeled feature representations to distinguish between ID and OOD samples. The classifier g , during deployment, flags out-of-distribution inputs.

an assumption rarely valid in practice. Activation-based methods, though effective, often compromise ID task performance and rely on hyperparameters tuned with hypothetical or substitute OOD samples. This reliance limits their effectiveness under evolving test-time conditions and shifting distributions.

Another line of OOD detection research uses clustering in feature space to separate ID and OOD samples. These methods exploit structural patterns in learned representations rather than direct scoring or activation manipulation. Gulati et al. [8] improve over k -means by using non-negative kernel regression for soft clustering. Sinhamahapatra et al. [27] explore clustering in self-supervised latent spaces, showing its potential in low-label settings. Closely related in spirit, NG-Mix [5] proposes pseudo-anomaly generation for semi-supervised anomaly detection.

The task of OOD detection is inherently challenging, and the EO domain is no exception due to the diverse and heterogeneous nature of satellite imagery. Despite the critical importance of OOD detection in EO, this area remains relatively underexplored, with only a limited number of studies addressing the issue. One such study develops a Dirichlet Prior Network to quantify distributional uncertainty in deep learning models for satellite image analysis [6]. This method assumes Dirichlet distributions for ID samples and employs various setups where classes, color channels, and environmental features are alternately treated as ID or OOD. Another study frames anomaly detection as a cumulative open-set detection and location separation task [3]. This approach also uses a Dirichlet prior, like the first study, to expand the representation space between ID (normal) and OOD (anomalous) samples by predicting the anticipated categorical distribution. The method suits scenarios where pre-event images are unavailable, subject to

radiation differences, or not recent enough to aid detection.

In this paper, we do not aim to optimize OOD detection performance but propose a method suited for real-world applications where the distribution is unknown, maintaining ID task performance is critical, and expensive methods are impractical. Our method detects shifts within the same satellite sources. This type of near-distribution shift presents a greater challenge than far-distribution shifts, such as those between natural and satellite imagery, as the differences are more subtle and harder to disentangle. By tackling these nuanced shifts, our method makes OOD detection more practical for real-world geospatial applications.

3. Problem Formulation

Let \mathcal{X} represent the set of all possible data the model may encounter, and \mathcal{Y} the set of class labels. The dataset on which a model is trained is defined as in-distribution (ID), denoted $D_{ID} \subset \mathcal{X}_{ID} \times \mathcal{Y}$. During inference, however, the model may encounter data from an unknown distribution, referred to as the WILD dataset $D_{WILD} \subset \mathcal{X}_{WILD}$, which may contain both ID and OOD samples.

We assume a pre-trained neural network $f : \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|}$, trained on D_{ID} , and our objective is to distinguish between ID and OOD samples within D_{WILD} . The network f extracts features $\mathbf{z} \in \mathbb{R}^F$, where F is the dimensionality of the feature space. Specifically, we fit a binary classifier $g : \mathbb{R}^F \rightarrow \{0, 1\}$, parametrized by θ , that operates on these features \mathbf{z} for a given sample in \mathcal{X}_{WILD} . We define $g(\mathbf{z}; \theta) = 0$ if $\mathbf{x} \sim D_{ID}$, and $g(\mathbf{z}; \theta) = 1$ if $\mathbf{x} \sim D_{OOD}$.

4. Method

Training a classifier to distinguish in-distribution (ID) from out-of-distribution (OOD) samples typically requires la-

Table 1. Overview of experimental setups for evaluating TARDIS under controlled distribution shifts. The table details the covariate and semantic shift scenarios, including the dataset used, split method, and the composition of the ID and out-of-distribution OOD sets. All conditions are post-event except where specified.

Type	Dataset	Split Method	ID (Training Set)	OOD (Test Set)
Covariate Shift	EuroSAT	Spatial Split (Longitude)	West Europe	East Europe
		Similar Disasters - Distant Locations	Nepal Flooding	Midwest Flooding
	xBD	Similar Disasters - Nearby Locations	Santa Rosa Wildfire	Woolsey Fire
		Different Disasters - Distant Locations	Hurricane Matthew	Nepal Flooding
		Different Disasters - Nearby Locations	Hurricane Matthew	Mexico Earthquake
		Temporal	Portugal Wildfire (Pre-Disaster)	Portugal Wildfire
Semantic Shift	EuroSAT	Unseen Class (Repeated for 10 Classes)	9 out of 10 classes (e.g., Annual Crop, Forest, Herbaceous Vegetation, Highway, Industrial, Pasture, Permanent Crop, Residential, River)	1 holdout class (e.g., Sea/Lake)

beled ID and OOD data, which is impractical in real-world settings. This motivates generating surrogate labels for samples with unknown distributions, referred to as WILD samples. Our method, illustrated in Figure 1, operates with a pre-trained model, known ID samples, and WILD samples. Using known ID samples, it assigns surrogate ID and OOD labels to WILD samples, which are then used to train a binary distribution classifier. TARDIS involves training a binary classifier to discriminate between ID and OOD samples based on internal activations, as shown in Figure 2. Given M ID and N WILD samples, we pass both through the frozen model and extract activations from a specified layer. After pooling for spatial downsampling, we obtain one-dimensional feature vectors, combine them, and assign corresponding ID or WILD labels. We shuffle this set and reserve 30% for validation. On the remaining 70%, we run k -means clustering to assign surrogate ID/OOD labels to WILD samples based on distance to ID samples. Using this labeled set, we train a binary classifier $g(\mathbf{z}; \theta)$ to map activations to distribution labels.

At deployment, we extract activations from a sample of unknown distribution and predict its label as $\hat{y} = g(\text{Downsample}(f(\mathbf{x})); \theta)$. The output can be interpreted as a probability of domain shift or thresholded for binary classification (0: ID, 1: OOD).

Surrogate Label Assignment. The core novelty of TARDIS lies in generating surrogate labels for WILD samples based on their proximity to known ID samples in fea-

ture space. We assume that ID and OOD features are separable enough for clustering to be effective. Specifically, we cluster the combined ID+WILD feature space into k groups (Step 3 in Figure 2). For each cluster, if the proportion of ID samples is threshold T , all samples in that cluster receive label 0 (ID); otherwise, label 1 (OOD). Distance to ID samples thus serves as a proxy for assigning surrogate labels.

Given k, T , and the feature vectors, we perform clustering and assign surrogate labels accordingly. To select the optimal parameters, we minimize the composite objective:

$$(k^*, T^*) = \arg \min_{k, T} H(S) + P_{\text{mis-ID}} - P_{\text{corr-ID}}$$

where $H(S)$ is the average cluster entropy (encouraging homogeneity), $P_{\text{mis-ID}}$ is the proportion of ID samples misclassified as OOD, and $P_{\text{corr-ID}}$ is the proportion of correctly classified ID samples. This balances three goals: coherent clusters, low ID misclassification, and high correct ID classification, promoting better surrogate label quality.

5. Experimental Setup

In this section, we describe experimental setups to evaluate TARDIS under controlled distribution shifts, mimicking real-world geospatial deployment challenges. We use the EuroSAT [12] dataset for patch-level classification and the

Table 2. Performance metrics comparing the oracle classifier g_{oracle} and the surrogate classifier g^* for EuroSAT and xBD datasets across various experimental setups. The g_{oracle} classifier acts as an upper bound for the g^* , which is the proposed method. The † notation indicates that over 10 measurements, the difference was not found to be statistically significant ($p < 0.05$). All conditions are post-event except where specified.

Type	Dataset	OOD (Test Set)	ID (Training Set)	AUROC↑		FPR95↓	
				g_{oracle}	g^*	g_{oracle}	g^*
Covariate Shift	EuroSAT	Eastern Europe	Western Europe	0.91 ± 0.04†	0.89 ± 0.06†	0.32 ± 0.11†	0.37 ± 0.10†
		Midwest Flooding	Nepal Flooding	1.00 ± 0.00†	1.00 ± 0.00†	0.01 ± 0.02†	0.01 ± 0.02†
	xBD	Woolsey Fire	Santa Rosa Wildfire	0.94 ± 0.02†	0.93 ± 0.02†	0.23 ± 0.08†	0.27 ± 0.09†
		Nepal Flooding	Hurricane Matthew	0.99 ± 0.01	0.98 ± 0.01	0.06 ± 0.04†	0.07 ± 0.02†
		Mexico Earthquake	Hurricane Matthew	0.96 ± 0.02†	0.94 ± 0.03†	0.11 ± 0.05†	0.19 ± 0.08†
		Portugal Wildfire (Pre)	Portugal Wildfire	0.99 ± 0.00	0.95 ± 0.01	0.06 ± 0.04	0.23 ± 0.20
		Santa Rosa Wildfire	Woolsey Fire	0.94 ± 0.02†	0.93 ± 0.02†	0.23 ± 0.08†	0.27 ± 0.09†
Semantic Shift	EuroSAT	Forest	Remaining 9 classes	0.99 ± 0.01†	0.97 ± 0.02†	0.08 ± 0.04†	0.08 ± 0.05†
		Herbaceous Vegetation		0.93 ± 0.03†	0.88 ± 0.07†	0.26 ± 0.11†	0.29 ± 0.14†
		Highway		0.63 ± 0.08†	0.56 ± 0.06†	0.78 ± 0.13†	0.88 ± 0.10†
		Industrial		0.97 ± 0.02†	0.96 ± 0.02†	0.13 ± 0.14†	0.22 ± 0.07†
		Pasture		0.98 ± 0.01	0.95 ± 0.01	0.09 ± 0.06†	0.28 ± 0.19†
		Permanent Crop		0.92 ± 0.03†	0.89 ± 0.03†	0.26 ± 0.08†	0.33 ± 0.11†
		Residential		0.87 ± 0.03†	0.82 ± 0.06†	0.46 ± 0.21†	0.43 ± 0.15†
		River		0.87 ± 0.04	0.81 ± 0.06	0.44 ± 0.14†	0.57 ± 0.10†
		Sea-Lake		1.00 ± 0.00†	1.00 ± 0.00†	0.00 ± 0.00†	0.00 ± 0.00†
		Annual Crop		0.92 ± 0.03†	0.91 ± 0.04†	0.25 ± 0.10†	0.32 ± 0.16†

xBD [10] dataset for semantic segmentation. EuroSAT contains 27,000 labeled images across ten land use/land cover classes. xBD includes pre/post-disaster imagery labeled for building damage assessment, covering hurricanes, floods, wildfires, and more. For our purposes, we formulate xBD as a building footprint segmentation task with two classes: building footprint and background. These datasets offer diverse evaluation conditions in terms of volume, task, sensor, spatial/temporal coverage, and geolocation, making them suitable test-beds for our OOD method.

To simulate controlled distribution shifts, we adjust the train, validation, and test splits in both datasets so that the model, trained on the modified training set, encounters realistic shifts during testing. We evaluate two shift types: covariate shift, where input distribution changes between train and test, and semantic shift, where unseen classes appear at test time. Covariate shifts include spatial (different locations), temporal (different time periods), and environmental (different landcover/disaster types) variance. For semantic shift, we train on nine of ten classes and test on the held-out class, rotating through all classes. Table 1 summarizes these setups.

These setups assume access to clear ID (train) and OOD (test) labels—an unrealistic assumption in real-world deployments. To address this, we evaluate our method in two ways. First, we use known ID and OOD labels to directly train and evaluate the binary classifier g , referred to as g_{oracle} . Second, we treat test-time OOD labels as unknown (i.e., WILD) and apply our surrogate label assignment to generate surrogate ID and OOD labels. This version, trained and evaluated using surrogate labels, is denoted as g^* . We consider g_{oracle} as an upper bound on performance, since it uses ground-truth labels, while g^* relies on inferred surrogates.

6. Experimental Results

In this section, we present ablation studies and experimental results. Note that ID task performance is not reported, as it remains unaffected by our method.

Ablation Studies. We conduct benchmark studies on the g_{oracle} classifier to explore the factors influencing TARDIS’s ability to detect test-time distribution shifts. This ability depends significantly on the layer from which internal ac-

tivations are extracted and the downsampling method applied. To address this, we test different layers and compare their performance (Table 4 and Table 5). For downsampling, we experiment with several methods, including mean and standard deviations, average pooling, max pooling, and PCA-based reduction. The max pooling-based downsampling method achieves the highest performance, likely due to its ability to retain the most salient activation patterns, which we argue is important for effective OOD detection (Table 6). We then evaluate various classifiers: KNeighbors, GaussianNB, DecisionTree, ExtraTrees, LogisticRegression, SVC, RandomForestUnbalanced, RandomForest, AdaBoost, and GradientBoosting. Results indicate that Logistic Regression provides the best tradeoff between classification performance and wall time (Table 7). We then compare g^* to g_{oracle} as we tune the parameters k and T required for surrogate label assignment. To determine k and T , we use a Tree-structured Parzen Estimator for sampling. Specifically, we set the search boundaries for k (the number of clusters) between 2 and $0.3 \times M$, where M is the total number of samples, and for T (the fraction of in-distribution samples within a cluster, below which the cluster is labeled as surrogate OOD) between 0.01 and 0.2. These values were empirically determined based on preliminary tuning for optimal performance. We then run 20 independent experiments and select the best-performing (k, T) pair. Fixing T based on the hyperparameter search results, we observe a recurring pattern that enables us to fix k to $0.3 \times M$ across all setups. This choice is based on the observation that the classifier g^* reaches performance levels close to g_{oracle} when k is set to 30% of the training samples. This trend holds for both datasets, with the xBD disaster scenario (semantic shift) and the EuroSAT Pasture experiment (covariate shift) showing similar behavior. In both cases, performance improves as more clusters are added, approaching that of g_{oracle} . In Figures 8 and 9, we show how we determine the optimal (k, T) pair that achieves high surrogate label assignment performance across all setups. The same pair generalizes to real-world deployment, removing the need for hyperparameter tuning.

Surrogate Sample Assignment Benchmark. Table 2 shows results for g^* and g_{oracle} across 17 setups. Overall, g^* closely approaches g_{oracle} , which serves as the upper bound. In 13 setups for AUROC and all 17 for FPR95, the performance gap is not statistically significant. In the few remaining cases, the gap is small. These results show that TARDIS can effectively assign surrogate labels without access to test-time labels.

ID/OOD Balance. As discussed in Section 4, splitting ID and WILD data into validation and clustering sets affects the ID/OOD ratio throughout the pipeline. In Table 2, WILD consists entirely of OOD samples, but in real deployments, it may contain both ID and OOD. While data is

shuffled, monitoring this ratio is key to understanding sensitivity. For example, in the *Pasture* semantic shift setup, the clustering and classifier stages see 17% and 15% OOD, respectively. In *Nepal Flood – Midwest Flood*, the ratios are 46% and 35%. This variation shows that TARDIS performs robustly across different ID/OOD balances, supporting the generalizability of our method.

Table 3. Performance comparison of post-hoc score-based and activation-based OOD detection methods across 10 semantic shift and 7 covariate shift experiments on EuroSAT and xBD. We report the mean \pm standard deviation across all 10 or 7 experiments for each method. TARDIS achieves comparable performance while requiring no hyperparameter tuning and preserving main task performance.

	Method	AUROC \uparrow	FPR95 \downarrow
Semantic Shift	Mahalanobis [19]	0.85 \pm 0.01	0.19 \pm 0.05
	MSP [13]	0.90 \pm 0.03	0.23 \pm 0.07
	Energy [21]	0.89 \pm 0.02	0.20 \pm 0.04
	ReAct [30]	0.92 \pm 0.04	0.13 \pm 0.01
	NAP [25]	0.93 \pm 0.03	0.21 \pm 0.05
	TARDIS (Ours)	0.95 \pm 0.02	0.16 \pm 0.03
Covariate Shift	Mahalanobis [19]	0.88 \pm 0.03	0.22 \pm 0.06
	MSP [13]	0.92 \pm 0.04	0.25 \pm 0.08
	Energy [21]	0.91 \pm 0.03	0.22 \pm 0.05
	ReAct [30]	0.94 \pm 0.05	0.15 \pm 0.02
	NAP [25]	0.95 \pm 0.02	0.23 \pm 0.06
	TARDIS (Ours)	0.98 \pm 0.01	0.11 \pm 0.04

Comparison with Existing OOD Detection Methods. We benchmark existing post-hoc activation-based and score-based methods across all semantic and covariate shift experiments, reporting mean and standard deviation (Table 3). Gradient-based methods involving back-propagation are excluded due to their high computational cost and increased wall time. Score-based methods (Mahalanobis, MSP, Energy) preserve ID performance as they require no model changes, but rely on potentially overconfident predictions and are less sensitive to near-domain shifts. Activation-based methods like ReAct and NAP modify internal activations—via clipping or purification—boosting OOD detection but at the cost of reduced ID performance, limiting their practicality.

Why Clustering Alone is Insufficient. Our method includes a distance-based clustering step followed by an additional classification step (Step 3 and 4 in Figure 2). Removing the classification step results in a drop in ROC AUC for the *River* experiment from 97% to 79%—a trend consistently observed across all experiments. This demonstrates that clustering alone (i.e., distance-based assessment) is insufficient for effective distribution detection, justifying the need for Step 4 in Figure 2.

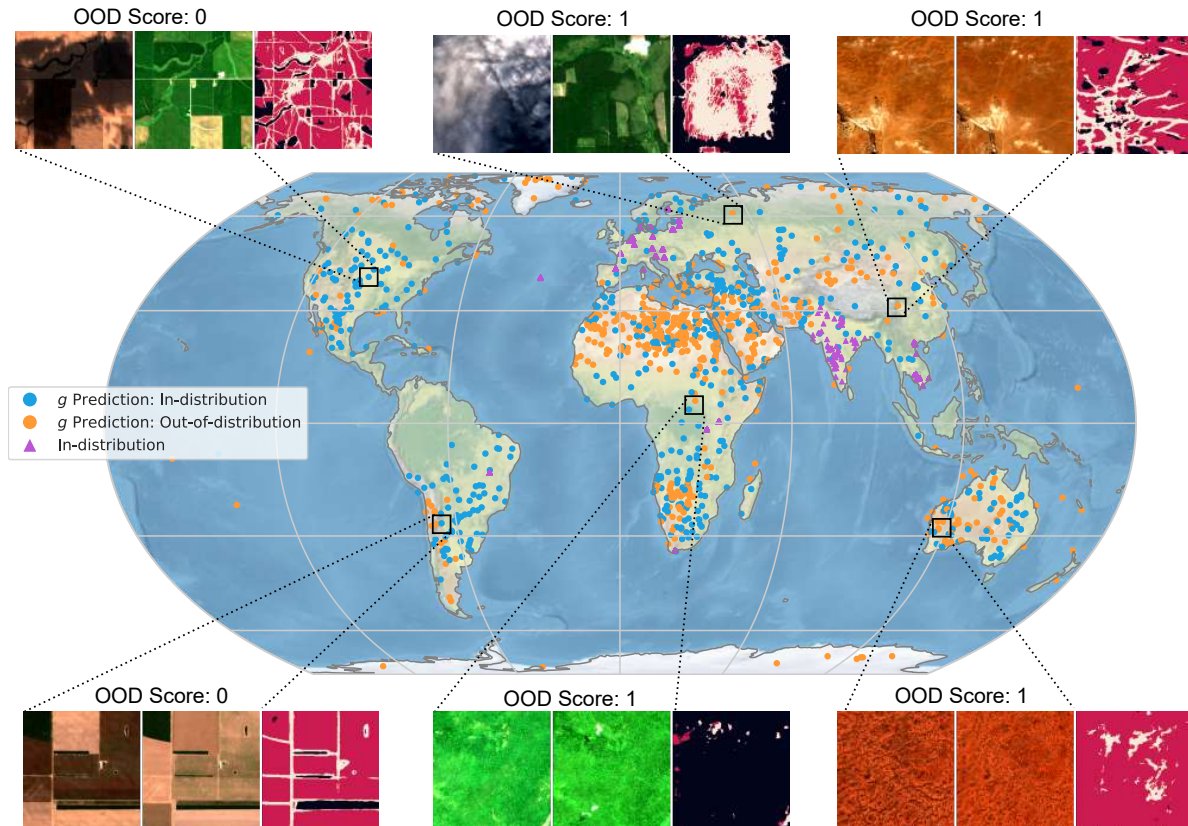


Figure 3. Geographical distribution of ID and WILD sets, containing 500 and 1200 samples, respectively. The ID set is sampled from the FTW dataset training set, while the WILD set is randomly sampled from the Microsoft Planetary Computer. Each Sentinel-2 patch, provided in two different time frames (planting and harvesting). The model f takes both Sentinel-2 images from different seasons as input, and the predictions are shown on the right. The OOD probability values are thresholded at 0.5.

7. OOD Detection Goes Global: Real-World Deployment

We use the Fields of the World (FTW) [14] training set and pre-trained models to demonstrate the capabilities of TARDIS in a large-scale deployment scenario. FTW is a geographically diverse dataset designed for agricultural field segmentation, covering 24 regions across Europe, Africa, Asia, and South America. The dataset contains approximately 70,000 samples, each consisting of multi-date, multi-spectral Sentinel-2 satellite patches paired with three-class semantic segmentation masks (field, field boundary, and background). The task involves segmenting these classes using a pair of Sentinel-2 images — one for the planting season and one for the harvesting season — as input. Field boundary data is crucial for global agricultural monitoring, however training large scale models to segment field boundaries from satellite imagery presents significant challenges due to the geographic diversity of fields, varying crop cycles, and agro-climatic conditions, all of which introduce substantial distribution shifts. This complexity,

along with corresponding distribution shifts found in real world imagery, makes the FTW dataset ideal for testing TARDIS’s ability to diagnose model performance during inference. Additionally, the multi-date nature of the dataset is particularly suitable for evaluating models that must handle spatiotemporal variations in satellite imagery.

Sampling ID Set. To form the ID set, we sample 50 patches from each country represented in the training set, ensuring a geographically diverse ID dataset that closely matches the data on which the model was originally trained.

Sampling WILD Set. We collect multispectral Sentinel-2 L2A satellite images using the Microsoft Planetary Computer [28]. The images are processed to Level-2A (bottom-of-atmosphere) and stored in cloud-optimized GeoTIFF (COG) format. We use four spectral bands: Red (B04), Green (B03), Blue (B02), and Near-Infrared (B08), each with a spatial resolution of 10 meters per pixel. We randomly query 1200 Sentinel-2 scenes from those over land and with cloud coverage $< 10\%$. Then, from each scene, we generate random patches of shape $256 \times 256 \times 4$ pixels (filtering out samples with $\geq 10\%$ zero or NaN pixels)

and query both the planting and harvesting seasons for the same locations. The planting and harvesting dates are fixed based on the hemisphere: April 1 to June 30 and September 1 to November 30 for the Northern Hemisphere, and October 1 to December 31 and March 1 to May 31 for the Southern Hemisphere. The two patches are concatenated along the channel dimension, resulting in a patch of shape $256 \times 256 \times 8$ pixels. Following the pre-trained model’s convention, the concatenated patch is upsampled to $512 \times 512 \times 8$, processed by the model, and then downsampled back to the original dimensions. For surrogate label assignment, we fix the hyperparameters k and T to 0.3×1200 and 0.1 , respectively, as suggested by the benchmark study.

The geographical distribution of ID and WILD samples is shown in Figure 3, with WILD samples further classified into surrogate ID and OOD. A clear pattern emerges: samples from arid biomes—such as the deserts of Inner Australia, the Sahara, and Patagonia—and polar regions, including Icelandic glaciers and the South Pole, are more frequently assigned as OOD. This is likely because the ID samples predominantly represent mesic environments: moderately moist, managed landscapes typical of agricultural areas. In contrast, arid and polar regions exhibit extreme environmental conditions that differ significantly from those in the ID set. This ecological dissimilarity likely drives their classification as surrogate OOD, reflecting the model’s sensitivity to environmental context and distribution shifts.

To evaluate scalability, we measure the time required to classify a WILD sample’s internal activation as surrogate ID or OOD using the logistic regression classifier g . A 256×256 input patch (10 m resolution, covering 6.5536 km^2) takes 0.003 seconds to process. Applied to mainland Africa ($29.77 \times 10^6 \text{ km}^2$), the method runs on a single GPU in under 4 hours—a fraction of the time needed for full f -model inference. This enables spatial diagnostics, identifying where the field boundary model is likely in- or out-of-distribution and where outputs may be unreliable. Such rapid screening is especially useful in low-data regions, where models are typically less robust and OOD detection can guide targeted interventions or data collection.

We assess the reliability of the OOD classifier g by comparing its output skewness on the training set to the performance of the model f on the FTW test set. High skewness indicates confident (mostly ID or OOD) predictions, while low skewness suggests uncertainty or misclassification. Plotting skewness against f ’s test performance reveals a trend: low skewness aligns with poor model performance (Figure 4), as seen in countries like Portugal, Cambodia, and Vietnam. This suggests that low skewness may signal the presence of OOD samples. Thus, TARDIS can also serve as a sanity check for detecting distribution shifts in test splits.




Figure 4. Correlation between model performance and OOD score skewness. Low performance of the f model on the FTW test set aligns with low skewness in the OOD classifier g ’s score distribution, suggesting the presence of OOD samples in the test set.

Limitations and Future Work. Our method assumes access to a non-trivial number of WILD samples and requires periodic updates to the clustering used for surrogate label assignment. Future work could explore density-based clustering methods instead of the current distance-based approach, as they may better handle variable sample densities and outliers, further improving surrogate label assignment.

8. Conclusion

We present TARDIS, a distribution shift detector that assigns surrogate ID and OOD labels to samples from unknown distributions by leveraging clustering in the feature space. These surrogate labels are assigned based on known ID samples and used to train a binary classifier that outputs a distribution score during inference. Rather than optimizing for OOD detection performance, TARDIS targets real-world scenarios where the distribution is unknown, ID task performance must be preserved, and computational overhead is a concern. Across 13 of 17 experiments mimicking real-world semantic and covariate distribution shifts, our method achieves near-upper-bound accuracy for surrogate label assignment while matching the performance of top post-hoc OOD detection methods. We demonstrate how TARDIS scales in a real-world application, providing interpretable insights into model behavior under distribution shifts and helping identify potential biases and limitations in the dataset. This contributes to the robustness and trustworthiness of models, making it particularly valuable in safety-critical, time-sensitive, and low-data settings.

References

- [1] Emily Aiken, Esther Rolf, and Joshua Blumenstock. Fairness and representation in satellite-based poverty maps: Evidence of urban-rural disparities and their impacts on downstream policy. *arXiv preprint arXiv:2305.01783*, 2023. 2
- [2] Georges Le Bellier and Nicolas Audebert. Detecting out-of-distribution earth observation images with diffusion models, 2024. 2
- [3] Mihai Coca, Iulia Coca Neagoe, and Mihai Datcu. Hybrid dnn-dirichlet anomaly detection and ranking: Case of burned areas discovery. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–16, 2022. 2, 3
- [4] Andrija Djuricic, Nebojsa Bozanic, Arjun Ashok, and Rosanne Liu. Extremely simple activation shaping for out-of-distribution detection. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [5] Hao Dong, Gaëtan Frusque, Yue Zhao, Eleni Chatzi, and Olga Fink. Nng-mix: Improving semi-supervised anomaly detection with pseudo-anomaly generation. *IEEE Transactions on Neural Networks and Learning Systems*, 2024. 3
- [6] Jakob Gawlikowski, Sudipan Saha, Anna Kruspe, and Xiao Xiang Zhu. Out-of-distribution detection in satellite image classification, 2021. 2, 3
- [7] Jakob Gawlikowski, Sudipan Saha, Julia Niebling, and Xiao Xiang Zhu. Handling unexpected inputs: incorporating source-wise out-of-distribution detection into SAR-optical data fusion for scene classification. *EURASIP J. Adv. Signal Process.*, 2023(1), 2023. 2
- [8] Aryan Gulati, Xingjian Dong, Carlos Hurtado, Sarath Shekizhar, Swabha Swayamdipta, and Antonio Ortega. Out-of-distribution detection through soft clustering with non-negative kernel regression. *arXiv preprint arXiv:2407.13141*, 2024. 3
- [9] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017. 2
- [10] Ritwik Gupta, Richard Hosfelt, Sandra Sajeev, Nirav Patel, Bryce Goodman, Jigar Doshi, Eric Heim, Howie Choset, and Matthew Gaston. xbd: A dataset for assessing building damage from satellite imagery, 2019. 5, 1
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015. 1
- [12] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification, 2019. 4, 1
- [13] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016. 1, 2, 6
- [14] Hannah Kerner, Snehal Chaudhari, Aninda Ghosh, Caleb Robinson, Adeel Ahmad, Eddie Choi, Nathan Jacobs, Chris Holmes, Matthias Mohr, Rahul Dodhia, Juan M. Lavista Ferres, and Jennifer Marcus. Fields of the world: A machine learning benchmark dataset for global agricultural field boundary segmentation, 2024. 7
- [15] Hannah Kerner, Catherine Nakalembe, Adam Yang, Ivan Zvonkov, Ryan McWeeny, Gabriel Tseng, and Inbal Becker-Reshef. How accurate are existing land cover maps for agriculture in sub-saharan africa? *Scientific Data*, 11(1):486, 2024. 2
- [16] Hannah Kerner, Saketh Sundar, and Mathan Satish. Multi-region transfer learning for segmentation of crop field boundaries in satellite images with limited labels. *arXiv preprint arXiv:2404.00179*, 2024. 2
- [17] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International conference on machine learning*, pages 5637–5664. PMLR, 2021. 2
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2012. 1
- [19] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks, 2018. 2, 6
- [20] Shiyu Liang, Yixuan Li, and R. Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks, 2020. 2
- [21] Weitang Liu, Xiaoyun Wang, John D. Owens, and Yixuan Li. Energy-based out-of-distribution detection, 2021. 6
- [22] Weitang Liu, Xiaoyun Wang, John D. Owens, and Yixuan Li. Energy-based out-of-distribution detection, 2021. 2
- [23] Matthias Minderer, Josip Djolonga, Rob Romijnders, Frances Hubis, Xiaohua Zhai, Neil Houlsby, Dustin Tran, and Mario Lucic. Revisiting the calibration of modern neural networks, 2021. 2
- [24] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 427–436, 2015. 1
- [25] Bartłomiej Olber, Krystian Radlak, Adam Popowicz, Michal Szczepankiewicz, and Krystian Chachuła. Detection of out-of-distribution samples using binary neuron activation patterns, 2023. 2, 6
- [26] Esther Rolf, Konstantin Klemmer, Caleb Robinson, and Hannah Kerner. Mission critical–satellite data is a distinct modality in machine learning. *arXiv preprint arXiv:2402.01444*, 2024. 2
- [27] Poulami Sinhamahapatra, Rajat Koner, Karsten Roscher, and Stephan Günnemann. Is it all a cluster game?—exploring out-of-distribution detection based on clustering in the embedding space. *arXiv preprint arXiv:2203.08549*, 2022. 3
- [28] Microsoft Open Source, Matt McFarland, Rob Emanuele, Dan Morris, and Tom Augspurger. microsoft/planetarycomputer: October 2022, 2022. 7
- [29] Yiyu Sun and Yixuan Li. Dice: Leveraging sparsification for out-of-distribution detection, 2022. 2

- [30] Yiyu Sun, Chuan Guo, and Yixuan Li. React: Out-of-distribution detection with rectified activations, 2021. [2](#), [6](#)
- [31] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. [1](#)
- [32] Sang Michael Xie, Ananya Kumar, Robbie Jones, Fereshte Khani, Tengyu Ma, and Percy Liang. In-n-out: Pre-training and self-training using auxiliary information for out-of-distribution robustness. *arXiv preprint arXiv:2012.04550*, 2020. [2](#)