

UC Berkeley

UC Berkeley Previously Published Works

Title

Understanding the Dependence of Perception Model Competency on Regions in an Image

Permalink

<https://escholarship.org/uc/item/0vh924tz>

ISBN

978-3-031-63796-4

Authors

Pohland, Sara

Tomlin, Claire

Publication Date

2024

DOI

10.1007/978-3-031-63797-1_8

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial License, available at <https://creativecommons.org/licenses/by-nc/4.0/>

Peer reviewed

Understanding the Dependence of Perception Model Competency on Regions in an Image

Sara Pohland^[0000-0003-2746-6372] and Claire Tomlin^[0000-0003-3192-3185]

University of California, Berkeley CA 94720, USA
{spohland,tomlin}@berkeley.edu

Abstract. While deep neural network (DNN)-based perception models are useful for many applications, these models are black boxes and their outputs are not yet well understood. To confidently enable a real-world, decision-making system to utilize such a perception model without human intervention, we must enable the system to reason about the perception model’s level of competency and respond appropriately when the model is incompetent. In order for the system to make an intelligent decision about the appropriate action when the model is incompetent, it would be useful for the system to understand *why* the model is incompetent. We explore five novel methods for identifying regions in the input image contributing to low model competency, which we refer to as image cropping, segment masking, pixel perturbation, competency gradients, and reconstruction loss. We assess the ability of these five methods to identify unfamiliar objects, recognize regions associated with unseen classes, and identify unexplored areas in an environment. We find that the competency gradients and reconstruction loss methods show great promise in identifying regions associated with low model competency, particularly when aspects of the image that are unfamiliar to the perception model are causing this reduction in competency. Both of these methods boast low computation times and high levels of accuracy in detecting image regions that are unfamiliar to the model, allowing them to provide potential utility in decision-making pipelines¹.

Keywords: Model Competency · Saliency Maps · Explainability.

1 Importance of Understanding Model Competency

Deep neural networks (DNNs) have been very successful in image classification tasks. However, DNN models are considered black boxes, and there currently exists only a shallow understanding of how these models obtain their outputs [1]. This greatly limits our ability to trust these models in real-world systems, where their failure could be detrimental without human intervention. To confidently employ DNN-based perception models in the real world, we need to develop more robust and generalizable models, but we also need to gain a deeper understanding

¹ The code for reproducing our methods and results is available on GitHub: <https://github.com/sarapohland/explainable-competency>.

of how these models work and when they might fail. There has been much work on quantifying uncertainty in perception systems [2], as well as identifying when an input is out-of-distribution (OOD) [3]. However, with only knowledge that the model is uncertain or that an input is OOD, the systems that utilize these perception models are limited in their ability to appropriately respond when the perception model lacks confidence. It would be useful to understand *why* a model is believed to be uncertain and *why* an input is deemed to be OOD before the output of a perception model is used for decision-making.

We explore these ideas in this paper, developing methods to identify particular regions in an image that lead to a lack of model confidence. Towards this end, we explore five approaches for identifying regions in an image contributing to low model competency. We begin by considering a naive approach, in which images are partitioned into regions and the competency associated with each cropped region is assessed. We then explore segment masking, pixel perturbation, and competency gradient approaches that are related to methods present in artificial intelligence (AI) explainability work. Finally, we present a novel method for understanding regions contributing to model incompetency that is based on autoencoder reconstruction loss. We develop methods for evaluating each of these approaches and compare their performance across three diverse datasets. We present an analysis of the results and conclude with avenues for future work.

2 Background & Related Work

In this work, we explore explanations for *model competency*—a generalized form of predictive uncertainty. Predictive uncertainty generally arises from three factors: data (or aleatoric) uncertainty, model (or epistemic) uncertainty, and distributional shift [4]. Data uncertainty refers to uncertainty arising from complexities of the data (i.e., noise, class overlap, etc.), while model uncertainty reflects the ability of the perception model to capture the true underlying model of the data [5]. These types of uncertainty are explored in work referred to as uncertainty quantification (Section 2.1). Distributional shift refers to mismatched training and test distributions [6]. Extensive work has explored methods to detect inputs that are out-of-distribution (OOD) (Section 2.2). This paper aims to expand upon work on uncertainty quantification and OOD detection with the goal of offering explanations for why a perception model’s prediction is uncertain. This idea of explaining perception model competency is related to previous work within the area of explainable image classification (Section 2.3), which seeks to offer explanations for the classification decisions of perception models. Our work differs from existing work in this area in that we seek to offer explanations for low model competency, rather than explanations for a model’s prediction.

2.1 Uncertainty Quantification

Extensive research has been done on methods to understand and quantify uncertainty in a neural network’s prediction. The modeling of these uncertainties can

generally be divided into methods based on (1) deterministic neural networks, (2) Bayesian neural networks (BNNs), and (3) ensembles of neural networks [2]. Approaches that use a single deterministic neural network estimate model uncertainty based on a single forward pass within a deterministic network. These approaches are often either distance-based [7, 8] or gradient-based [9, 10], although many other approaches exist [2]. A BNN is a stochastic model whose output is a probability distribution over its predictions [11, 12]. Approaches that employ BNNs extract uncertainty as a statistical measure, often using sampling methods [13], variational inference [14–16], or Laplace approximation [17]. Ensemble methods combine the predictions of multiple deterministic networks to form a probability density function [18]. While all of these methods provide an estimate of the uncertainty associated with a model’s prediction, none of these offer rationale for *why* the model is uncertain.

2.2 Out-of-Distribution (OOD) Detection

Many recent approaches have focused on quantifying distributional uncertainty—uncertainty that is caused by a change in the input data distribution [6]. Approaches that are specifically focused on determining if an input falls outside of the input-data distribution are referred to as OOD detection methods [3]. These approaches are generally either (1) classification-based, (2) density-based, (3) distance-based, or (4) reconstruction-based. Classification-based methods is a broad category that captures methods that largely rely on classifier predictions. These methods generally seek to revise the overly confident softmax scores at the output of neural networks to detect OOD samples more robustly [19–21]. Density-based methods model the training distribution with some probabilistic model and flag test data in low-density regions as OOD. These approaches tend to rely on Gaussian mixture models (GMMs) [22], kernel density estimation (KDE) [23], normalizing flows [24, 25], or likelihood ratios [26]. Distance-based methods use distance metrics in the feature space with the assumption that OOD samples should lie relatively far from the centroids or prototypes of the training classes. Many popular approaches utilize the Mahalanobis distance [27, 28], cosine similarity [29], or a nearest-neighbor distance [30]. Finally, reconstruction-based methods rely on the reconstruction loss of autoencoders (AEs) [31, 32], variational autoencoders (VAEs) [33], convolutional autoencoders (CAEs) [34], or generative adversarial networks (GANs) [35–37]. While all these methods have been used to successfully identify OOD samples, they do not offer a clear explanation as to *why* the samples are deemed OOD.

2.3 Explainable Image Classification

The field of explainable artificial intelligence (XAI) seeks to offer explanations for why a model makes the decisions that it does with the goal of developing more interpretable and trustworthy models [1, 38]. Within the area of explainable image classification, there exist (1) posthoc techniques, which explain the

behavior of an AI system after it has been trained without modifying the underlying model architecture, and (2) *antehoc* methods, which incorporate the aspect of explainability into the model design process to generate classification models that are inherently more interpretable [39]. Our work relates more closely to *posthoc* techniques, which employ various methods to understand the workings of black-box models and explain their decision-making process.

The most common *posthoc* approach for explaining an image classifier is to identify key image regions contributing to the model’s predictions and display them using a saliency map. These methods display the regions salient toward the prediction of a class through gradient-based approaches [40–42], perturbation-based approaches [43–45], layerwise-relevance propagation [46–48], masking-based approaches [49], or deconvolution methods [50]. Other approaches aim to construct an inherently interpretable pseudo-classifier that approximates the working mechanism of the black box classifier, offering either local interpretations around a particular instance [51,52] or global interpretations [53,54]. There are also approaches that provide explanations based on abstract vector representations, which can be linked to human-understandable concepts. These methods tend to use concept bottleneck models [55] or concept activation vectors [56] if the categories of interest are provided, or they focus on discovering task-relevant concepts automatically [57, 58]. Finally, there are methods that offer counterfactual explanations by generating alternative scenarios to explain the behavior of the classifier [59]. While all of these methods enhance the understanding of image classifiers, none of them explicitly deals with model *competency*. These methods do not offer an estimate of or explanation for how confident a model is in its prediction (as we demonstrate in Appendix B).

2.4 Explainable Competency Estimation

While there has been work on quantifying model uncertainty, identifying OOD inputs, and explaining model predictions, to the best of our knowledge, there has not been an effort to offer explanations for the lack of competency associated with a model’s prediction. We seek to bridge this gap by bringing together work on uncertainty quantification, OOD detection, and explainability. We see this as an important step in developing decision-making systems that can reason about model competency and respond safely in the face of uncertainty.

3 Approach for Understanding Model Competency

Recall that *model competency* is a generalized form of predictive uncertainty that captures data (or aleatoric) uncertainty, model (or epistemic) uncertainty, and distributional shift [4]. To gain a deeper understanding of perception model competency, we explore five methods for visualizing regions in the input image that contribute most to a lack of competency (similar to saliency methods discussed in Section 2.3). Towards this end, we first compute a competency score, estimating the confidence in a model’s prediction for a given input image (Section 3.1).

If the model is deemed incompetent (i.e., the competency score is below some threshold), we seek to understand which regions in the image most lead to this lack of competency by estimating the dependence of the competency score on segmented regions in the image (Section 3.2).

3.1 Estimating Model Competency

Let f be the true underlying model from which our images are drawn and \hat{f} be the predicted model (or the perception model). For an input image, \mathbf{X} , the perception model aims to estimate the true class of the image, $f(\mathbf{X})$, from the set of all classes. The competency of the model for this image is given by

$$C(\mathbf{X}) := P(\{\hat{f}(\mathbf{X}) = f(\mathbf{X})\}|\mathbf{X}). \quad (1)$$

To simplify our notation, let \hat{c} be the class predicted by the deterministic perception model (i.e., $\hat{f}(\mathbf{X}) = \hat{c}$), so we can express the model competency as

$$C(\mathbf{X}) = P(\{f(\mathbf{X}) = \hat{c}\}|\mathbf{X}). \quad (2)$$

It is difficult to estimate this probability directly because we are limited by the data contained in the training sample. Let D be the event that the input image is in-distribution (i.e., drawn from the same distribution as the training samples) and write the following lower bound on competency:

$$C(\mathbf{X}) \geq P(D \cap \{f(\mathbf{X}) = \hat{c}\}|\mathbf{X}) \quad (3)$$

$$= P(\{f(\mathbf{X}) = \hat{c}\}|\mathbf{X}, D)P(D|\mathbf{X}). \quad (4)$$

These probabilities are not readily available but can be approximated [4]. To estimate the first probability, we can fit a calibrated transfer classifier, such as the logistic regression classifier, and use its class probability output to obtain an estimate, $\hat{p}_{\hat{c}|D}(\mathbf{X})$. To estimate the probability that an image is from a class distribution, we can model the training distribution as a Gaussian mixture model (GMM) and compute the Mahalanobis distance of the feature vector to each of the class distributions. We then use these distances as input to another logistic regression classifier to estimate the probability that the input image came from the training distribution, $\hat{p}_D(\mathbf{X})$. Using these approximations, the lower bound on the probability of model competency can be approximated as

$$C(\mathbf{X}) \geq \hat{p}_{\hat{c}|D}(\mathbf{X})\hat{p}_D(\mathbf{X}). \quad (5)$$

We refer to this expression as the *competency score*. This metric ties together work on uncertainty estimation (Section 2.1) and OOD detection (Section 2.2), using a deterministic classification-based approach to estimate model uncertainty assuming a sample is in-distribution and a distance-based approach to estimate the probability that a sample is in-distribution. Note that the competency score estimates a probability and is thus always between zero and one. If the model is 100% competent on a given image, the probability that the model's prediction

is correct is one. If the model is entirely incompetent, the probability that the prediction is correct is zero. The competency score is thus between zero and one, with higher scores corresponding to higher levels of competency. Another important aspect of this metric, that distinguishes it from many existing uncertainty and OOD estimates, is that it is differentiable with respect to the input image.

3.2 Identifying Regions Contributing to Low Competency

We explore five approaches for estimating the dependence of the competency score on regions in the input image to understand the extent to which different regions in the image contribute to low model competency.

Approach 1: Image Cropping The first naive approach we consider is referred to as *Cropping* and is described in Figure 1. For this approach, we simply compute the competency score for each region in the image within a grid. We tune the height and width of the grid cells to achieve the best performance.




Fig. 1. Image cropping approach for identifying low competency regions. We partition the image into grid cells, crop the image around each grid cell to obtain a new image, X_j^i , and compute the competency score, C_j^i , for each cropped image. Regions with lower scores are associated with lower levels of competency and are said to contribute more to the overall low model competency, resulting in a higher dependency score, d_j^i .

Approach 2: Segment Masking We refer to the second approach considered as *Masking*. In this approach, we compute the competency score for each segment in the image, while masking out the rest of the image, as shown in Figure 2. We explore various methods for masking in an attempt to capture “feature missingness” [60]—blurring regions that are not of interest, adding noise to those regions, or replacing those pixels with zeros, ones, average values, uniform random values, or Gaussian random values—and select the highest-performing method.

Approach 3: Pixel Perturbation The third approach, which is closely related to the second, is referred to as *Perturbation* and is described in Figure 3. In this approach, we measure the increase in competency achieved by perturbing segments of the image. This approach is related to perturbation approaches in

XAI that predict the importance of pixels on a model’s prediction by perturbing those pixels and observing the resulting change in output [43]. We consider various methods for perturbing regions of interest, as discussed for *Masking*.




Fig. 2. Segment masking approach for identifying low competency regions. We begin by segmenting the image using the Felzenszwalb segmentation algorithm [61]. For each segment determined by this algorithm, we mask out the rest of the image to obtain a new image, X_i , and compute the competency score, C_i , of that masked image. Segments with lower corresponding competency scores are said to contribute more to the overall low model competency for that image, resulting in a higher dependency score, d_i .




Fig. 3. Pixel perturbation approach for identifying low competency regions. We begin by segmenting the image using the Felzenszwalb algorithm [61]. We then perturb each segment determined by this algorithm to obtain a new image, X_i , and compute the competency score, C_i , for each of these perturbed images. If the competency score increases (above C) after perturbing a given region, that region is believed to contribute more to the overall low model competency, resulting in a higher dependency score, d_i .

Approach 4: Competency Gradients The fourth approach, which we refer to as *Gradients*, relies on gradient information relating the input image to the estimated competency score and is described in Figure 4. The competency score (defined as C in Section 3.1) is a differentiable function of the input image, \mathbf{X} . Therefore, we can compute the partial derivative of the competency score with respect to each of the pixel values in the input image. If the derivative is large, then small changes in this pixel value will significantly affect the competency score, and we say that model competency is more dependent on this pixel. This approach is similar to gradient-based approaches in XAI [41] but considers

the dependence of the competency score on segmented regions, rather than the dependence of the model output on individual pixels.




Fig. 4. Competency gradients approach for identifying low competency regions. We begin by computing the partial derivative of the overall competency score with respect to each of the pixel values in the input image. We then calculate the average derivative over segmented regions in the image, using the Felzenszwalb algorithm [61].

Approach 5: Reconstruction Loss The fifth and final approach we considered is referred to as *Reconstruction* because it depends on the reconstruction loss of an autoencoder trained to fill in missing regions of an input image, as shown in Figure 5. The autoencoder is designed to reconstruct the original image from the feature vector used for prediction by the perception model. We assume that the ability of the autoencoder to accurately reconstruct a region of an image reflects the familiarity of the model with that image region and thus the dependence of low model competency on that region. Reconstruction-based approaches have been explored for anomaly detection [32] and localization [62] but have not yet been explored as an explainability tool. For determining the reconstruction loss across segmented regions in an image, we considered different methods and loss functions and selected those that achieved the best performance.




Fig. 5. Reconstruction loss approach for identifying low competency regions. We begin by segmenting the image using the Felzenszwalb algorithm [61], then generate masked images for each segment. Given an image with a segment masked out using ones, an autoencoder aims to predict the pixels of the original image. The difference between the original image and the predicted image is the reconstruction loss. If the reconstruction loss is high for a given image segment, this segment is believed to be unfamiliar to the perception model, and the overall low model competency is said to be more dependent on this region in the image, resulting in a higher dependency score, d_i .

4 Method Evaluation & Analysis

We provide a thorough analysis of the performance and usefulness of each of these five approaches for identifying regions contributing to low model competency. We conduct analysis across three unique datasets that assess the ability to identify unfamiliar objects (Section 4.2), detect regions associated with unseen classes (Section 4.3), and recognize unexplored areas in an environment (Section 4.4).

4.1 Metrics for Comparison

Following much of the work in explainable image classification (Section 2.3), we provide a *visual* comparison of the outputs of the five approaches. This visual analysis allows us to better understand the method outputs and determine the extent to which we trust each of the methods for explaining model competency. It also helps to determine whether the outputs of these methods offer explanations that are helpful to and interpretable by a human user. While we cannot display every example image in this paper, we include a few representative images.

While visual analysis is useful, it may be misleading to rely solely on a qualitative comparison [60]. In an effort to provide a more rigorous quantitative comparison, we develop a metric for *accuracy* that evaluates the ability of each of the five approaches to identify regions that are unfamiliar to the model. Each image in the test set is manually labeled to indicate which regions contain features that were not present in the training set. A pixel that corresponds to a feature that was not present in the training set is considered a positive sample, and a pixel that corresponds to a feature that was present is considered a negative sample. We then assess the ability of the five approaches to label the dataset in the same way, measuring the overall accuracy, true positive rate (TPR), true negative rate (TNR), positive predictive value (PPV), and negative predictive value (NPV) averaged across all of the images in the test set.

Finally, in order for any of these approaches to truly be useful for decision-making, we must consider the *computation time* required to identify regions contributing to low competency. While it is valuable to provide users a better understanding of how their model works, when it is incompetent, and what leads to incompetency, it is desirable for systems to recognize their level of competency and respond appropriately without the need for human intervention. In order for this autonomy to be realistic, the system needs to reason about its competency fast enough that doing so does not interrupt the decision-making process.

4.2 Dataset 1: Lunar Environment

The first dataset we consider was obtained from a lunar environment developed by the Johns Hopkins University Applied Physics Laboratory in the Gazebo simulator, in which the training data contains images from an uninhabited moon and the test data contains images from a habited moon. While the training images only contain the lunar surface, the sky, and shadows, the test images additionally contain astronauts and human-made structures. This dataset assesses

the ability of the regional competency approaches discussed in Section 3.2 to identify unfamiliar objects. A visual comparison of these approaches is shown in Figure 6 and a quantitative comparison is given in Table 1. We also consider the performance of existing class activation mapping techniques for identifying unfamiliar objects in Figure 11 and Table 4 of Appendix B.1.




Fig. 6. Visual comparison of regional competency explanation methods for a lunar environment. Rows correspond to different input images and columns to different techniques. More red regions are more associated with low levels of model competency.

Upon visual inspection, it is apparent that the *Cropping* method can roughly identify regions in the lunar images with unfamiliar objects, but this method is quite noisy, limiting its utility. This method has the worst overall accuracy, as it is quite conservative and often classifies regions that should be familiar to the model as ones leading to low model competency. In addition, this method is the slowest by far and may be too slow for many decision-making problems.

The four approaches that use segmentation appear more useful for identifying unfamiliar objects. The *Masking* method does well on some examples in the lunar dataset, but it also often inverts the expected order of dependency, giving low dependency scores to many regions that contain unfamiliar objects. The outputs of the *Perturbation* method more often align with what is expected, but this method also fails to identify several unfamiliar objects. The *Gradients* method tends to perform quite well when the unfamiliar structures look significantly different from elements of the uninhabited moon, but its performance deteriorates when the image contains dark structures that blend into the background.

These three methods perform comparably in terms of overall accuracy, but the *Gradients* method more often identifies regions as containing unfamiliar features, resulting in a higher TPR and lower TNR compared to the *Masking* and *Perturbation* methods. Visually, the *Reconstruction* method appears to be the

most useful and reliable; it consistently assigns high scores to unfamiliar structures and generally always seems to outperform the other methods. Quantitatively, it also distinguishes between familiar and unfamiliar features most accurately. This method achieves the highest overall accuracy, along with the highest TPR, TNR, PPV, and NPV. In terms of computation time, the *Gradients* method is fastest, but the *Masking*, *Perturbation*, and *Reconstruction* methods are also reasonably fast, obtaining their outputs in around 0.15 seconds or less.

Table 1. Quantitative comparison of regional competency explanation methods for a lunar environment. A true positive occurs when a method correctly identifies a pixel as one corresponding to an unfamiliar object (i.e., an astronaut or human-made structure), and a true negative occurs when a method identifies a pixel as one corresponding to a feature present in the training set (i.e., the sky, the lunar surface, or a shadow).

Method	Average Time (sec)	Average Accuracy (%)				
		Overall	TPR	TNR	PPV	NPV
Cropping	1.539	71.27	70.01	72.97	26.48	92.03
Masking	0.151	84.01	33.07	92.99	65.06	87.62
Perturbation	0.116	83.53	29.74	92.57	59.22	87.02
Gradients	0.070	82.16	51.32	87.59	64.44	91.01
Reconstruction	0.081	96.16	88.87	98.83	93.53	95.95

4.3 Dataset 2: Speed Limit Signs

The second dataset contains speed limit signs in Germany [63]. While the training dataset is composed of common speed limit signs (30 km/hr and higher), the test dataset set also contains an uncommon speed limit (20 km/hr). This dataset assesses the ability of the regional competency approaches to identify regions associated with unseen classes. A visual comparison of these approaches is shown in Figure 7 and a quantitative comparison is given in Table 2. We also consider the performance of existing class activation mapping techniques for identifying regions associated with unseen classes in Figure 12 and Table 5 of Appendix B.2.

Table 2. Quantitative comparison of regional competency explanation methods for a speed limit dataset. A true positive occurs when a method correctly identifies a pixel as one corresponding to a feature associated with an unseen class (i.e., the number 2), and a true negative occurs when a method identifies a pixel as one corresponding to a feature present in the training set (i.e., other parts of the traffic sign and background).

Method	Average Time (sec)	Average Accuracy (%)				
		Overall	TPR	TNR	PPV	NPV
Cropping	24.27	22.21	54.37	21.08	2.30	93.06
Masking	5.17	47.57	20.32	48.44	4.55	92.14
Perturbation	5.16	83.76	32.24	85.51	12.77	97.33
Gradients	0.319	83.98	60.99	84.85	12.76	98.41
Reconstruction	0.109	92.19	58.14	93.43	26.20	98.44

Inspecting outputs across examples in the speed limit signs dataset, it appears that the *Cropping* method is not useful for this dataset, offering very little information through its outputs. The *Masking* method also does not appear very useful, as the outputs generally do not align with our expectations about the dependence of low model competency on elements of the unseen class (i.e., the number 2). Both of these methods have very low overall accuracy.

Compared to the first two methods, *Perturbation* and *Gradients* appear to offer more useful information. However, while both methods will sometimes assign higher dependency to the digit associated with the unseen class, the *Perturbation* method is not the most reliable and the *Gradients* method often places a higher dependency on the zero digit than it does on the two. While these two methods appear very similar in terms of their overall accuracy, the *Gradients* method has a much higher TPR, while maintaining a similar TNR.

Both visually and quantitatively, the *Reconstruction* method appears to perform the best. This method often indicates that low model competency is highly dependent on the digit 2, which is associated with an unseen class. It also achieves the highest overall accuracy, TNR, PPV, and NPV. In addition, this method has the lowest average computation time, maintaining a time of around 0.1 seconds.

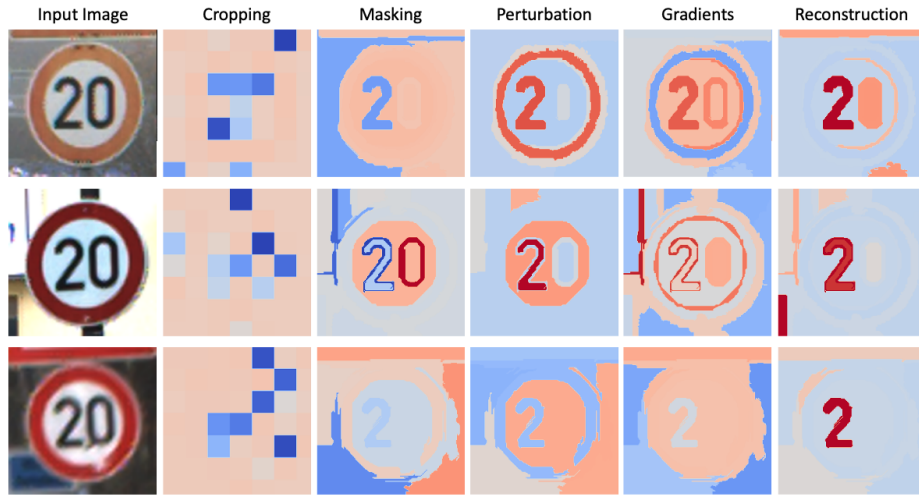


Fig. 7. Visual comparison of regional competency explanation methods for a speed limit dataset. Rows correspond to different input images and columns to different techniques. More red regions are more associated with low levels of model competency.

4.4 Dataset 3: Outdoor Park

The third and final dataset contains images from regions in a park. While the training dataset only contains images from forested and grassy regions of the park, the test dataset additionally includes images from around the park’s pavilion. This dataset assesses the ability of the regional competency approaches to

identify unexplored areas in an environment. A visual comparison of these approaches is shown in Figure 8 and a quantitative one is given in Table 3. We also consider the performance of existing class activation mapping techniques for identifying unexplored regions in Figure 13 and Table 6 of Appendix B.3.

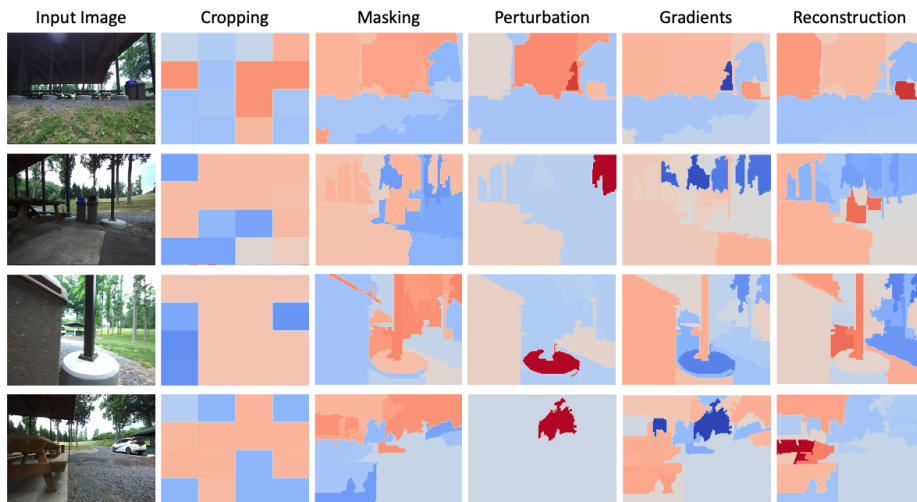


Fig. 8. Visual comparison of regional competency explanation methods for an outdoor park. Rows correspond to different input images and columns to different techniques. More red regions are more associated with low levels of model competency.

As with the speed limit signs dataset, the *Cropping* method appears to offer little explanation for low model competency with the park dataset. The *Masking* and *Perturbation* methods sometimes identify unexplored regions as expected, but they do not consistently assign these regions higher incompetency dependency scores. These three methods all have low overall accuracy scores.

The *Gradients* and *Reconstruction* methods perform comparably, assigning high dependency scores to many of the image segments associated with unexplored areas and achieving similar levels of accuracy. Comparing these two methods, *Gradients* achieves a higher overall accuracy, TPR, and NPV, but *Reconstruction* achieves a higher TNR and PPV. The *Reconstruction* method is faster than *Gradients* in terms of computation time, but both methods are reasonably fast and could be used in most decision-making pipelines.

Because neither the *Gradients* nor *Reconstruction* method clearly outperforms the other, it is interesting to consider how they might be used in combination. The average accuracy achieved using the average of the outputs of these two methods is shown in the last row of Table 3. Combining these two methods increases the overall accuracy to over 80% and the TPR to over 90%. Using this technique, the NPV is also higher, but the TNR and PPV are lower. While using

two methods would require more computational effort, the computation time is still within reason to be useful for decision-making.

Table 3. Quantitative comparison of regional competency methods for an outdoor park environment. A true positive occurs when a method correctly identifies a pixel as one corresponding to an unexplored region of the environment (i.e., areas around the pavilion), and a true negative occurs when a method identifies a pixel as one corresponding to a feature present in the training set (i.e., forested and grassy regions).

Method	Average Time (sec)	Average Accuracy (%)				
		Overall	TPR	TNR	PPV	NPV
Cropping	0.261	37.82	48.48	17.20	53.51	15.73
Masking	0.301	34.51	29.00	43.60	44.73	24.73
Perturbation	0.330	37.67	22.47	67.14	27.76	30.48
Gradients	0.130	77.13	81.13	69.52	81.99	71.86
Reconstruction	0.053	75.56	79.35	73.81	84.53	67.49
Grad. + Reco.	0.183	80.09	90.26	62.33	81.22	79.73

4.5 Analysis of Results

Comparing all of our results across the three datasets in Figure 9, we see that the *Cropping* method may provide a rough idea of image regions contributing to low model competency for simple datasets, but this method is unlikely to provide utility for more complex datasets. The *Masking* and *Perturbation* methods may be useful for some datasets, but it is unlikely that they would provide more explanatory power than the *Gradients* and *Reconstruction* methods. Among all of the regional competency methods, the *Reconstruction* method appears the most promising for identifying regions contributing to low model competency, but the *Gradients* method also has merit. While the *Reconstruction* method more reliably informs us of regions in the image that are unfamiliar to the classifier, the *Gradients* method more directly tells us which pixels the competency score actually depends on. Because these two methods inform us of slightly different aspects of model confidence, there may be value in combining their outputs.

From Figure 10, both the *Gradients* and *Reconstruction* methods can obtain their outputs in around 0.1 seconds or less on a 2.3 GHz Quad-Core Intel Core i7 processor. While these methods are fast enough to be usable in most decision-making pipelines, their computation times are dependent on image size, classifier architecture, and segmentation algorithm parameters, so computation time can vary across systems. While the increase in computation time seen for the speed limit signs dataset was not terribly significant for the *Gradients* and *Reconstruction* methods, it was quite significant for the *Cropping*, *Masking*, and *Perturbation* methods. The *Cropping* method is unlikely to run fast enough for most systems, and the *Masking* and *Perturbation* methods may be too slow for systems with larger images and a greater numbers of image segments.

From a visual comparison of the performance of the regional competency methods across the three datasets, it appears that these methods may be better




Fig. 9. Comparison of the overall accuracy of the five regional competency explanation methods in identifying regions of an image that are unfamiliar to the perception model.

at detecting unfamiliar objects and structures in an image (as with the lunar and park datasets), rather than identifying elements of unseen classes (as with the speed limit dataset). For the speed limit dataset, the ‘2’ associated with the unseen class looks very similar to the ‘3’ associated with a seen class. It is very possible that elements of the background appear more unfamiliar to the model than does the unseen digit. All five methods are much more precise when tasked with simply identifying features that are unfamiliar to the model, rather than those that the classifier should depend on but has not seen before.

Furthermore, these methods are generally more successful at detecting objects and structures that do not look very similar to elements in the training set. The *Gradients* and *Reconstruction* methods were generally effective for the park dataset but tended to assign low dependency scores to the gravel in unexplored areas of the park, which is unfamiliar to the model but looks similar to dirt and short grass that the classifier has seen. While the *Reconstruction* method was quite successful for the lunar dataset, the *Gradients* method struggled to identify dark and shadowy objects that blended into the lunar background.

In addition, these methods were generally most successful for the lunar dataset, which is relatively simple and can be more clearly segmented into a few distinct regions. The Felzenszwalb segmentation algorithm [61] is useful because it is fast and does not need to be trained on a specific dataset to work effectively. However, it often produced strange results, particularly for the speed limit and park datasets, affecting the performance of the four approaches that utilized this algorithm. These methods may be improved by using a different segmentation algorithm—especially one that is domain-specific.

Overall, the *Gradients* and *Reconstruction* methods show promise for identifying regions associated with low model competency, particularly when aspects of the image that are unfamiliar to the perception model cause this reduction in




Fig. 10. Comparison of the computation time of the five regional competency methods.

competency. These methods are generally accurate and fast enough to be used for decision-making, but additional work may be done to improve these methods.

5 Conclusions

In this paper, we sought to develop tools for understanding when DNN-based image classification models fail and why. We expand upon uncertainty quantification and OOD detection work to offer explanations for why a perception model lacks confidence in its prediction for a given input image by identifying specific regions in the image that lead to low model competency. Rather than offering explanations for classification decisions, as is typically done in explainable AI work, we aim to explain why a model is incompetent in its prediction.

With these goals in mind, we developed five distinct methods for identifying regions in an image contributing to low model competency, which we refer to as (1) *Image Cropping*, (2) *Segment Masking*, (3) *Pixel Perturbation*, (4) *Competency Gradients*, and (5) *Reconstruction Loss*. Each of these methods outputted an incompetency dependency score for each pixel in a set of input images. We evaluated each of these five approaches across three diverse datasets with the aim of determining which methods were most suitable for identifying unfamiliar objects, regions associated with unseen classes, and unexplored areas in an environment. We compared the dependency outputs of the five approaches visually, evaluated their ability to accurately identify image features that are unfamiliar to the classification model, and compared their computation times.

We found the *Reconstruction Loss* method to be the most effective and promising approach overall but also found merit in the *Competency Gradients* method. We believe both methods offer valuable information about why an image classification model is incompetent for a given input image, and their outputs can be obtained fast enough to be used in decision-making problems.

6 Limitations & Future Work

Additional work could be done to improve the performance of the regional competency estimation methods presented in this work. First, it may be worthwhile to explore different segmentation algorithms, particularly ones that are domain-specific or trained to work well with a particular dataset or environment of interest. For the *Reconstruction* method, it would be interesting to consider additional reconstruction loss functions, as well as ensembles of autoencoders. While the *Reconstruction* method is independent of the competency metric, the other four methods are directly dependent on the chosen metric. In this work, we only considered a single metric for competency, which incorporates different facets of uncertainty into a single score. It would be interesting to disentangle these sources of uncertainty or consider different metrics for competency. While the *Cropping*, *Masking*, and *Perturbation* methods do not have strict requirements for the chosen competency metric, the *Gradients* method is only useful if the competency score is differentiable with respect to the input image.

Beyond these more immediate next steps for improving the proposed methods, there are several other avenues for future work. First, this work focused on explaining low model competency for OOD images with regional features that are clearly unfamiliar to the perception model. Additional work would need to be done to understand low model competency for more nuanced cases, particularly with in-distribution images, and to understand high model competency. One could explore other methods for explaining model competency that are not focused on identifying particular regions contributing to low competency. It would be worthwhile to identify other types of image features that might contribute to low competency, such as image brightness. This work also seeks to understand perception model competency of trained models without modifying the underlying model architecture. It would be interesting to also explore ad hoc methods from the XAI literature to better understand model competency.

The approaches explored in this paper can be used to provide users a deeper understanding of how a model works and when it might fail. Additional work could be done with this goal in mind to combine visual explanations with language. Future work should also explore how the regional competency information presented in this paper should play into decision-making problems. While these methods provide perception-based systems with more information about why an image classification model is incompetent, it is still not clear how the system should respond. Additional work would need to be done to decide the appropriate response to model incompetency for different applications. Finally, the work presented in this paper focused on image classification problems. Additional work should be done to understand how these approaches could be adapted for regression, semantic segmentation, object detection, and other problems.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

A Data Sources & Algorithm Parameters

To reproduce this work, data is available on Drive, code is available on GitHub, and the chosen model architectures and training parameters are available on Drive. For more information on reproducing this work, please see this README.

B Comparison to Class Activation Maps

The goal of the work presented in this paper is to identify regions in an image contributing to low model competency and to display this incompetency dependency in a saliency map. This is closely related to work on developing class activation maps (CAMs), which are class-discriminative saliency maps that identify key regions in an image used by the model to make a particular class prediction. There are various methods for generating CAMs [64], but we choose to focus on a few representative approaches: Grad-CAM [41], Guided Grad-CAM [41], Grad-CAM++ [65], Integrated Gradients [42], SmoothGrad [66], DeepLIFT [48], Score-CAM [45], Ablation-CAM [67], Eigen-CAM [68], and LayerCAM [69]. These methods were implemented with the help of the PyTorch library for CAM methods [70] and the Captum model interpretability library for PyTorch [71].

The chosen mapping approaches are designed to explain model predictions, but we evaluate their ability to recognize unfamiliar objects (Section B.1), identify regions associated with unseen classes (Section B.2), and determine unexplored regions in an environment (Section B.3). We find that these approaches are not suited for generating competency maps, as is done in our work.

B.1 Dataset 1: Lunar Environment

Table 4. Quantitative comparison of CAM methods for a lunar environment.

Method	Average Time (sec)	Average Accuracy (%)				
		Overall	TPR	TNR	PPV	NPV
Grad-CAM	0.03	66.18	32.50	71.24	25.32	84.34
Guided Grad-CAM	0.04	76.24	30.97	83.88	35.94	88.89
Grad-CAM++	0.04	60.80	35.76	65.88	22.19	85.76
Integrated Gradients	0.46	71.88	21.26	79.34	20.47	87.58
SmoothGrad	2.68	77.51	18.19	85.88	19.32	87.92
DeepLIFT	0.05	73.36	24.01	80.58	20.67	87.96
Score-CAM	0.13	62.40	38.94	67.26	23.42	87.18
Ablation-CAM	0.23	62.39	35.44	67.20	21.52	85.76
Eigen-CAM	16.20	71.90	36.22	78.21	31.17	89.05
LayerCAM	0.11	73.81	19.87	82.31	15.68	87.25

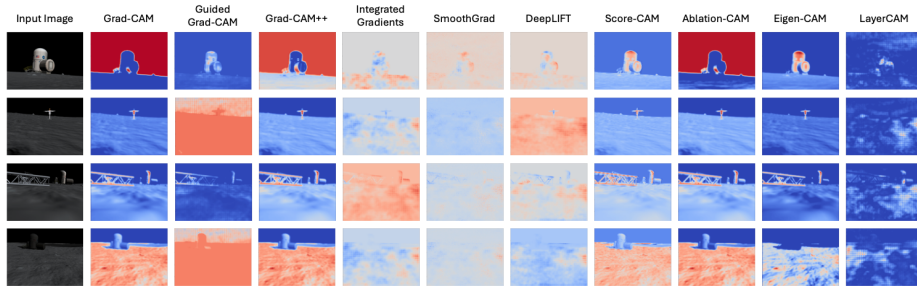


Fig. 11. Visual comparison of class activation maps (CAMs) for a lunar environment. Rows correspond to different input images and columns to different techniques for generating saliency maps. More red regions contribute more to the model’s prediction.

B.2 Dataset 2: Speed Limit Signs

Table 5. Quantitative comparison of CAM methods for a speed limit dataset.

Method	Average	Average Accuracy (%)				
	Time (sec)	Overall	TPR	TNR	PPV	NPV
Grad-CAM	0.41	69.88	36.78	71.01	4.53	97.07
Guided Grad-CAM	0.60	83.15	17.39	85.39	5.65	96.59
Grad-CAM++	0.41	51.01	18.57	52.03	1.71	94.87
Integrated Gradients	18.06	51.43	49.46	51.50	3.33	96.78
SmoothGrad	80.19	50.30	49.89	50.31	3.28	96.78
DeepLIFT	0.87	52.90	48.20	53.06	3.35	96.80
Score-CAM	4.18	–	–	–	–	–
Ablation-CAM	4.11	51.72	25.95	52.54	2.17	95.30
Eigen-CAM	88.97	71.66	62.17	71.99	6.78	98.27
LayerCAM	0.62	87.59	54.00	88.78	14.09	98.21

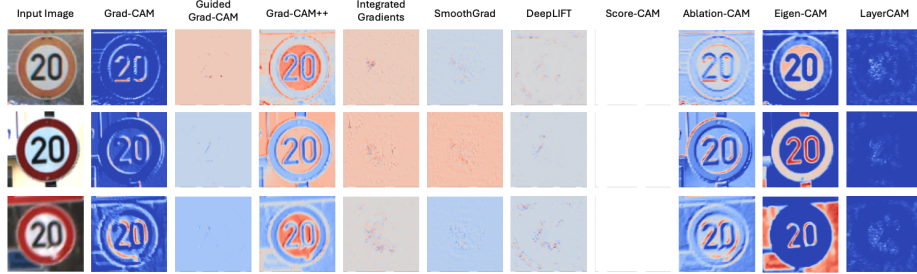


Fig. 12. Visual comparison of class activation maps (CAMs) for a speed limit dataset. Rows correspond to different input images and columns to different techniques for generating saliency maps. More red regions contribute more to the model’s prediction.

B.3 Dataset 3: Outdoor Park

Table 6. Quantitative comparison of CAM methods for an outdoor park environment.

Method	Average Time (sec)	Average Accuracy (%)				
		Overall	TPR	TNR	PPV	NPV
Grad-CAM	0.03	46.62	33.22	74.19	72.27	35.69
Guided Grad-CAM	0.04	34.43	13.30	78.22	60.28	30.69
Grad-CAM++	0.04	30.60	26.10	49.83	45.23	22.01
Integrated Gradients	0.46	30.69	13.25	67.67	34.40	28.66
SmoothGrad	2.65	33.98	22.23	59.38	43.78	28.68
DeepLIFT	0.05	31.57	18.97	57.25	45.10	26.82
Score-CAM	0.13	29.05	21.72	42.66	40.99	22.07
Ablation-CAM	0.22	32.16	20.99	55.19	46.16	25.71
Eigen-CAM	15.79	27.22	16.82	47.69	36.88	23.51
LayerCAM	0.10	31.27	14.96	63.28	44.79	27.66




Fig. 13. Visual comparison of class activation maps (CAMs) for an outdoor park. Rows correspond to different input images and columns to different techniques for generating saliency maps. More red regions contribute more to the model's prediction.

References

1. T. R auker, A. Ho, S. Casper, and D. Hadfield-Menell, “Toward Transparent AI: A Survey on Interpreting the Inner Structures of Deep Neural Networks,” 2023.
2. J. Gawlikowski, C. R. N. Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruspe, R. Triebel, P. Jung, R. Roscher, M. Shahzad, W. Yang, R. Bamler, and X. X. Zhu, “A survey of uncertainty in deep neural networks,” *Artificial Intelligence Review*, vol. 56, no. S1, October 2023.
3. J. Yang, K. Zhou, Y. Li, and Z. Liu, “Generalized Out-of-Distribution Detection: A Survey,” 2022.
4. V. Rajendran and W. LeVine, “Accurate layerwise interpretable competence estimation,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
5. Yarin Gal, “Uncertainty in Deep Learning,” Doctor of Philosophy, University of Cambridge, Cambridge, England, September 2016.
6. M. Sugiyama, J. Qui onero-Candela, A. Schwaighofer, and N. Lawrence, Eds., *Dataset shift in machine learning*. Cambridge, Mass: MIT Press, December 2008.
7. T. Ramalho and M. Miranda, “Density estimation in representation space to predict model uncertainty,” in *Engineering Dependable and Secure Machine Learning Systems (EDSMLS)*, 2020.
8. J. Z. Liu, Z. Lin, S. Padhy, D. Tran, T. Bedrax-Weiss, and B. Lakshminarayanan, “Simple and principled uncertainty estimation with deterministic deep learning via distance awareness,” in *34th International Conference on Neural Information Processing Systems (NIPS)*, 2020.
9. P. Oberdiek, M. Rottmann, and H. Gottschalk, “Classification uncertainty of deep neural networks based on gradient information,” in *Artificial Neural Networks in Pattern Recognition (ANNPR)*, 2018.
10. J. Lee and G. AlRegib, “Gradients as a measure of uncertainty in neural networks,” in *IEEE International Conference on Image Processing (ICIP)*, 2020.
11. R. M. Neal, “Bayesian learning via stochastic dynamics,” in *5th International Conference on Neural Information Processing Systems (NIPS)*, 1992.
12. R. M. Neal, *Bayesian Learning for Neural Networks*, P. Bickel, P. Diggle, S. Fienberg, K. Krickeberg, I. Olkin, N. Wermuth, and S. Zeger, Eds. Springer New York, 1996, vol. 118.
13. M. Welling and Y. W. Teh, “Bayesian learning via stochastic gradient langevin dynamics,” in *28th International Conference on International Conference on Machine Learning (ICML)*, 2011.
14. A. Graves, “Practical variational inference for neural networks,” in *24th International Conference on Neural Information Processing Systems (NIPS)*, 2011.
15. D. J. Rezende and S. Mohamed, “Variational inference with normalizing flows,” in *32nd International Conference on International Conference on Machine Learning (ICML)*, 2015.
16. Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *33rd International Conference on Machine Learning (ICML)*, 2016.
17. H. Ritter, A. Botev, and D. Barber, “A scalable laplace approximation for neural networks,” in *6th International Conference on Learning Representations (ICLR)*, 2018.
18. B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” in *31st International Conference on Neural Information Processing Systems (NIPS)*, 2017.

19. S. Liang, Y. Li, and R. Srikant, "Enhancing the reliability of out-of-distribution image detection in neural networks," in *6th International Conference on Learning Representations (ICLR)*, 2018.
20. Y.-C. Hsu, Y. Shen, H. Jin, and Z. Kira, "Generalized ODIN: Detecting Out-of-Distribution Image Without Learning From Out-of-Distribution Data," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
21. W. Liu, X. Wang, J. D. Owens, and Y. Li, "Energy-based out-of-distribution detection," in *34th International Conference on Neural Information Processing Systems (NIPS)*, 2020.
22. B. Zong, Q. Song, M. R. Min, W. Cheng, C. Lumezanu, D. Cho, and H. Chen, "Deep autoencoding gaussian mixture model for unsupervised anomaly detection," in *6th International Conference on Learning Representations (ICLR)*, 2018.
23. W. Hu, J. Gao, B. Li, O. Wu, J. Du, and S. Maybank, "Anomaly detection using local kernel density estimation and context-based regression," *IEEE Transactions on Knowledge and Data Engineering*, vol. 32, no. 2, November 2020.
24. D. J. Rezende and S. Mohamed, "Variational Inference with Normalizing Flows," in *32nd International Conference on Machine Learning (ICML)*, 2015.
25. D. P. Kingma and P. Dhariwal, "Glow: Generative Flow with Invertible 1x1 Convolutions," in *32nd Conference on Neural Information Processing Systems (NeurIPS)*, 2018.
26. J. Ren, P. J. Liu, E. Fertig, J. Snoek, R. Poplin, M. A. DePristo, J. V. Dillon, and B. Lakshminarayanan, "Likelihood ratios for out-of-distribution detection," in *33rd International Conference on Neural Information Processing Systems (NIPS)*, 2019.
27. K. Lee, K. Lee, H. Lee, and J. Shin, "A simple unified framework for detecting out-of-distribution samples and adversarial attacks," in *32nd International Conference on Neural Information Processing Systems (NIPS)*, 2018.
28. A. Zaeemzadeh, N. Bisagno, Z. Sambugaro, N. Conci, N. Rahnavard, and M. Shah, "Out-of-distribution detection using union of 1-dimensional subspaces," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
29. E. Techapanurak, M. Suganuma, and T. Okatani, "Hyperparameter-free out-of-distribution detection using cosine similarity," in *Asian Conference on Computer Vision (ACCV)*, 2020.
30. Y. Sun, Y. Ming, X. Zhu, and Y. Li, "Out-of-Distribution Detection with Deep Nearest Neighbors," in *39th International Conference on Machine Learning (ICML)*, 2022.
31. Y. Xia, X. Cao, F. Wen, G. Hua, and J. Sun, "Learning discriminative reconstructions for unsupervised outlier removal," in *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015.
32. D. Gong, L. Liu, V. Le, B. Saha, M. R. Mansour, S. Venkatesh, and A. v. d. Hengel, "Memorizing Normality to Detect Anomaly: Memory-augmented Deep Autoencoder for Unsupervised Anomaly Detection," 2019.
33. J. An and S. Cho, "Variational autoencoder based anomaly detection using reconstruction probability," in *Special Lecture on IE, SNU Data Mining Center*, 2015.
34. Z. Chen, C. K. Yeo, B. S. Lee, and C. T. Lau, "Autoencoder-based network anomaly detection," in *2018 Wireless Telecommunications Symposium (WTS)*, 2018.
35. H. Zenati, C. S. Foo, B. Lecouat, G. Manek, and V. R. Chandrasekhar, "Efficient GAN-Based Anomaly Detection," 2019.
36. M. Sabokrou, M. Khalooei, M. Fathy, and E. Adeli, "Adversarially learned one-class classifier for novelty detection," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

37. P. Perera, R. Nallapati, and B. Xiang, "Ocgan: One-class novelty detection using gans with constrained latent representations," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
38. S. Ali, T. Abuhmed, S. El-Sappagh, K. Muhammad, J. M. Alonso-Moral, R. Con-falonieri, R. Guidotti, J. Del Ser, N. Díaz-Rodríguez, and F. Herrera, "Explainable Artificial Intelligence (XAI): What we know and what is left to attain Trustworthy Artificial Intelligence," *Information Fusion*, vol. 99, November 2023.
39. V. Kamakshi and N. C. Krishnan, "Explainable Image Classification: The Journey So Far and the Road Ahead," *AI*, vol. 4, no. 3, August 2023.
40. K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," in *Workshop at International Conference on Learning Representations*, 2014.
41. R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *IEEE International Conference on Computer Vision (ICCV)*, 2017.
42. M. Sundararajan, A. Taly, and Q. Yan, "Axiomatic attribution for deep networks," in *34th International Conference on Machine Learning (ICML)*, 2017.
43. R. C. Fong and A. Vedaldi, "Interpretable explanations of black boxes by meaningful perturbation," in *IEEE International Conference on Computer Vision (ICCV)*, 2017.
44. L. M. Zintgraf, T. S. Cohen, T. Adel, and M. Welling, "Visualizing deep neural network decisions: Prediction difference analysis," in *International Conference on Learning Representations (ICLR)*, 2017.
45. H. Wang, Z. Wang, M. Du, F. Yang, Z. Zhang, S. Ding, P. Mardziel, and X. Hu, "Score-cam: Score-weighted visual explanations for convolutional neural networks," in *Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2020.
46. S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek, "On Pixel-Wise Explanations for Non-Linear Classifier Decisions by Layer-Wise Relevance Propagation," *PLOS ONE*, vol. 10, no. 7, July 2015.
47. G. Montavon, S. Lapuschkin, A. Binder, W. Samek, and K.-R. Müller, "Explaining nonlinear classification decisions with deep Taylor decomposition," *Pattern Recognition*, vol. 65, May 2017.
48. A. Shrikumar, P. Greenside, and A. Kundaje, "Learning important features through propagating activation differences," in *34th International Conference on Machine Learning*, 2017.
49. P. Dabkowski and Y. Gal, "Real time image saliency for black box classifiers," in *31st International Conference on Neural Information Processing Systems (NIPS)*, 2017.
50. M. D. Zeiler and R. Fergus, "Visualizing and Understanding Convolutional Networks," in *Computer Vision – ECCV 2014*, 2014.
51. M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should i trust you?: Explaining the predictions of any classifier," in *22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016.
52. M. T. Ribeiro, S. Singh, and C. Guestrin, "Anchors: High-Precision Model-Agnostic Explanations," in *AAAI Conference on Artificial Intelligence*, 2018.
53. S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *31st International Conference on Neural Information Processing Systems (NIPS)*, 2017.
54. S. M. Lundberg, G. G. Erion, and S.-I. Lee, "Consistent feature attribution for tree ensembles," in *34th International Conference on Machine Learning (ICML)*, 2017.

55. P. W. Koh, T. Nguyen, Y. S. Tang, S. Mussmann, E. Pierson, B. Kim, and P. Liang, "Concept bottleneck models," in *37th International Conference on Machine Learning (ICML)*, 2020.
56. B. Kim, M. Wattenberg, J. Gilmer, C. Cai, J. Wexler, F. Viegas, and R. Sayres, "Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV)," in *35th International Conference on Machine Learning (ICML)*, 2018.
57. A. Ghorbani, J. Wexler, J. Zou, and B. Kim, "Towards automatic concept-based explanations," in *33rd International Conference on Neural Information Processing Systems (NIPS)*, 2019.
58. C.-K. Yeh, B. Kim, S. O. Arik, C.-L. Li, T. Pfister, and P. Ravikumar, "On completeness-aware concept-based explanations in deep neural networks," in *34th International Conference on Neural Information Processing Systems (NIPS)*, 2020.
59. Y. Goyal, Z. Wu, J. Ernst, D. Batra, D. Parikh, and S. Lee, "Counterfactual visual explanations," in *36th International Conference on Machine Learning*, 2019.
60. P. Sturmfels, S. Lundberg, and S.-I. Lee, "Visualizing the Impact of Feature Attribution Baselines," *Distill*, vol. 5, no. 1, January 2020.
61. P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient Graph-Based Image Segmentation," *International Journal of Computer Vision*, vol. 59, no. 2, September 2004.
62. V. Zavrtanik, M. Kristan, and D. Skočaj, "Reconstruction by inpainting for visual anomaly detection," *Pattern Recognition*, vol. 112, April 2021.
63. S. Houben, J. Stallkamp, J. Salmen, M. Schlipsing, and C. Igel, "Detection of traffic signs in real-world images: The German Traffic Sign Detection Benchmark," in *International Joint Conference on Neural Networks*, 2013.
64. M. He, B. Li, and S. Sun, "A survey of class activation mapping for the interpretability of convolution neural networks," in *Signal and Information Processing, Networking and Computers*, 2023.
65. A. Chattopadhyay, A. Sarkar, P. Howlader, and V. N. Balasubramanian, "Gradcam++: Generalized gradient-based visual explanations for deep convolutional networks," in *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2018.
66. D. Smilkov, N. Thorat, B. Kim, F. Viégas, and M. Wattenberg, "Smoothgrad: removing noise by adding noise," 2017.
67. S. Desai and H. G. Ramaswamy, "Ablation-cam: Visual explanations for deep convolutional network via gradient-free localization," in *Winter Conference on Applications of Computer Vision (WACV)*, 2020.
68. M. Bany Muhammad and M. Yeasin, "Eigen-cam: Visual explanations for deep convolutional neural networks," *SN Computer Science*, vol. 2, no. 1, January 2021.
69. P.-T. Jiang, C.-B. Zhang, Q. Hou, M.-M. Cheng, and Y. Wei, "Layercam: Exploring hierarchical class activation maps for localization," *IEEE Transactions on Image Processing*, vol. 30, 2021.
70. J. Gildenblat, "Pytorch library for cam methods," <https://github.com/jacobgil/pytorch-grad-cam>, 2021.
71. N. Kokhlikyan, V. Miglani, E. W. Miguel Martin, B. Alsallakh, J. Reynolds, A. Melnikov, N. Kliushkina, C. Araya, S. Yan, and O. Reblitz-Richardson, "Captum: A unified and generic model interpretability library for pytorch," <https://github.com/pytorch/captum>, 2020.