

Supporting Safety Analysis of Image-processing DNNs through Clustering-based Approaches

MOHAMMED OUALID ATTAOUI, SnT Centre, University of Luxembourg, Luxembourg

HAZEM FAHMY, SnT Centre, University of Luxembourg, Luxembourg

FABRIZIO PASTORE, SnT Centre, University of Luxembourg, Luxembourg

LIONEL BRIAND*, Lero Centre, University of Limerick, Ireland and School of EECS, University of Ottawa, Canada

The adoption of deep neural networks (DNNs) in safety-critical contexts is often prevented by the lack of effective means to explain their results, especially when they are erroneous. In our previous work, we proposed a white-box approach (HUDD) and a black-box approach (SAFE) to automatically characterize DNN failures. They both identify clusters of similar images from a potentially large set of images leading to DNN failures. However, the analysis pipelines for HUDD and SAFE were instantiated in specific ways according to common practices, deferring the analysis of other pipelines to future work.

In this paper, we report on an empirical evaluation of 99 different pipelines for root cause analysis of DNN failures. They combine transfer learning, autoencoders, heatmaps of neuron relevance, dimensionality reduction techniques, and different clustering algorithms. Our results show that the best pipeline combines transfer learning, DBSCAN, and UMAP. It leads to clusters almost exclusively capturing images of the same failure scenario, thus facilitating root cause analysis. Further, it generates distinct clusters for each root cause of failure, thus enabling engineers to detect all the unsafe scenarios. Interestingly, these results hold even for failure scenarios that are only observed in a small percentage of the failing images.

CCS Concepts: • **Software and its engineering** → **Software defect analysis**; • **Computing methodologies** → **Machine learning**.

Additional Key Words and Phrases: DNN Explanation, DNN Functional Safety Analysis, DNN Debugging, Clustering, Transfer Learning

ACM Reference Format:

Mohammed Oualid Attaoui, Hazem Fahmy, Fabrizio Pastore, and Lionel Briand. 2024. Supporting Safety Analysis of Image-processing DNNs through Clustering-based Approaches. 1, 1 (March 2024), 51 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Deep neural networks (DNNs) have achieved extremely high predictive accuracy in various domains, such as computer vision [3, 73], autonomous driving [51, 89], and natural language processing [22,

*This work was done while affiliated with University of Luxembourg and University of Ottawa

Authors' addresses: Mohammed Oualid Attaoui, SnT Centre, University of Luxembourg, JFK 29, Luxembourg, Luxembourg, mohammed.attaoui@uni.lu; Hazem Fahmy, SnT Centre, University of Luxembourg, JFK 29, Luxembourg, Luxembourg, hazem.fahmy@uni.lu; Fabrizio Pastore, SnT Centre, University of Luxembourg, JFK 29, Luxembourg, Luxembourg, fabrizio.pastore@uni.lu; Lionel Briand, Lero Centre, University of Limerick, Tierney building, Limerick, Ireland, School of EECS, University of Ottawa, Ottawa, Canada, lionel.briand@uni.lu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Association for Computing Machinery.

XXXX-XXXX/2024/3-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

63]. Despite their superior performance, the lack of explainability of DNN models remains an issue in many contexts. While they can approximate complex and arbitrary functions, studying their structure often provides little or no insight into the underlying prediction mechanisms. There seems to be an intrinsic tension between Machine Learning (ML) performance and explainability. Often the highest-performing methods (for example, Deep Learning) are the least explainable, and the most explainable (for example, decision trees) are the least accurate [35].

For DNNs to be trustworthy, in many critical contexts where they are used, we must understand why they behave the way they do [7]. Explanation methods aim at making neural network decisions trustworthy [32]. Several explanation methods are proposed in the literature (see Section 5). In our work, because of our focus on safety analysis, we focus on explanation methods for root cause analysis, that is identifying the underlying reason of a DNN failure (root cause) which is, in our context, an incorrect DNN prediction. More precisely, we aim to identify root causes in terms of characteristics of the input images leading to failures; in other words, we are interested in identifying the different scenarios in which the DNN may fail. Such characterization is the first step towards retraining the DNN.

Root cause analysis techniques based on unsupervised learning have proven their effectiveness [87, 97]. These methods group failure samples (e.g., data collected during hardware testing) without requiring diagnostic labels, such that the samples in each cluster share similar root causes.

Our previous work is the first application of unsupervised learning to perform root cause analysis targeting DNN failures. Precisely, we proposed two DNN explanation methods:

SAFE (Safety Analysis based on Feature Extraction) [4] and *HUDD* (Heatmap-based Unsupervised Debugging of DNNs) [24]. They both process a set of failure-inducing images and generate clusters of similar images. Commonalities across images in each cluster provide information about the root cause of the failure. Further, the identified root causes support safety analysis because they help identify possible countermeasures to address the problem. For example, applying our approaches to failure-inducing images for a DNN that classifies car seat occupancy may include a cluster of images with child seats containing a bag; such cluster may help engineers determine that bags inside child seats are likely to be misclassified. Possible countermeasures could be to retrain the DNN using more child seats with objects or, if it does not work, integrating additional components that make the approach safer (e.g., radar technology [44]). Both *SAFE* and *HUDD* also support the identification of additional images to be used to retrain the DNN.

HUDD and *SAFE* differ with respect to the kind of data used to perform clustering and the pipeline of steps they rely on. *HUDD* applies clustering based on internal DNN information; precisely, for all failure-inducing images, it generates heatmaps capturing the relevance of DNN neurons on the DNN output. Finally, it applies a hierarchical clustering algorithm relying on a distance metric based on the generated heatmaps. *SAFE* is black-box as it does not rely on internal DNN information. It generates clusters based on the visual similarity across failure inducing images. To this end it relies on feature extraction based on transfer learning, dimensionality reduction, and the DBSCAN clustering algorithm.

SAFE and *HUDD* rely on a pipeline that has been configured in specific ways according to best practices. However, several variants exist for each component of both approaches (e.g., different transfer learning models, different clustering algorithms).

In this paper, we aim to evaluate these pipeline variants for both *SAFE* and *HUDD*. Therefore, we propose an empirical evaluation of 99 alternative configurations for *SAFE* and *HUDD* (pipelines). These pipelines were obtained using different combinations of feature extraction methods, clustering algorithms, and dimensionality reduction techniques; in addition, we assessed the effect of fine tuning the transfer learning models used by feature extraction methods. Consistent with *HUDD* and *SAFE*, our pipelines support the characterization of DNNs tested at the level of models, not systems.

Model-level testing, also called offline testing [38], concerns testing DNN models in isolation, whereas system-level testing, also called online testing [38], targets the system integrating the DNN (e.g., an autonomous driving system tested within a simulator [38]). Supporting system-level testing is part of future work.

For our empirical evaluation we considered six case study subjects, two of which were provided by our industry partner in the automotive domain, IEE Sensing [45]. Our subjects' applications include head pose classification, eye gaze detection, drowsiness detection, steering angle prediction, unattended child detection, and car position detection.

We present a systematic and extensive evaluation scheme for these pipelines, which entails generating failure causes that resemble realistic scenarios (e.g., poor lighting conditions or camera misconfiguration). Since in these scenarios the causes of failures are known a priori, such an evaluation scheme enables us to objectively analyze and evaluate the performance of pipelines while controlling the frequency of such failure scenarios.

Our empirical results suggest that the best pipelines support and facilitate the process of functional safety analysis such that they 1) can generate root-cause clusters (RCCs) that group together a very high proportion of images capturing a same root cause (94.3%, on average), 2) can capture most of the root causes of failures for all case study subjects (96.7%, on average), and 3) can perform well (i.e., are reliable) in the presence of rare failure instances in a dataset (i.e., when some causes of failures affect less than 10% of the failure-inducing images). In our approach, the root causes of failures are determined by engineers after inspecting the identified clusters. Although such a solution still requires human involvement, it simplifies an engineer's task¹; indeed, it is unlikely that a human can manually identify similarities across a large set of images leading to DNN failures. Further, though our previous work (i.e., SEDE [27]) aims to improve the degree of automation by automatically deriving expressions capturing commonalities in failure-inducing images, in this paper, we tackle an orthogonal problem: assessing which pipelines lead to clusters with better purity and coverage. One possible future work is the integration of the best analysis pipeline with SEDE.

The remainder of this paper is organized as follows. In Section 2, we briefly present the main features and limitations of SAFE and HUDD, along with other feature extraction models (Autoencoders and Backpropagation-based Heatmaps). In Section 3, we describe the different models and algorithms we use in our evaluated pipelines. In Section 4, we present the research questions, the experiment design and results, including a comparison between 99 pipelines. In Section 5, we discuss and compare related work. Finally, we conclude this paper in Section 6.

2 BACKGROUND

This section provides an overview of our previous work that inspired this research. We focus on clustering methods, heatmap-based DNN Explanations, the HUDD and SAFE DNN explanation methods, and Autoencoders.

2.1 Clustering

Clustering is a data analysis method that mines essential information from a dataset by grouping data into several groups called *clusters*. In clustering, similar data points are grouped into the same cluster, while non-similar data points are put into different clusters. There are two main objectives in data clustering; the first objective is to minimize the dissimilarity within the cluster, and the second objective is to maximize the inter-cluster dissimilarity. HUDD and SAFE rely on hierarchical

¹In our previous work [4], we conducted a user study demonstrating that the inspection of five random images per cluster is sufficient for an analyst to correctly identify the root causes of failures.

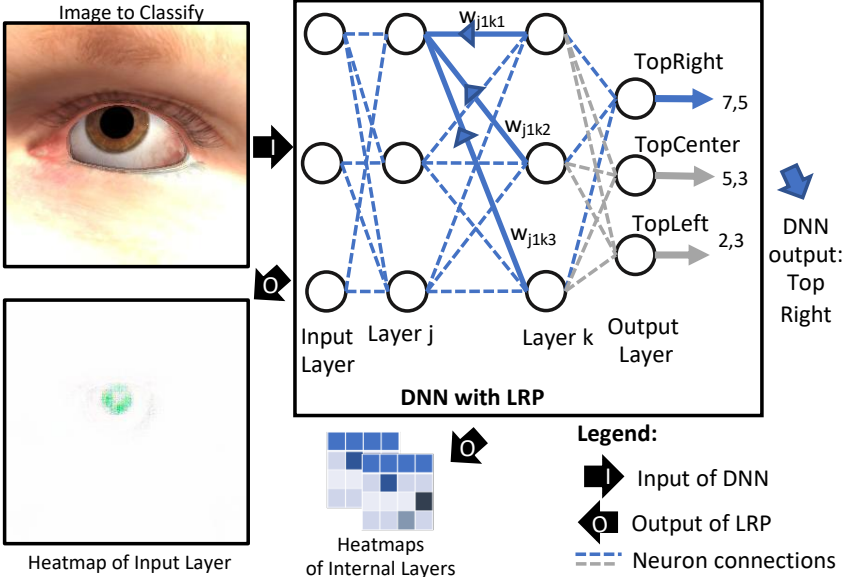


Fig. 1. Layer-Wise Relevance Propagation.

agglomerative clustering (HAC [72]) and density-based clustering (DBSCAN [23]), respectively. In HAC, each observation starts in its own cluster and pairs of clusters are iteratively merged to minimize an objective function (e.g., error sum of squares [95]). DBSCAN works by considering dense regions as clusters; it is detailed in Section 3.

2.2 Heatmap-based DNN Explanations

Approaches that aim to explain DNN results have been developed in recent years [31]. Most of these concern the generation of heatmaps that capture the importance of pixels in image predictions. They include black-box [15, 69] and white-box approaches [60, 77, 81, 100, 102]. Black-box approaches generate heatmaps for the input layer and do not provide insights regarding internal DNN layers. White-box approaches rely on the backpropagation of the relevance score computed by the DNN [60, 77, 81, 100, 102].

In this Section, we focus on a white-box technique called Layer-Wise Relevance Propagation (LRP) [60] because it has been integrated into HUDD. LRP was selected because it does not present the shortcomings of other heatmap generation approaches [24].

LRP redistributes the relevance scores of neurons in a higher layer to those of the lower layer. Figure 1 illustrates how LRP operates on a fully connected network used to classify inputs. In the forward pass, the DNN receives an input and generates an output (e.g., classifies the gaze direction as TopLeft) while recording the activations of each neuron. In the backward pass, LRP generates *internal heatmaps* for a DNN layer k , which consists of a matrix with the relevance scores computed for all the neurons of layer k .

The heatmap in Figure 1 shows that the pupil and part of the eyelid, which are the non-white parts in the heatmap, had a significant effect on the DNN output. Furthermore, the heatmap in Figure 2 shows that the mouth and part of the nose are the input pixels that mostly impacted on the DNN output.




Fig. 2. An example image of HPD subject (on the left) and applied LRP (on the right) showing that the mouth had a large influence on the DNN behavior.

A heatmap is a matrix with entries in \mathbb{R} , i.e., it is a triple (N, M, f) where $N, M \in \mathbb{N}$ and f is a map $[N] \times [M] \rightarrow \mathbb{R}$. We use the syntax $H[i, j]_x^L$ to refer to an entry in row i (i.e., $i < N$) and column j (i.e., $j < M$) of a heatmap H computed on layer L from an image x . The size of the heatmap matrix (i.e., the number of entries) is $N \cdot M$, with N and M are determined by the dimensions of the DNN layer L . For convolution layers, N represents the number of neurons in the feature map, whereas M represents the number of feature maps. For example, the heatmap for the eighth layer of AlexNet has size 169×256 (convolution layer), while the the heatmap for the tenth layer has size 4096×1 (linear layer).

2.3 Heatmap-based Unsupervised Debugging of DNNs (HUDD)

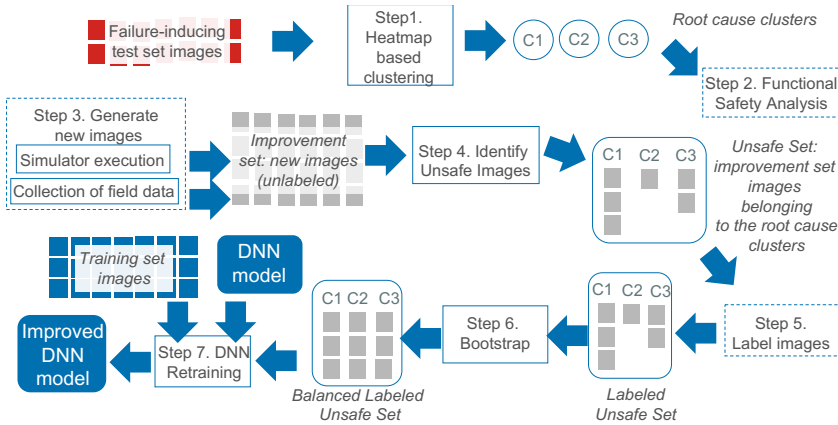


Fig. 3. Overview of HUDD.

Although heatmaps may provide useful information to determine the characteristics of an image that led to an erroneous result from the DNN, they are of limited applicability because, to determine the cause of all DNN errors observed in the test set, engineers may need to visually inspect all the error-inducing images, which is practically infeasible. To overcome such limitations, we recently developed HUDD [24], a technique that facilitates the explanation and removal of the DNN errors

observed in a test set. HUDD generates clusters of images that lead to a DNN error because of the same root cause. The root cause is determined by the engineer who visualizes a subset of the images belonging to each cluster and identifies the commonality across each image (e.g., for a Gaze detection DNN, all the images present a closed eye). To further support DNN debugging, HUDD automatically retrains the DNN by selecting a subset from a pool of unlabeled images that will likely lead to DNN errors because of the same root causes observed in the test set.

Figure 3 provides an overview of HUDD, which consists of six steps. In Step 1, root cause clusters are identified by relying on a hierarchical clustering algorithm applied to heatmaps generated for each failure inducing image. Step 2 involves a visual inspection of clustered images. In this step, engineers visualize a few representative images for each RCC; the inspection enables the engineers to determine which are the commonalities across the images in each cluster and, therefore, determine the failure root cause. Example root causes include the presence of an object inside a child seat (as reported in the Introduction) or a face turned left thus making an eye not visible and causing misclassification in a gaze detection system. HUDD's Step 2 supports functional safety analysis because each failure root cause represents a usage scenario in which the DNN is likely to fail, and, based on domain knowledge, engineers can determine the likelihood of each failure scenario, its safety impact, and possible countermeasures, as required by functional safety analysis standards [46, 47]. For example, objects inside child seats might be very common but they lead to false alarms not hazards; misclassified gaze may instead prevent the system from determining that the driver is not pay attention to the road. Countermeasures include the retraining of the DNN, which is supported by HUDD's Step 3. In Step 3, a new set of images, referred to as the *improvement set*, is provided by the engineers to retrain the model. In Step 4, HUDD automatically selects a subset of images from the improvement set called the *unsafe set*. The engineers label the images in the unsafe set in Step 5. Finally, in Step 6, HUDD automatically retrains the model to enhance its prediction accuracy.

Heatmap-based Clustering in HUDD. Clustering based on heatmaps is a key component of HUDD, an its functioning is useful to understand some of the pipelines considered in this paper. HUDD relies on LRP to generate an heatmap for every internal layer of the DNN, for each failure-inducing image. However, since distinct DNN layers lead to entries defined on different value ranges [61], to enable the comparison of clustering results across different layers, we generate normalized heatmaps by relying on min-max normalization [36].

For each DNN layer L , a distance matrix is constructed using the generated heatmaps; it captures the distance between every pair of failure-inducing image in the test set. The distance between a pair of images $\langle a, b \rangle$, at layer L , is computed as follows:

$$\text{heatmapDistance}_L(a, b) = \text{EuclideanDistance}(\tilde{H}_a^L, \tilde{H}_b^L) \quad (1)$$

where \tilde{H}_x^L is the heatmap computed for image x at layer L . *EuclideanDistance* is a function that computes the euclidean distance between two $N \times M$ matrices according to the formula

$$\text{EuclideanDistance}(A, B) = \sqrt{\sum_{i=1}^N \sum_{j=1}^M (A_{i,j} - B_{i,j})^2} \quad (2)$$

where $A_{i,j}$ and $B_{i,j}$ are the values in the cell at row i and column j of the matrix.

HUDD applies the HAC clustering algorithm multiple times, once for every DNN layer. For each DNN layer, HUDD selects the optimal number of clusters using the knee-point method applied to the weighted average intra-cluster distance (*WICD*). *WICD* is defined according to the following

formula:

$$WICD(L_l) = \frac{\sum_{j=1}^{|L_l|} \left(ICD(L_l, C_j) * \frac{|C_j|}{|C|} \right)}{|L_l|} \quad (3)$$

where L_l is a specific layer of the DNN, $|L_l|$ is the number of clusters in the layer L_l , ICD is the intra-cluster distance for cluster C_i belonging to layer L_l , $|C_j|$ represents the number of elements in cluster C_j , whereas $|C|$ represents the number of images in all the clusters.

In Formula 3, $ICD(L_l, C_j)$ is computed as follows:

$$ICD(L_l, C_j) = \frac{\sum_{i=0}^{N_j} heatmapDistance_{L_l}(p_i^a, p_i^b)}{N_j} \quad (4)$$

where p_i is a unique pair of images in cluster C_j , and N_j is the total number of pairs it contains. The superscripts a and b refer to the two images of the pair to which the distance formula is applied.

HUDD then select the layer L_m with the minimal $WICD$. By definition, the clusters generated for layer L_m are the ones that maximize cohesion and we therefore anticipate that they will group together images that exhibit similar characteristics.

In our study, we rely on HUDD as a feature extraction method; precisely, we use the heatmaps generated by the layer selected by HUDD as features.

2.4 Safety Analysis based on Feature Extraction (SAFE)

SAFE is based on a combination of a transfer learning-based feature extraction method, a clustering algorithm, and a dimensionality reduction technique. The workflow of SAFE matches HUDD's, except for Step 1 and Step 4. In SAFE's Step 1 RCCs are identified by relying on non-convex clustering (DBSCAN) applied to features extracted from failure-inducing images; HUDD, instead, applies hierarchical clustering to heatmaps. In Step 4, SAFE selects the improvement step using a procedure that relies on DBSCAN's outputs.

The pipelines evaluated in this paper had been inspired by the pipeline implemented in SAFE's Step 1, which consists of three stages (see Figure 4): *Feature Extraction*, *Dimensionality Reduction*, and *Clustering*. In this paper we investigate variants of the SAFE pipeline using different combinations of these components. Additionally, we introduce a fine-tuning stage where we fine-tune the pre-trained transfer learning models to generate more domain-specific models. Excluding clustering, which was introduced in Section 2.1, the components of SAFE's pipeline are briefly described below.

2.4.1 Transfer Learning and Feature Extraction. To maximize the accuracy of image-processing DNNs in a cost-effective way, engineers often rely on the transfer learning approach, which consists of transferring knowledge from a generic domain, usually ImageNet [82], to another specific domain,




Fig. 4. Generation of root cause clusters with SAFE

(e.g., safety analysis, in our case). In other terms, we try to exploit what has been learned in one task and generalize it to another task. Researchers have demonstrated the efficiency of transfer learning from ImageNet to other domains [86].

Transfer learning-based *Feature Extraction* is an efficient method to transform unstructured data into structured raw data to be exploited by any machine learning algorithm. In this method, the features are extracted from images using a pre-trained CNN model [18].

The standard CNN architecture comprises three types of layers: convolutional layers, pooling layers, and fully connected layers. The convolutional layer is considered the primary building block of a CNN. This layer extracts relevant features from input images during training. Convolutional and pooling layers are stacked to form a hierarchical feature extraction module. The CNN model receives an input image of size $(224, 224, 3)$. This image is then passed through the network's layers to generate a vector of features. The feature extraction process, for each image, generates raw data represented by a 2D matrix (denoted as X) formalized below:

$$X = \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1m} & l_1 \\ x_{21} & x_{22} & \dots & x_{2m} & l_2 \\ \dots & \dots & \dots & \dots & \dots \\ x_{k1} & x_{k2} & \dots & x_{km} & l_k \end{bmatrix}, l_i \in \{C_1, C_2, \dots, C_c\} \quad (5)$$

where C_i represent the class categories, c is the number of categories, $m = N \times N$ is the number of features, and k is the size of the dataset.

SAFE uses the VGG16 model pre-trained on the ImageNet dataset as a feature extraction method.

2.4.2 Dimensionality Reduction. Dimensionality reduction aims at approximating data in high-dimensional vector spaces [34]. It is important in our context since we extract a high number of features from the images (512 to 2048). In SAFE, we used the Principal Component Analysis (PCA) dimensionality reduction method to reduce the number of features from 2048 to 100.

2.5 Autoencoders

Autoencoders (AE) are unsupervised artificial neural networks that learn how to compress and encode the data before reconstructing it from the compressed encoded version to a representation that resembles the original input as much as possible. AEs can extract features that can be used to improve downstream tasks, such as clustering or supervised learning, that benefit from dimensionality reduction and higher-level features. In other words, AEs try to learn an approximation to the identity function and, by placing various constraints on the network's architecture and activation functions, they extract useful representations [28].

Figure 5 illustrates the neural network architecture of a simple AE. It consist of four main components:

- **Encoder:** learns how to compress the input data and reduce its dimensions into an encoded representation.
- **Bottleneck:** contains the encoded representation of the input data (i.e., the extracted features vector).
- **Decoder:** reconstructs the input data from the encoded version (retrieved from the *Bottleneck*) such that it resembles the original input data as much as possible.
- **Reconstruction Loss:** the difference between the *Encoder's* input and the reconstructed version (the *Decoder's* output). The objective is to minimize such loss during training.

The objective of an AE's training process is to minimize its reconstruction loss, measured as either the mean-squared error or the cross-entropy loss between original inputs and its constructed inputs.




Fig. 5. Autoencoder Architecture

3 THE PROPOSED PIPELINES

This section presents the different pipelines that can be used to implement variants of SAFE and HUDD. The evaluated pipelines differ from the original SAFE and HUDD variants with respect to four components: Feature Extraction, Dimensionality Reduction, Clustering, and Fine-Tuning. Each pipeline is a combination of a feature extraction method (FE), a dimensionality reduction technique (DR), and a clustering algorithm (CA). When feature extraction is based on transfer learning, we distinguish between models that are fine-tuned and not fine-tuned (FT/NoFT); feature extraction approaches not based on transfer learning cannot be fine-tuned. We refer to each pipeline with the pattern $FE/[FT, NoFT]/DR/CA$, with each keyword being replaced with the name of the selected method. We depict in Figure 6 all the pipelines evaluated in our study; the different components are described in the following sections.

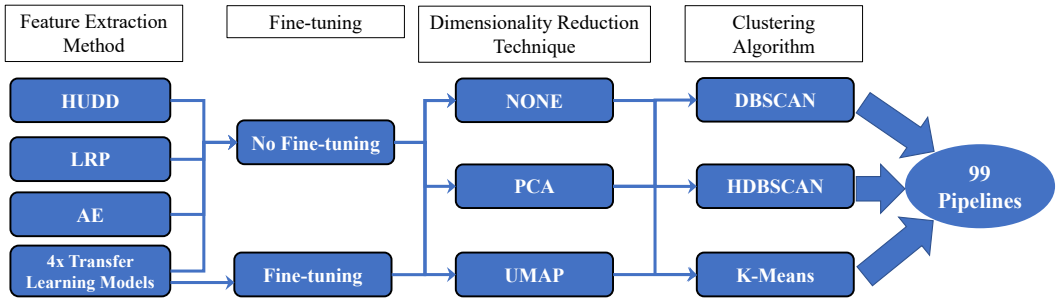


Fig. 6. Pipelines evaluated in our experiments.

3.1 Feature Extraction

3.1.1 Feature Extraction based on Transfer Learning. Several DNN architectures to extract features based on transfer learning have been proposed: Inception-V3 [84], VGGNet [79], ResNet-50 [40], and Xception [11]. These DNNs were trained on ImageNet [16], which is a dataset with more than 14 millions annotated images. The number of extracted features depends on the selected DNN architecture; Inception-V3, VGGNet-16, ResNet-50, and Xception generate 2048, 512, 2048, and 2048 features, respectively. They are described in the following paragraphs.

- **VGG-16:** VGG-16 is a Convolution Neural Network (CNN) architecture and the winner of the ILSVR (Imagenet) competition in 2014. VGG-16 focuses on convolution layers of 3×3 filters

with a stride of 1 and always uses the same padding and maxpooling layer of 2×2 with a stride of 2 instead of having a large number of hyper-parameters. It follows this arrangement of convolution and max pool layers consistently throughout the whole architecture. VGG-16 has two fully connected layers followed by a softmax layer as an output. The network has an image input size of 224×224 .

- **ResNet-50:** ResNet [40] is a CNN based on residual blocks. This architecture aims to solve vanishing gradient problems in deep neural networks. During the backpropagation process, the gradient diminishes dramatically in deep networks. Small values of gradients prevent the weights from changing their values, which slows the training process. To solve this issue, ResNet introduces residual blocks. These building blocks present skip connections between the previous convolutional layer's input and the current convolutional layer's output. Similar to VGG-16, the network has an image input size of 224×224 .
- **Inception-V3:** Inception-V3 is a refined version of Inception [85]. This network proposes additional variants of Inception blocks to reduce the number of multiplications in the convolution and minimize computational complexity. These variants are based on two factorizations: factorization into smaller convolutions and factorization into asymmetric convolutions. The network has an image input size of 224×224 .
- **Xception:** Xception is a pre-trained CNN that is 71 layers deep and can classify images into 1,000 different classes such as animals, objects, and humans. This allowed the model to learn various feature representations for a wide range of images. The Xception's input size is 299×299 .

3.1.2 Fine-tuning. Fine-tuning is a strategy to improve the performance of pre-trained models [19]. It aims to benefit from the knowledge gained from a source task and generalize it to a target task. Fine-tuning requires a pre-trained model to be prepared for reuse (i.e., the final layers may be removed and replaced with more appropriate ones), and configured to enable knowledge transfer. Knowledge transfer is achieved by freezing the shallow layers (close to the input), which learn more generic features (edges, shapes, and textures), and retraining the deeper layers (i.e., we let the DNN algorithm update the weights of the layers close to the output), which learn more specific features from the input data [19].

To fine-tune a pretrained DNN model, we follow four steps:

- (1) Create a new model whose layers (along with their weights) are cloned from the pre-trained model, except for the output layer.
- (2) Add a new fully connected output layer with a number of outputs equal to the number of classes in the target dataset, and initialize its weights with random values.
- (3) Freeze shallow layers in the network, which are responsible for the feature extraction process (to guarantee that all the important features, previously learned by the pre-trained model, are not eliminated).
- (4) Start training the new model on the target dataset, where the weights of all the non-frozen layers will keep updating using the backpropagation process. For the termination criterion, we use 100 epochs or until the loss stops improving (whichever criterion is met first).

3.1.3 Feature Extraction based on Autoencoders. As explained in Section 2.5, an *Encoder* plus a *Decoder* make up an autoencoder (AE). The input is compressed by the *Encoder*, and the *Decoder* reconstructs the input using the *Encoder*'s compressed version (at the *Bottleneck*).

Since AEs extract only the few input features necessary to aid the reconstruction of the output, the encoder might ignore other features which are not prioritized. For example in case of face images, the AE can discard the color of the skin because it is a non-prioritized feature to the AE.

However, the encoder often learns useful properties of the data [33]. The model can then receive input data from any domain, and a fixed-length feature vector obtained at the *Bottleneck* can be used for clustering. Such a vector offers a compressed version of the input data representation containing sufficient information about this data.

3.1.4 Feature Extraction based on Heatmaps. In our work, we rely on heatmaps as an additional method for feature extraction. Since heatmaps represent the relevance of each neuron on DNN outputs, failure-inducing inputs sharing the same underlying cause should show similar heatmaps.

Precisely, we rely on two different methodologies for extracting features using heatmaps, we refer to them as *LRP* and *HUDD*, according to the name of the technique driving feature extraction. *LRP* and *HUDD* have been introduced in Section 2. Feature extraction based on *LRP*, which generates heatmaps for internal layers but does not integrate a mechanism to select the most informative layer, considers the heatmap computed by the *LRP* technique for the input layer. Feature extraction based on *HUDD*, instead, considers the heatmap generated for the DNN internal layer selected by *HUDD* as the best for clustering.

3.2 Dimensionality Reduction

Several dimensionality reduction techniques exist in the literature. In this paper, we rely on two state-of-the-art techniques: Principal Component Analysis (PCA) [66] and Uniform Manifold Approximation and Projection (UMAP) [58]. PCA is used for its simplicity of implementation and because it does not require much time and memory resources. UMAP is used for its effectiveness when applied before clustering. UMAP groups data points based on relative proximity, which optimizes the clustering results. PCA and UMAP are described below.

3.2.1 Principal Component Analysis (PCA). To reduce dimensionality, PCA creates a 2-dimensional matrix of variables and observations. Then, for this matrix, it constructs a variable space with a dimension corresponding to the number of variables available. Finally, it projects each data point onto the first few maximum variance directions in the variable space. This procedure allows PCA to obtain a lower-dimensional data representation while maximizing data variation. The first principal component can equivalently be defined as the direction that maximizes the variance of the projected data [78]. In our context, we reduce the features for all our evaluated pipelines to 10 components. We empirically obtained this number in a preliminary investigation conducted with one of our case study subjects (i.e., HPD). Precisely, we executed a clustering algorithm (K-means) multiple times; each execution was performed with a set of features obtained by applying PCA with a different number of components. We evaluated all the clustering solutions using the Silhouette Index [75] (see Section 3.3) and chose the number of components yielding the highest index value.

3.2.2 Uniform Manifold Approximation and Projection (UMAP). Uniform Manifold Approximation and Projection (UMAP) is a dimension reduction technique that can be used for visualization but also for general non-linear dimension reduction. UMAP is fast, and scaling well in terms of both dataset size and dimensionality. The main limitation of UMAP is that it doesn't preserve the density of the data, which is, instead, better preserved by PCA.

First, UMAP forms a weighted graph representation between each pair of data points, where the edge weights are the probability of two data points being connected to each other. This graph is obtained by extending a radius outward each data point such that two data points are connected if their radii overlap. However, since an underestimation of such a radius can lead to the generation of small, isolated clusters, and its overestimation can lead to connecting all data points together, UMAP selects such a radius locally. The radius selection is performed based on the distance from each data point to its ' $n - th$ ' neighbor. Finally, UMAP decreases the likelihood of two data points

getting connected past the first neighbor (as the radius grows larger), which preserves the balance between the high-dimensional and low-dimensional representations. Once the high-dimensional graph is constructed, UMAP optimizes the layout of a low-dimensional representation to be as similar as possible. The general idea is to initialize the low-dimensional data points and then move them around until they form clusters that have the same structure as the high-dimensional data, preserving the connectedness of the data points. UMAP calculates Similarity Scores (distances) in the high dimensional graph to help identify the clustered points and tries to preserve that clustering in the low dimensional graph.

Since UMAP can keep the structure of the data, even in a 2-dimensional (2D) space, we reduce the number of features to two. We ran an ablation study where we evaluated clustering results obtained for one of our case study subjects (HPD presented in Section 4.1), considering both the original high-dimensional feature space and its reduced 2D counterpart. The average silhouette index, which we employ to measure the quality of the clusters, was similar in both cases. Such consistency indicates that the 2D representation sufficiently captures the relevant structure of the data, yielding cluster qualities comparable to those achieved in higher dimensions.

3.3 Clustering algorithms

In this study, we rely on three well-known clustering algorithms, K-means [55], DBSCAN [23], and HDBSCAN [57] described below. These three clustering algorithms were chosen after preliminary experiments including also the Hierarchical Agglomerative Clustering (HAC) [62] and the Mean Shift algorithm [30]. For our study, we selected algorithms that were already integrated in HUDD and SAFE (i.e., HAC and DBSCAN) but also natural extensions of SAFE (i.e., relying on HDBSCAN) and mean-based algorithms (i.e., K-means and Mean Shift). When generating clusters for one of our subjects (i.e., HPD, see Section 4.1), HAC and Mean Shift yielded much lower values of the Silhouette Index than the DBSCAN, HDBSCAN, and K-means algorithms; therefore, we discarded HAC and Mean Shift from our selection. Since a clustering algorithm may require the manual selection of parameters' values, such as the number of clusters (K-means) or the minimum distance between data points (DBSCAN), we rely on an internal evaluation metric (the Silhouette Index [75]) and the *knee-point method* [76] to automate the selection of such values.

The *Silhouette Index* is a standard practice in cluster analysis that maximizes cohesion (i.e., how closely related objects are in a cluster) and separation (i.e., how well-separated a cluster is from other clusters).

The *knee-point method* automates the *elbow method* heuristics [88] by fitting a spline to the raw data using univariate interpolation, normalizing min/max values of the fitted data, and selecting the knee-points at which the curve most significantly deviates from the straight line segment that connects the first and last data point. We rely on the *knee-point method* to automatically select the optimal number of clusters for the K-means algorithm.

3.3.1 K-means: K-means is a well-known clustering algorithm. It takes a number K as input and divides the data into K clusters based on the distance calculated from the data points to the center of the cluster. This algorithm's main function is to minimize the distance between the data points and their cluster center as much as possible. In the original K-means algorithm, the number of clusters (K) is set manually, which can affect the quality of the clusters since we don't have any prior knowledge of the data (i.e., in our context, engineers cannot know in advance how many root causes of failures should be identified).

To select an optimal value of K , we rely on the knee-point method. Precisely, we cluster the data with different values of K (in our evaluation, we consider the range $[5 - 35]$). For each clustering result, we compute the within-cluster sum of squared errors (SSD), which is the sum of the distances




Fig. 7. Approximating the optimal number of clusters K using the Knee-point method. In this case the optimal K is equal to 6.

of each point to its cluster center. We then apply the knee-point approach to these SSDs and their respective K . Figure 7 shows an example of K -approximation using the knee-point method.

3.3.2 DBSCAN: DBSCAN (Density-Based Spatial Clustering of Applications with Noise) [23], is an algorithm that defines clusters using local density estimation. It can be divided into four steps:

- (1) The ϵ -neighborhood of a data point is determined as the set of data points that are at most ϵ distant from it.
- (2) If a data point has a number of neighbors, above a configurable threshold (called *MinPts*), it is then considered a core point, and a high-density area has been detected.
- (3) Since core points can be in each other's neighborhoods, a cluster consists of the set of core points that can be reached through their ϵ -neighborhoods and all the data points in these ϵ -neighborhoods.
- (4) Any data point that is not a core point and does not have a core point in its neighborhood is considered noise.

To obtain clusters using DBSCAN, we need to select two configuration parameters: (1) the distance threshold, ϵ , to determine the ϵ -neighborhood of each data point, and (2) the minimum number of neighbors, *MinPts*, needed for a data point to be considered a core point. For the identification of the values for ϵ and *MinPts*, we rely on the same strategy integrated in SAFE, described below.

We determine the optimal value for ϵ by first computing the Euclidean distance from each data point to its closest neighbor. Then, we identify the optimal ϵ value as the knee-point of the curve obtained by considering those distances in ascending order.

To select an optimal *MinPts* value, we execute DBSCAN multiple times with varying *MinPts* values and with an ϵ equal to the optimal value determined above. We then select the clustering configuration that corresponds to the highest Silhouette Index value.

3.3.3 HDBSCAN: HDBSCAN (Hierarchical Density-Based Spatial Clustering of Applications with Noise) is an extension of DBSCAN to solve its main limitation: selecting a global ϵ . DBSCAN uses a single global ϵ value to determine the clusters. When the clusters have varying densities, using one global value can lead to a suboptimal partitioning of the data. Instead, HDBSCAN overcomes such a limitation by relying on different ϵ values for each cluster, thus finding clusters of varying densities.

HDBSCAN first builds a hierarchy using varying ϵ to figure out which clusters end up merging together and in which order. Based on the hierarchy of the clusters, HDBSCAN selects the most persisting clusters as final clusters. Cluster persistence represents how long a cluster stays without splitting when decreasing the value of ϵ . After selecting a cluster, all its descendants are ignored.

Figure 8 shows an example of the clusters' hierarchy found by HDBSCAN. The y -axis represents the values of ϵ . Vertical bars represent clusters; the color and width of each vertical bar depict the size of the cluster. We can notice that certain clusters split after the value of ϵ is increased, while others persist. HDBSCAN decides which subclusters to select based on their persistence. The persistence of a subcluster is captured by the length of the colored vertical bars in the plot. HDBSCAN selects the clusters having the highest persistence. The unselected data points are considered noise. In our example, only 6 clusters are selected (circled bars); they are the longest vertical bars in the hierarchy.




Fig. 8. Example clusters selected by HDBSCAN.

4 EMPIRICAL EVALUATION

In this section, we aim to evaluate the pipelines presented in Section 3. A pipeline leads to the generation of clusters of images that are visually inspected by safety engineers to determine the root cause captured by each. We assume that a root cause can be described in terms of the commonalities across the images in a cluster; each root cause is thus a distinct scenario in which the DNN may fail (hereafter, *failure scenario*). The pipeline that best support such process should be the one requiring

minimal effort towards accurate identification of root causes. Therefore, the best pipeline is the one that generates clusters having a high proportion of similar images (to facilitate the identification of the root cause, based on analyzing similarities across images in a cluster), enable the detection of all the root causes of failures, and is reliable in the presence of rare root causes of a particular root cause (to avoid ignoring infrequent but unsafe failure scenarios). Based on the above, we defined three research questions to drive our empirical evaluation:

RQ1 *Which pipeline generates root cause clusters with the highest purity?* We define a *pure* cluster as one that contains only images representing the same failure scenario. Such clusters are expected to be easier to interpret; indeed, the engineer should more easily determine the root cause of failures if all the images share the same characteristics. Therefore, the best pipeline is the one that leads to clusters with the highest degree of purity. The purity of a cluster is computed as the maximum proportion of images belonging to a same failure scenario in this cluster.

RQ2 *Which pipelines generate root cause clusters covering more failure scenarios?* This research question investigates to which extent the different pipelines miss failure scenarios. Ideally, all failure scenarios should be captured by one or more clusters. We say that a failure scenario is covered by a cluster if a majority of the images in the cluster belong to the scenario; indeed, commonalities shared by most of the images in a cluster should be noticed by engineers during visual inspection. We aim to determine which pipeline maximizes such coverage.

RQ3 *How is the quality of the generated root cause clusters affected by infrequent failure scenarios?* Some failure scenarios may be infrequent but are nevertheless important to identify as they may lead to severe hazards once the DNN is deployed in the field. Ideally, a pipeline should be able to produce high-quality clusters even when a small number of images belong to one or more failure scenarios. In this research question, we vary the number of images belonging to failure scenarios and study how the effectiveness of pipelines—purity and coverage of the generated clusters—is affected.

RQ4 *How do pipelines perform with failure scenarios that are not synthetically injected?* The only way to know what are the failures scenarios affecting our subject DNNs, for RQ1 to RQ3, is to rely on test set images presenting alterations (e.g., blurriness) that DNN cannot process (e.g., because it was not trained on such images). However, the results observed with injected failure scenarios may not generalize to pre-existing failure scenarios (i.e., scenarios that the original DNN cannot properly handle despite being trained for them). This research question assesses if the pipelines that perform best with injected failure scenarios also perform best with pre-existing failure scenarios and vice-versa.

To perform our empirical evaluation, we implemented our pipelines' components using different libraries. Feature extraction based on LRP was implemented with PyTorch [71], Tensorflow [1], and Keras [12] as an extension of the DNNs under test whereas transfer learning models were implemented using Tensorflow and Keras. The clustering algorithms and the dimensionality reduction methods rely on the Scikit-Learn library [67]. All the experiments were carried out on an Intel Core i9 processor running macOS with 32GB RAM. Additionally, in our experiments, we relied on the LRP implementation provided by LRP authors [59] for well-known types of layers (i.e., MaxPooling, AvgPooling, Linear, and Convolutional layers).

4.1 Subjects of the study

To evaluate our pipelines, we consider four different DNNs that process synthetic images in the automotive domain. These DNNs support gaze detection, drowsiness detection, headpose detection, and unattended child detection, which are subjects of ongoing innovation projects at IEE Sensing, our industry partner developing sensing components for automotive. Additionally, we consider two

DNNs that process real-world images to support autonomous driving: steering angle prediction and car position detection.

The gaze detection DNN (GD) performs gaze tracking; it can be used to determine a driver's focus and attention. It divides gaze directions into eight categories: TopLeft, TopCenter, TopRight, MiddleLeft, MiddleRight, BottomLeft, BottomCenter, and BottomRight. The drowsiness detection DNN (OC) has the same architecture as the gaze detection DNN and relies on the same dataset, except that it predicts whether the driver's eyes are open or closed.

The head-pose detection DNN (HPD) is an important cue for scene interpretation and computer remote control, such as in driver assistance systems. It determines the pose of a driver's head in an image based on nine categories: straight, rotated left, rotated right, rotated top left, rotated bottom right, rotated right, rotated top right, tilted, and headed up.

The unattended child detection DNN is trained with the *Synthetic dataset for Vehicle Interior Rear seat Occupancy* detection (SVIRO) [14]. SVIRO is a dataset generated by IEE Sensing that represents scenes in the passenger compartment of ten different vehicles. The dataset has been used to train DNNs performing rear seat occupancy detection using a camera system. The original IEE's DNN classifies SVIRO images into seven classes: adult, child, infant, child seat (empty or occupied), and infant seat (empty or occupied). However, the trained IEE DNN cannot be made publicly available for replication studies; therefore, in our study, we use SVIRO to retrain IEE's DNN from scratch with only three output classes (i.e., *empty seats*, *children/infants not accompanied by adults*, and *the presence of an adult*). For our classification task, we relabeled the SVIRO dataset as follows. A seat is empty when there is an object, an empty child/infant seat, or nothing. We consider the presence of a child/infant and the presence of an adult as distinct classes. IEE's DNN architecture is opensource [17], it follows a VGG-19 architecture and we retrained it for 2,000 epochs, with a batch size of 64.

Steering Angle Prediction (SAP) datasets are commonly used in autonomous driving or vehicle control systems [20]. These datasets are designed to train machine learning models to predict the appropriate steering angle for a given input image. The steering angle is a crucial parameter that determines the direction in which a vehicle should turn. The images can represent different perspectives of the road ahead, including images from a front-facing camera, multiple camera angles, or even side or rear cameras. For example, an image in the dataset could show the view of the road ahead from the driver's perspective.

For Steering Angle Prediction, we rely on the pre-trained Autumn DNN model [70], which follows the DAVE-2 architecture [6] provided by NVIDIA. It is a DNN to automate steering commands of self-driving vehicles [90]; it predicts the angle at which the wheel should be turned. It has been trained on a dataset of road images captured by a dashboard camera mounted in the car.

Car Position Detection (CPD) DNNs are used by most Advanced-Driver Assistance Systems (ADAS) to predict the positions of the cars in the scene [94]. For example, a dataset could include images captured from different angles or heights, representing various driving scenarios like urban environments, highways, or parking lots. The goal is to predict the position of each car on the scene. We rely on the CenterNet DNN [21], which is an accurate DNN used by most competition-winning approaches for object detection [91]. It has been trained on images from the ApolloScape dataset [42] collected using a dashboard camera to estimate the absolute position of vehicles with respect to the ego-vehicle.

For each subject DNN, we apply our pipelines to a set of failure-inducing images. Such sets consist of (1) images belonging to a provided test set and leading to a DNN failure and (2) test set images that were not leading to a DNN failure but had been modified to cause a DNN failure; the latter are images with injected root causes of failures and are described in Section 4.2. In classifier DNNs (i.e., OC, GD, HPD, and SVIRO) a failure occurs in the presence of an image being incorrectly

Table 1. Case Study Systems

DNN	Data Source	Training Set Size	Test Set Size (Accuracy)	Failure inducing images	# M ¹	# N ²	# H ³	# B ⁵	# SG ⁶	# EG ⁷	# EO ⁸	# S ⁹	# D ¹⁰	# NF ¹¹
GD	UnityEyes	61,063	132,630 (96%)	5,371	-	80	-	80	-	-	-	80	80	80
OC	UnityEyes	1,704	4,232 (88%)	506	-	20	-	20	-	-	-	20	20	20
HPD	Blender	16,013	2,825 (44%)	1,580	90	90	90	90	90	90	-	90	90	90
SVIRO	Blender	15,489	3,427 (74%)	884	-	30	-	30	-	-	30	30	30	30
SAP	Autopilot [8]	33,808	45,406 (84%)	7,169	-	90	-	90	-	-	-	90	90	90
CPD	Apollo [42]	5,208	4,996 (91%)	444	-	90	-	90	-	-	-	90	90	90

¹ Mask ² Noise ³ Hand ⁵ Blurriness ⁶ SunGlasses ⁷ EyeGlasses ⁸ Everyday Object ⁹ Scaling ¹⁰ Darkness ¹¹ No Injected Fault

classified. For SAP and CPD, which are regression DNNs, we set a threshold to determine DNN failures. For SAP, we observe a DNN failure when the squared error between the predicted and the true angle is above 0.18 radian (10.3°), which is deemed to be an acceptable error in related work [89]. For CPD, since it tackles a multi-object detection problem, we report a DNN failure when the result contains at least one *false positive* (i.e., the distance between the predicted position of the car and the ground truth is above 10 pixels [80]).

In Table 1, we provide details about the case study subjects used to evaluate our pipelines. For each subject, we report the source of the dataset (e.g., the simulator used to generate the data), the training and test set sizes, the accuracy of the DNN on the original test set, the number of failure-inducing images and the number of images for each injected root cause (they are detailed in Section 4.2).

We fine-tune the pipelines relying on transfer learning using the test sets of the respective case studies. We use the resulting fine-tuned model to extract the features from the failure-inducing sets. We train on the test sets because the number of images in each set is sufficient for the model to learn the features. Also, we train the autoencoders on the training set, and use the test set of the respective case study to validate the results. The termination criterion is 50 epochs unless we reach an early stopping point (the model stops improving). After training, we use only the encoder part to extract the features from the images in the failure inducing set.

4.2 Injected Failure Scenarios

To assess the ability of different pipelines to generate clusters that are pure and cover all the root causes of failures, we need to know the root causes of failures in the test set. Since such root causes may vary (e.g., lack of sufficient illumination, presence of a shadow in a specific part of the image) and it is not possible to objectively demonstrate that a failure cause has been correctly captured by a cluster (e.g., some readers may not agree that certain images show lack of sufficient illumination), to avoid introducing bias and subjectivity in our results, we modify a subset of the provided test set images so that they will fail because of known root causes of failures. In total, we considered nine different root causes to be injected in our test set images and refer to them as *injected failure scenarios* (i.e., failure scenarios with injected root causes).

We derive an image belonging to an injected failure scenario by modifying a test set image according to the specific root cause we aim to inject; for example, by covering the mouth of a person with a mask. To ensure that a modified image leads to a DNN failure because of the injected root cause, we modify only test set images that, before modification, lead to a correct DNN output.

Figure 9 illustrates the different injected failure scenarios. Below, we describe the nine root causes considered in our study:

- **Hand:** The presence of a hand blocking the full view of the driver's face could affect the DNN result, leading it to mispredict the driver's head direction. We simulate a hand that is partially covering the face appearing in the image.
- **Mask:** Similar to *Hand*, the presence of a mask covering the nose or the mouth may affect a DNN that recognizes the driver's head pose. Using image key points, we drew the shape of a white mask to simulate a mask covering the nose and the mouth.
- **Sunglasses:** As for the *Mask*, we use the eyes' key points to draw sunglasses covering the driver's eyes.
- **Eyeglasses:** Different from the *Sunglasses*, we draw glasses with the eyes being still visible.
- **Noise:** A noisy image is one that contains random perturbations of colors or brightness. In real-world automotive systems, such a failure scenario occurs due to a defective camera or a high signal-to-noise ratio (SNR) in the communication channel between different electronic control units (ECUs), resulting in a noisy input. Also, some image compression algorithms, particularly those used in certain file formats like JPEG, can introduce artifacts and noise into the image during the compression process [39, 53]. Related work has considered this failure scenario to assess the fault tolerance of DNNs [101]. We use the Scikit-Image library [93] to add Gaussian Noise, a statistical noise with a probability density function equal to a normal distribution, also known as Gaussian Distribution.
- **Blurriness:** This scenario can occur because of camera shake, especially when the camera is integrated into the car. Motion blur can also happen when capturing moving objects such as cars and pedestrians. This failure scenario was used to evaluate DNN robustness [89]. We use the Pillow library [13] to add blurriness to images using a radius of 30 pixels.
- **Darkness:** In practice, poor lighting conditions could make the DNN fail because it cannot clearly recognize what is depicted in a relatively dark image. This failure scenario was used in a related work to evaluate DNN robustness [68]. We use the Pillow library [13] to decrease the brightness of images by a factor of 0.3; we selected 0.3 because it is the lowest value introducing failures in our subject results.
- **Scaling:** Such a failure scenario mimics the situation where a camera is misconfigured, leading to rescaled images being fed to the DNN. We reduce the size of an image by a value based on the image size (i.e., large 1200px × 1200px images are scaled by 400px, small 320px × 320px images by 70px). and insert a black background using the Pillow library [13]. Camera malfunctions or technical issues with the zooming mechanism can result in a scaling failure. Scaling was also used in the literature for data augmentation [10].
- **Everyday Object:** For the SVIRO dataset, we introduce, in the car's rear seat, an everyday object (e.g., a washing machine or a handbag) never observed in the training set, thus simulating the effect of an incomplete training dataset. Such objects capture the case of an unseen label during the training, which is a commonly used faulty scenario [43].



Fig. 9. Injected failure scenarios in our study.

For regression DNNs (SAP and CPD), we randomly selected 90 images to be copied and modified for each failure scenario. For classifier DNNs, for each failure scenario, we randomly selected 10 images for each class label.

4.3 Pre-existing Failure Scenarios

Since it is usually not possible to achieve perfect accuracy through training, our DNNs, like any machine learning model, are affected by failure scenarios whose effects are visible in the original test set (e.g., borderline images that are misclassified because they are very similar to the ones belonging to another class). In other words, some failure scenarios could already be identified in the original test set and we refer to such scenarios as *pre-existing failure scenarios*.

Unfortunately, it is not possible to identify pre-existing failure scenarios in a test set because commonalities across failure-inducing images might be partially perceptible (e.g., shadows on faces) and, consequently, it might be difficult to precisely determine the causes for such failures. Therefore, we cannot perform an accurate assessment of our pipelines on pre-existing failure scenarios. However, for some of the subject DNNs classifying simulator images, it is possible to make assumptions on some of the possible causes of DNN errors; such causes can be expressed in terms of simulator parameters leading to borderline cases that are likely hard to classify by a DNN. We refer to such parameters as *failure-inducing parameters*. For each failure-inducing parameter, it is possible to identify one or more *unsafe values*. We then generate images that are likely to cause a DNN failure by configuring the simulator with a value for a failure-inducing parameter close to an unsafe value.

Table 2. Failure-inducing parameters considered to address RQ4

DNN	Parameter	Unsafe values
GD,OC	Gaze Angle	Values used to label the gaze angle in eight classes (i.e., 22.5°, 67.5°, 112.5°, 157.5°, 202.5°, 247.5°, 292.5°, 337.5°).
	Openness	Value used to label the gaze openness in two classes (i.e., 20 pixels).
	H_Headpose	Values indicating a head turned completely left or right (i.e., 160°, 220°)
	V_Headpose	Values indicating a head looking at the very top/bottom (i.e., 20°, 340°)
	DistToCenter	Value below which the eye is looking middle center (i.e., 11.5 pixels).
	PupilToBottom	Value below which the pupil is mostly under the eyelid (i.e., -16 pixels).
	TopToPupil	Value below which the pupil is mostly above the eyelid (i.e., -16 pixels).
HPD	Headpose-X	Boundary cases (i.e., -28.88°, 21.35°), values used to label the headpose in nine classes (-10°, 10°).
	Headpose-Y	Boundary cases (i.e., -88.10°, 74.17°), values used to label the headpose in nine classes (-10°, 10°).

In our previous work [4], we have identified a set of failure-inducing parameters affecting the HPD, OC, and GD DNNs; they are listed in Table 2. For GD, we identified unsafe values related to the angle of the eye gaze (8 values) and the openness of the eye (1 value) because they all may affect gaze detection results. For OC, we consider the openness of the eye (1 unsafe value), which directly affects classification, and values characterizing an unrealistic image, with a pupil below the eyelid (i.e., a distance between the pupil and the bottom eyelid below -16 pixels) or above the eyelid (i.e., a distance between the top eyelid and the pupil below -16 pixels). For HPD, we consider the Horizontal and Vertical Headpose parameters which represent the classification classes of the DNN (8 unsafe values). For Gaze Angle, Openness, Headpose-X, and Headpose-Y, the value of a failure-inducing parameter is considered close to an unsafe value if the difference between them is below 25% of the length of the subrange including the average value. For PupilToBottom, and TopToPupil, the value of a failure-inducing parameter is considered close to an unsafe value if it is below or equal to it. Table 3 provides the list of failure inducing scenarios for each subject DNN; basically, we have one failure scenario for each unsafe value except for the unsafe values of PupilToBottom and TopToPupil, which capture the same unsafe scenario (i.e., unrealistic image). Table 3 also reports the number of failure-inducing test set images belonging to each pre-existing failure scenario; note that an image can belong to one or more pre-existing failure scenarios and it was not possible to associate every image to a failure scenario. For example, this may happen because the DNN failure is due to the rendering of the image (e.g., a shadow may affect how the shape of the nose is perceived by the DNN), which is not controllable through simulator parameters but is the result of complex interactions among them (e.g., illumination direction, head orientation, light intensity).

For the HPD, OC, and GD DNNs, we could determine unsafe values for each failure-inducing parameter, because we know what are the simulator parameter values used to generate each image. For the SVIRO case study, we could not identify failure-inducing parameters because we only have access to the dataset, not the parameters associated with each image. Therefore, the possible reasons for misclassification (e.g., object size) cannot be directly mapped to the information provided to us, which is coarse grained (e.g., presence of an object on the seat).

Since we cannot know what are all the failure scenarios in our case study subjects, we do not compare our pipelines based on pre-existing failure scenarios. However, for completeness, in Section 4.7, we report on the performance of our pipelines with such failure scenarios affecting the OC, GD, and HPD DNNs.

For the experiments with injected failure scenarios (i.e., experiments assessing RQ1 to RQ3), we still include images belonging to pre-existing failure scenarios into the dataset since they

Table 3. Size of the failure-inducing set and the distribution of pre-existing failure scenario for each case study.

	GD	OC	HPD
# of failure inducing images	4937	283	865
Unrealistic	/	279	/
Openness 25	795	192	/
Border 337.5	726	/	/
Border 22.5	619	/	/
Border 67.5	386	/	/
Border 112.5	564	/	/
Border 157.5	635	/	/
Border 202.5	685	/	/
Border 247.5	486	/	/
Border 292.5	665	/	/
Headpose-X: -10	/	/	55
Headpose-X: +10	/	/	319
Headpose-X: -28	/	/	6
Headpose-X: +21	/	/	16
Headpose-Y: -10	/	/	431
Headpose-Y: +10	/	/	163
Headpose-Y: -88	/	/	0
Headpose-Y: +74	/	/	14

are usually observed for any DNN and, therefore, should be considered when generating RCCs. However, clusters that do not include any image belonging to an injected failure scenario are assumed to capture root causes related to pre-existing failure scenarios and, therefore, are ignored for computing purity and coverage (details are provided in the next Sections).

For RQ1-3, since we cannot make assumptions about the distribution of pre-existing and other failure scenarios, we include the same number of images for pre-existing failure scenarios and injected failure scenarios (see Table 1). For the experiments assessing pipelines with pre-existing failure scenarios (i.e., RQ4), instead, to be realistic, we consider the whole set of failure-inducing test images belonging to a pre-existing failure scenario.

4.4 RQ1: Which pipeline generates root cause clusters with the highest purity?

4.4.1 Design and measurements. A pure cluster includes only images presenting the same root cause (i.e., common cause leading to a DNN failure); for example, a hand covering a person's mouth. Pure clusters simplify root cause analysis because they should make it easier for an engineer to determine the commonality across images and therefore the cause of failures.

Since the likely root cause of the failure in our *injected failure scenarios* is known, we focus on these scenarios to respond to RQ1. For each RCC, we compute the proportion of images belonging to each injected failure scenario. Therefore, we measure the purity P of a cluster C (hereafter, P_C) as the highest proportion of images belonging to one injected failure scenario $f \in F$ assigned to cluster C , where F is the set of all failure scenarios. P_C is computed as follows:

$$P_C = \max_{f \in F} \left(\frac{C_f}{|C|} \right) \quad (6)$$

The proportion of a failure scenario f in a cluster C is computed as the number of images belonging to f assigned to cluster C (C_f), divided by the size of cluster C .

Clusters that do not include any image belonging to an injected failure scenario are assumed to capture root causes due to pre-existing failure scenarios and, consequently, are excluded from our analysis.

We study the purity distribution across RCCs generated for the different case study subjects. Since, ideally, we would like to obtain pure clusters, the best pipeline is the one that maximizes the average purity across the generated RCCs.

4.4.2 Methodology. We use the Conditional Inference Tree (CTree) algorithm [41] to generate a decision tree with a maximum depth set to 4 (we have four components in a pipeline) and a minimum split set to 10 (i.e., the weight of a node to be considered for splitting). The dataset used to build the tree consists of the components of each pipeline as attributes, and the purity of the generated clusters as the predicted outcome. The dataset size is equal to 99, the number of pipelines. We rely on decision trees because they enable us to determine how the different pipeline components contribute to the results (i.e., precision); the manual inspection of the configurations leading to the highest precision would not have enabled us to determine which components contribute most to precision.

Each node of the tree represents a feature of the pipeline. Leaves (terminal nodes) depict box plots representing distributions of the average purity across RCCs generated by the pipelines belonging to each leaf. Each point in the box plot is the average purity of one pipeline (i.e., the average of the purity of all the RCCs generated across all case study subjects). To split a node, the CTree algorithm first identifies the feature with the highest association (covariance) with the response variable (purity, in our case). Precisely, it relies on a permutation test of independence (null hypothesis) between any of the features and the response [83]; the feature with the lowest significant p-value is then selected ($\alpha = 0.05$, in our case). Once a feature has been selected, a binary split is then performed by identifying the value that maximizes the test statistics across the two subsets. Since we are in the presence of multiple hypotheses (assume m , for each node), to prevent a Type I error, for each feature j , CTree computes its Bonferroni-adjusted [99] p -value $_j$ as

$$\text{adjusted } p\text{-value}_j = 1 - (1 - p\text{-value}_j)^m$$

4.4.3 Results. Figure 10 depicts a regression tree illustrating how the different components of a pipeline (feature extraction methods, fine-tuning, dimensionality reduction techniques and clustering algorithms) determine the purity of the clusters generated by a pipeline. We notice that the pipelines with fine-tuned models (Node 3 and 4) generate lower-purity clusters than those without any fine-tuning (Node 6 and 7), which can be explained by the fine-tuning dataset not including the injected failure scenarios. For our approach, the objective of fine-tuning is to learn features that are specific for the context of use; please recall that our transfer learning models are based on ImageNet and we rely on them for feature extraction. However, we perform fine-tuning using the test set, which is smaller than the training set and thus leading to a quicker process. Further, to simulate a realistic usage, we did not include the injected failures into the dataset used for fine-tuning; indeed, since our injected root causes aim to capture scenarios not foreseen at training time, it would be unrealistic to consider such scenarios for fine-tuning. Finally, fine-tuning with images including injected failures (e.g., noise) may affect the quality of fine-tuning. Because of the choices above, fine-tuning leads to features that do not capture the injected faults but the characteristics of images without faults. As a result, in our experiments, images are clustered based only on their pre-existing fault (e.g., borderline class) instead of the injected faults. ImageNet models, instead, may capture features that are useful to cluster injected faults (e.g., the presence of everyday objects in SVIRO), but such features are then forgotten as an effect of catastrophic

forgetting during fine-tuning [9], thus leading to clustering results that are worse for fine-tuned transfer-learning models.

The pipelines using non-fine-tuned transfer learning models as a feature extraction method (Node 7) generate purer clusters (min = 50%, median = 80%, max = 96%) than the pipelines using an autoencoder model, HUDD, or LRP (Node 6) (min = 50%, median = 65%, max = 70%). The purpose of the Autoencoder model is to provide a condensed representation of the image to be used for reconstruction. This is done by ignoring the features that the model considers insignificant and only keeping the features that help the encoder reconstruct the image accurately. Therefore, a possible explanation for our result is that since the autoencoder is trained on the training set, the injected faults are ignored. Given that clustering is based on the condensed representation, the generated clusters are less pure than the ones generated by the pipelines with transfer learning models. Note that without empirical assessment, it is not possible to know in advance how autoencoders support clustering; indeed, injected faults may mask certain autoencoder features (e.g. presence of non-black pixels around the borders for scaled images) that turn out to be useful for clustering.

As for HUDD and LRP, it seems that their main limitation is that heatmaps cannot capture the presence of root causes affecting all the pixels in an image (i.e., the result of noise, blurriness, darkness, scaling). Heatmaps mainly capture which pixels of the image drive the DNN output, thus leading clustering to group images where the same pixels affected the output. For instance, the DNN's response to a blurred image with a shadow on the mouth could be different from that of another blurred image with a shadow on the eyes, thus leading to different clusters for these images although they represent the same injected failure scenario (blurriness).

Finally, we notice that the pipelines using HDBSCAN and DBSCAN (Node 3) as a clustering algorithm yield purer clusters (min = 25%, median = 40%, max = 80%) than those using K-means (Node 4, min = 22%, median = 27%, max = 29%). This is because K-means faces difficulty dealing with non-convex clusters. A cluster is convex if, for every pair of points belonging to it, it also includes every point on the straight line segment between them [50], which gives the cluster a hyperspherical form. Nevertheless, in many practical cases, the data leads to clusters with arbitrary, non-convex shapes. Such clusters, however, cannot be appropriately detected by a centroid-based algorithm (e.g., K-means), as they are not designed for arbitrary-shaped clusters.

DBSCAN and HDBSCAN are density-based clustering algorithms. They consider high-density regions as clusters (see Section 2). The root cause clusters generated by DBSCAN and HDBSCAN are arbitrary-shaped and more homogeneous (i.e., clusters with higher within-cluster similarity) with very similar images. In contrast, a convex cluster generated by K-means tends to be less dense and can group rather dissimilar images. As a result, a convex cluster is less pure than a non-convex one.

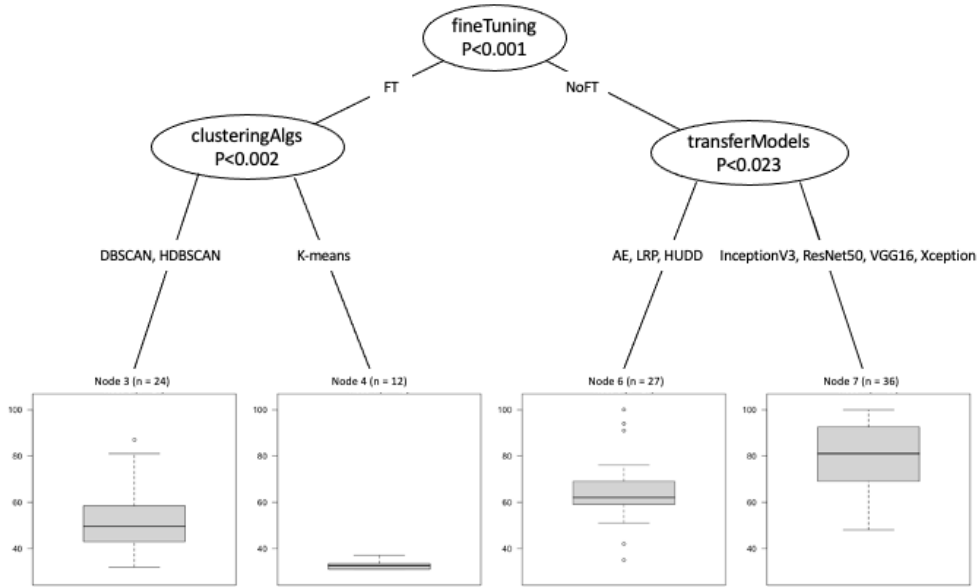


Fig. 10. Decision Tree illustrating how the different features of a pipeline determine the average purity of root-cause clusters.

We report the significance of these results in Table 4, including the values of the Vargha and Delaney’s \hat{A}_{12} effect size and the p -values resulting from performing a Mann-Whitney U-test to compare the average purity of the pipelines using transfer learning models (Node 7 in the decision tree) and the pipelines represented by the other nodes. Typically, an \hat{A}_{12} effect size above 0.56 is considered practically significant with higher thresholds for medium (0.64) and large (0.71) effects [48], thus suggesting the effect sizes between the pipelines using transfer learning models and other pipelines are large. Further, p -values suggest these differences are statistically significant.

Table 4. RQ1: p -values and and effect size values when comparing the results of the pipelines with the best purity of clusters (according to the decision tree) to the other pipelines.

	Node 3	Node 4	Node 6
p -value	7e-11	2e-7	5e-6
\hat{A}_{12}	1.00	1.00	0.80

Finally, in Table 5, we report the pipelines that generated clusters with an average purity above 90% across all case study subjects, along with the purity obtained for each subject; the complete results obtained for all pipelines appear in Appendix A, Table 14. An average purity of 100% means that all the clusters generated by the pipeline are pure. Interestingly, all the pipelines in Table 5 belong to Node 7 in Figure 10, thus confirming our main finding. Five of these seven best pipelines, rely on UMAP, without fine-tuning but with a transfer learning model, which is therefore our suggestion to perform root cause analysis. The best result is obtained with ResNet-50 combined with UMAP and DBSCAN.

Table 5. RQ1: Pipelines with a purity greater than 90%. The last column represents the average of averages.

Pipelines					Avg. purity across RCCs						
#	FE	FT	DR	CA	GD	OC	HPD	SVIRO	SAP	CPD	Avg.
19	VGG-16	NO	None	K-Means	91.7%	92.1%	95.5%	82.5%	97.3%	99.7%	93.2%
25	VGG-16	NO	UMAP	K-Means	97.6%	84.4%	93.7%	82.4%	90.3%	97.8%	91.0%
26	VGG-16	NO	UMAP	DBSCAN	99.0%	93.0%	99.6%	79.7%	98.1%	96.6%	94.3%
39	ResNet-50	NO	None	HDBSCAN	96.4%	100.0%	100.0%	78.8%	87.5%	100.0%	93.8%
43	ResNet-50	NO	UMAP	K-Means	99.4%	93.0%	82.3%	79.6%	99.6%	97.7%	91.9%
44	ResNet-50	NO	UMAP	DBSCAN	100.0%	95.8%	95.8%	79.0%	99.7%	99.3%	94.9%
62	Inception-V3	NO	UMAP	DBSCAN	93.4%	95.2%	98.1%	76.6%	97.4%	83.1%	90.7%

FE Feature Extraction *FT* Fine-tuning *DR* Dimensionality Reduction *CA* Clustering Algorithm

4.5 RQ2: Which pipelines generate root cause clusters covering more failure scenarios?

4.5.1 Design and measurements. This research question investigates the extent to which our pipelines identify all failure scenarios. We compare pipelines in terms of the percentage of injected failure scenarios being covered by at least one RCC. A failure scenario is covered by an RCC if it enables the engineer to determine the root cause of the failure. Precisely, when images belonging to a failure scenario f represent a sufficiently large share of images in a cluster C , it is easier for an engineer to notice that f is a likely root cause. Therefore, we assume that an injected failure scenario f is covered by a cluster C if the cluster C contains at least 90% of images with f . Since this threshold is relatively high, our results can be considered conservative.

Given that our injected failure scenarios are represented by the same number of images in the failure-inducing test set, every failure scenario has the same likelihood of being observed. Therefore, we expect to obtain RCCs corresponding to each failure scenario.

4.5.2 Methodology. We follow the same methodology as for RQ1 (see Section 4.4.2) but we construct a decision tree considering, for each pipeline, the average coverage achieved across case study subjects instead of the average purity.

4.5.3 Results. Figure 11 shows a decision tree illustrating how the different components of a pipeline determine the coverage of failure scenarios.

Each leaf node depicts a box plot with the distribution of the percentages of failure scenarios covered by the set of pipelines that include the components listed in the decision nodes.

For instance, Node 9 provides the distribution of the percentage of failure scenarios covered by the RCCs generated by pipelines using UMAP as a dimensionality reduction technique and non-fine-tuned transfer learning models as feature extraction methods (12 pipelines). Ideally, the root-cause clusters generated by a pipeline should cover 100% of the failure scenarios.

The decision tree in Figure 11 confirms RQ1 results. The pipelines without fine-tuning (Nodes 6, 8 and 9) outperform the pipelines with fine-tuning (Nodes 3 and 4). The pipelines with transfer learning models (Nodes 8 and 9) generate clusters that cover more failure scenarios than those generated by the pipelines using HUDD, LRP, and AE (Node 6). Also, the pipelines using the DBSCAN and HDBSCAN clustering algorithms (Node 3) yield better results than the ones using K-means (Node 4).

Further, the decision tree in Figure 11 gives us more insights into which dimensionality reduction method is more effective. We notice that the root-cause clusters generated by the pipelines using UMAP (Node 9) lead to a better distribution (min = 45%, median = 85%, max = 100%) than the pipelines using PCA or not using any dimensionality reduction (Node 8, min = 25%, median = 55%, max = 90%). This is because UMAP yields a better separation of the clusters (i.e., less clusters

overlap) compared to PCA. When using UMAP, all the data points converge towards their closest neighbor (the most similar data point). Therefore, neighboring data points in higher dimensions end up in the same neighborhood in lower dimensions, resulting in a compact and well-separated clusters where it is easier for the clustering algorithms to distinguish them.

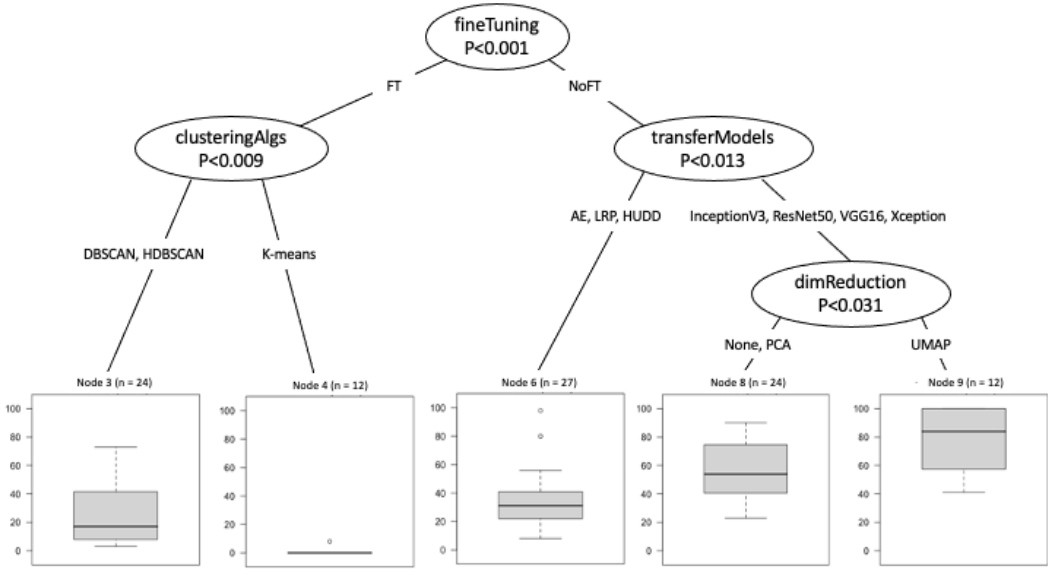


Fig. 11. Decision Tree illustrating how the different features of a pipeline determine the coverage of the failure scenarios.)

We report the significance of these results in Table 6, including the values of the Vargha and Delaney’s \hat{A}_{12} effect size and the p -values resulting from performing a Mann-Whitney U-test to compare the percentages of covered failure scenarios resulting from the pipelines using UMAP (Node 9 in the decision tree in Figure 11), and the other pipelines. Table 6 shows that the p -values, when comparing the pipelines using UMAP to the other pipelines, are always below 0.05. This implies that in all the cases, differences are statistically significant with large effect sizes (above 0.77).

Table 6. RQ2: p -values and and effect size values when comparing the results of the pipelines with the best coverage of the faulty scenarios (according to the decision tree) to the other pipelines.

	Node 3	Node 4	Node 6	Node 8
p-value	1e-5	1e-5	4e-5	8e-3
\hat{A}_{12}	0.95	1.00	0.91	0.77

In Table 7, we report the pipelines that generated clusters covering at least 90% of the failure scenarios across all case study subjects, along with the coverage obtained for each case study subject (complete results for all the pipelines are reported in Appendix B, Table 15). If the coverage is equal to 100%, all the failure scenarios are covered by the RCCs. Unsurprisingly, the pipelines in Table 7 belong to Node 7 in Figure 11: they rely on a non-fine-tuned transfer learning model

Table 7. RQ2: Pipelines with a coverage greater than 90%. The last column represents the average of averages.

Pipelines					Percentage of covered faulty scenarios						
#	FE	FT	DR	CA	GD	OC	HPD	SVIRO	SAP	CPD	Avg.
26	VGG-16	None	UMAP	DBSCAN	100.0%	100.0%	100.0%	80.0%	100.0%	100.0%	96.7%
44	ResNet-50	None	UMAP	DBSCAN	100.0%	100.0%	100.0%	60.0%	100.0%	100.0%	93.3%
62	Inception-V3	None	UMAP	DBSCAN	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
80	Xception	None	UMAP	DBSCAN	100.0%	100.0%	100.0%	40.0%	100.0%	100.0%	90.0%

FE Feature Extraction *FT* Fine-tuning *DR* Dimensionality Reduction *CA* Clustering Algorithm

for feature extraction, and UMAP for dimensionality reduction. Further, they all use DBSCAN for clustering. These pipelines consistently yielded the best results for all individual case studies (confirming the results obtained in RQ1).

Such findings are further supported by the results in Table 14 and Table 15, where we notice that the combination of UMAP with DBSCAN always achieves higher purity and coverage (in bold) than its alternatives, regardless of the used feature extraction method.

4.6 RQ3: How is the quality of root cause clusters generated affected by infrequent failure scenarios?

4.6.1 Design and measurements. We study the effect of infrequent failure scenarios on the quality of the RCCs generated by the pipelines. Indeed, infrequent scenarios may not be properly captured by clustering algorithms. With K-means, the number of clusters depends on within-cluster SSD (see Section 3.3.1) but the exclusion of small clusters may lead to unnoticeable changes in the computed SSD. With DBSCAN, small clusters may be treated as noise. With HDBSCAN, small clusters, which have a limited persistence (ϵ cannot be higher than the number of datapoints, see Section 3.3.3), may not be identified.

We consider a failure scenario infrequent when it is observed in a low proportion of the images in the failure-inducing set. To be practically useful, a good pipeline should be able to generate root-cause clusters even for infrequent failure scenarios; indeed, in safety-critical contexts, infrequent failure scenarios may lead to hazards and thus should be detected when testing the system. For instance, if only five out of hundred failure-inducing images belong to a failure scenario and we have three failure scenarios in total, a reliable pipeline should still generate an RCC containing only the images of the infrequent failure scenario.

4.6.2 Methodology. We generate 10 different failure-inducing sets for each case study subject (a total of 60 failure-inducing sets). To construct a failure-inducing set, for each root cause that might affect the case study (see Table 1, Page 17), we generate a number n of images affected by the injected root cause. We randomly select a number n that is lower than the number of images selected for the same root cause in RQ1 (see Table 1). Further, for classifier DNNs, we select a value higher than the number of classes of the corresponding case study (we enforce one root cause of failures for one image per class, at least); for regression DNNs, we select a value above 2. Since n is randomly selected (uniform distribution), we obtained failure-inducing sets containing failure scenarios whose number vary. Table 16, Appendix C, provides the details for each case study; for instance, the number of images representing a failure scenario for each failure-inducing set of the HPD case study (9 classes) is randomly selected between 9 and 90.

In addition, we also include a randomly selected number of images belonging to pre-existing failure scenarios, to mimic what happens in practice (see RQ1). The number of images belonging to pre-existing failure scenarios varies between two and the total number of injected failure scenario images.

Since we aim to study the effect of infrequent failure scenarios on the quality of the generated RCCs, we categorize our 290 failure scenarios into *infrequent* and *frequent*. Infrequent failure scenarios are the ones that include a proportion of injected images that is lower than the median proportion in all the generated failure-inducing sets (equals to 18% in our study). For example, noise is frequent in the dataset GD_1 ($64 > 18$) but infrequent in the dataset OC_2 ($4 < 18$).

We consider only the best pipelines resulting from the experiments in RQ1 and RQ2 (i.e., having purity or coverage above 90% as shown in Tables 5 and 7); they are pipeline 26 (VGG16/DBSCAN/UMAP/NoFT), 44 (ResNet50/DBSCAN/UMAP/NoFT), 62 (InceptionV3/DBSCAN/UMAP/NoFT), 19 (VGG16/K-means/None/NoFT), 25 (VGG16/K-means/UMAP/NoFT), 39 (ResNet50/HDBSCAN/None/NoFT), 43 (ResNet50/K-means/UMAP/NoFT), and 80 (Xception/DBSCAN/UMAP/NoFT). The first three pipelines (i.e., 26, 44, 62) were the best for both RQ1 and RQ2, the next four (i.e., 19, 25, 39, 43) were selected based on RQ1 results while the latter (i.e., 80) based those of RQ2. We compute the purity and coverage of the RCCs generated by each of these pipelines, following the same procedures adopted for RQ1 and RQ2. We then compare the distribution of purity and coverage for infrequent and frequent failure scenarios. The most reliable pipelines are the ones being affected the least, in terms of purity and coverage, by infrequent failure scenarios.

4.6.3 Results. In Figure 12, for each selected pipeline, we report the average purity across all the RCCs² with the injected failure scenarios having a certain frequency. The x -axis reports the proportion of images for failure scenarios whereas the y -axis reports the average purity of the RCCs associated to each failure scenario.

Figure 12 shows that when the frequency of the failure scenarios is below the median (infrequent scenario), the cluster purity obtained by pipelines tends to significantly lower and decrease rapidly as the frequency decreases. This is expected because when a failure scenario is infrequent, the clustering algorithm tends to either cluster its images as noise or distribute them over the other clusters. For density-based clustering algorithms, images belonging to infrequent scenarios may not become core points when the identification of a core point requires more data points in their neighborhood. In such case, images belonging to infrequent scenarios will be either labeled as noise points or border points (belonging to other clusters). The same is true for K-means, where these points are usually spread across other clusters because they cannot form a cluster.

To strengthen our findings, in Table 8, we report the results when comparing the purity of the selected pipelines for frequent and infrequent failure scenarios; further, we report the Vargha and Delaney's \hat{A}_{12} effect size and the p -values resulting from performing a Mann-Whitney U-test. We notice that for all pipelines, the difference between frequent and infrequent scenarios are significant (p -value < 0.05). However, the effect sizes for Pipelines 26, 62, 45, and 80 are small, while they are medium for Pipelines 19 and 44, which indicates that pipelines including DBSCAN (i.e., Pipelines 26, 62, 45, and 80) are much more reliable with infrequent scenarios than others (i.e., the difference between frequent and infrequent scenarios is less pronounced). Actually, the pipelines using DBSCAN fare better than the rest also in the general case. Indeed, almost all the injected failure scenarios with frequency above 18% have 100% purity (see Figure 12); further for infrequent failure scenarios they include less data points below 100% than the other pipelines. This is because DBSCAN tends to find clusters with different sizes if these clusters are dense enough; K-means, instead, derives clusters that are of similar size.

Further, we notice that the purity of the clusters generated by Pipeline 26 (VGG16/Dbscan/UMAP/NoFT), for infrequent failure scenarios, is higher (average is 94%) than the purity of the clusters

²As discussed in Section 4.4.1. The red vertical line represents the median frequency of failure scenarios. We say that an RCC is associated with (or captures) an injected failure scenario f when the majority of the images in the cluster belong to scenario f .

Table 8. RQ3: p-values and effect size values when comparing the purity of the best pipelines with the frequent and infrequent failure scenarios.

Pipelines	26	44	62	39	80	19	25	43
Average Purity for infrequent failure scenarios	94%	87%	91%	79%	87%	76%	70%	65%
Average Purity for frequent failure scenarios	100%	100%	100%	92%	99%	96%	96%	93%
p-value	4e-6	2e-10	1e-6	2e-9	8e-9	8e-5	2e-10	3e-14
\hat{A}_{12}	0.58	0.64	0.59	0.60	0.63	0.68	0.70	0.75

Table 9. RQ3: p-values and effect size values when comparing the best pipeline in Table 8 (i.e., Pipeline 26, VGG16/Dbscan/UMAP/NoFT) to the other pipelines based on the average purity of the clusters associated to infrequent failure scenarios.

Pipelines	44	62	39	80	19	25	43
p-value	0.002	0.51	4e-5	0.006	3e-12	2e-14	4e-21
\hat{A}_{12}	0.57	0.51	0.60	0.56	0.69	0.71	0.77

Table 10. RQ3: Fisher exact test values when comparing the coverage of the lowly represented and highly represented faulty scenarios by the clusters generated by the best pipelines.

Pipelines	26	44	62	39	80	19	25	43
Average Coverage for infrequent failure scenarios	85%	71%	82%	66%	73%	51%	46%	34%
Average Coverage for frequent failure scenarios	100%	98%	99%	86%	98%	86%	87%	77%
Fisher's Exact test	1e-5	1e-5	1e-5	1e-5	1e-5	2e-4	1e-5	1e-5

generated by the other pipelines; differences are significant (see Table 9), thus suggesting Pipeline 26 might be the best choice.

Concerning *coverage*, Figure 13 shows, for each pipeline, histograms with the average coverage obtained for failure scenarios having proportions of failure inducing images within specific ranges. In general, we observe that coverage is higher for frequent scenarios. This is due to the correlation between pure clusters and coverage; the less pure the generated clusters, the fewer failure scenarios they cover. When the failure scenarios are infrequent, their images are distributed over the other clusters, reducing their purity and, thus, reducing the probability of these scenarios being covered. To demonstrate the significance of the difference between coverage results obtained with frequent and infrequent scenarios, we apply the Fisher's Exact test³ to compare the coverage of frequent and infrequent scenarios for the clusters generated by the selected pipelines. We report the *p*-values resulting from the Fisher's Exact test in Table 10 and observe that differences are statistically significant thus indicating that pipelines perform better with frequent failure scenarios.

Further, Figure 13 shows that pipeline 62 (*InceptionV3/DBSCAN/UMAP/NoFT*) is the one performing best with the least frequent scenarios (i.e., range 0-5%) but no pipeline fares well in that range. Pipeline 26 (*VGG16/DBSCAN/UMAP/NoFT*) is the one performing best with infrequent scenarios in

³The Fisher's Exact test [92] is a statistical test used to determine if there is a non-random association between two categorical variables [96].

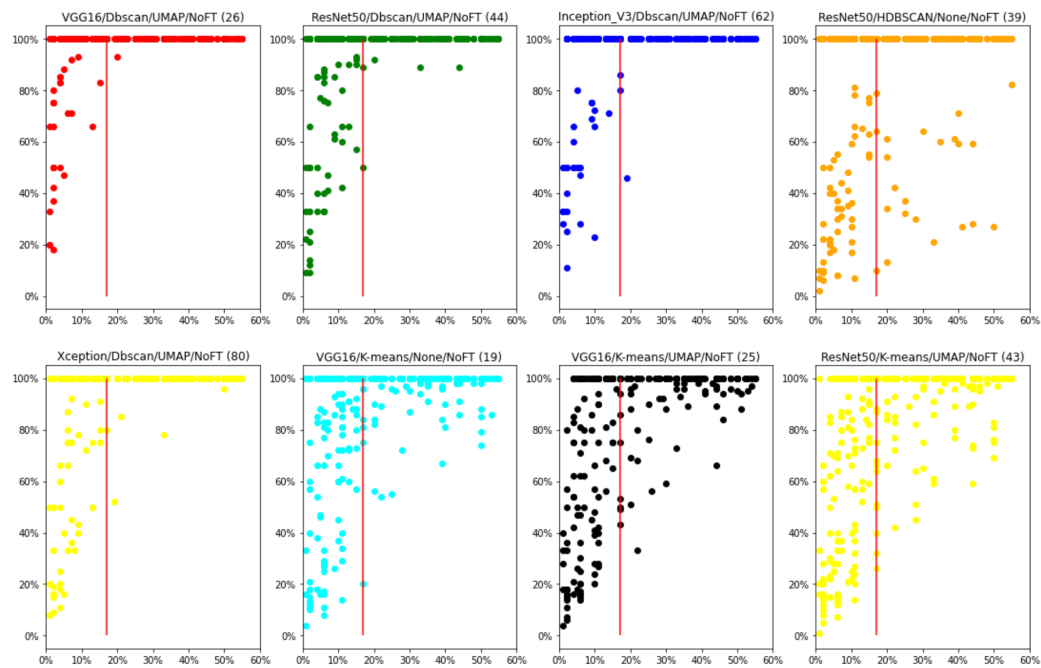


Fig. 12. Purity of the clusters associated with frequent and infrequent failure scenarios. The x-axis captures the frequency of a failure scenario (i.e., proportion of failure-inducing images for a failure scenario). Each data point is the average of all the RCCs associated to one distinct failure scenario. The red vertical line represents the median frequency of failure scenarios.

Table 11. RQ3: Fisher exact test values when comparing the best pipeline "VGG16/Dbscan/UMAP/NoFT" to the other pipelines based on the coverage of the infrequent failure scenarios.

Pipelines	44	62	39	80	19	25	43
Fisher's Exact test	4e-2	0.55	1e-5	2e-3	0.018	1e-5	1e-5

the range 5% to 20%; indeed, it is the only pipeline providing an average coverage above 90% for that range. To further demonstrate the significance of the difference in performance between Pipeline 26 and the other pipelines, we apply Fisher's exact test to the coverage obtained for infrequent scenarios. We report the p -values resulting from this test in Table 11. We notice that all the p -values are below 0.05 except when Pipeline 26 is compared to Pipeline 62; indeed, the results of these two pipelines are similar as visible in Figure 13), even though Pipeline 26 performs slightly better on average.

In conclusion, infrequent failure scenarios affect both purity and coverage; pipelines tend to perform worse when the failure scenarios are infrequent (their frequency is below the median). However, some pipelines fare better than others. Our results suggest that the pipeline relying on a non-fine-tuned VGG16 model, with UMAP and DBSCAN (Pipeline 26) is the best choice because it yields significantly higher purity and coverage than the other pipelines. Pipeline 26 is also less negatively affected by infrequent failure scenarios since its coverage is above 90% when the frequency is above 5%, which is not the case for all the other pipelines.

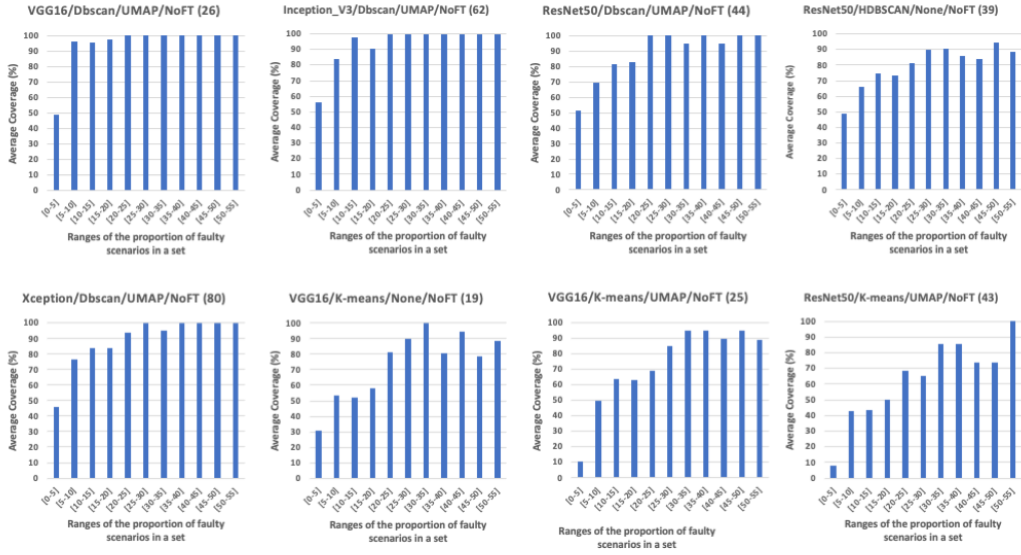


Fig. 13. Comparing the percentage of coverage across different ranges of proportions of failure scenarios in each set.

4.7 RQ4: How do pipelines perform with failure scenarios that are not synthetically injected?

4.7.1 Design and measurements. Our objective is to determine if the best pipelines identified in RQ1, RQ2, and RQ3 perform best also with pre-existing failure scenarios. As stated in Section 4.3, to address this research question we considered only the subject DNNs for which it is possible to determine the pre-existing failure scenarios each failure-inducing image may belong to; the selected DNNs are OC, GD, and HPD. The list of pre-existing failure scenarios is shown in Table 3 (page 21).

A pipeline should, ideally, identify all the pre-existing failure scenarios (i.e., generate at least one cluster for each pre-existing failure scenario thus maximizing coverage). Also, the generated clusters should be pure, that is, include only images belonging to a same pre-existing failure scenario. Consequently, as for RQ1 to RQ3, we compare pipelines based on the purity and coverage of the generated clusters.

4.7.2 Methodology. For each subject DNN, we applied all our pipelines to the set of failure-inducing images in the original test set and belonging to a pre-existing failure scenario.

As per RQ1 to RQ3, we compute coverage and purity of each cluster as follows. For each image, we know the pre-existing failure scenarios it belongs to. Therefore, for each generated cluster, we can determine the number of images belonging to each pre-existing failure scenario. Each cluster is considered to cover the pre-existing failure scenario with the largest number of clustered images; indeed, being the most frequent, the commonalities in those images are likely to be noticed by the engineer inspecting the results. For each cluster, purity is computed as the proportion of clustered images belonging to the scenario covered by the cluster.

We consider the pipelines leading to the best results for purity and coverage for RQ1, RQ2, and RQ3, and compare them with the pipelines leading to the best purity and coverage results when applied to the failure inducing images described above, across the three selected subject DNNs.

4.7.3 Results. In Table 12, we report the pipelines leading to the best purity and coverage when applied to the datasets with injected (RQ1, RQ2, RQ3) and pre-existing failure scenarios (i.e., RQ4). The values in parentheses capture the ranking of a pipeline for each dataset. For both purity and coverage, for each RQ, we rank our pipelines after sorting them in a decreasing order based on the average of the metric value computed for the OC, HPD, and GD DNNs; pipelines having the same average are assigned the same rank.

The results in Table 12 show that the pipeline with the highest coverage for pre-existing failure scenarios is Pipeline 26 (see column *Coverage-RQ4*), which confirms our findings for RQ3 (Section 4.6.3) where pipeline 26 leads to the highest coverage results when failure scenarios do not occur with the same frequency; the results observed for RQ4 can thus be explained by the fact that, in the original test set, failure scenarios do not have the same frequency. Further, Pipeline 26 achieves high purity with pre-existing failure scenarios; indeed it is ranked 4th in column *Purity-RQ4*. Interestingly, a white-box pipeline (i.e., Pipeline 8, combining HUDD, DBSCAN and UMAP) leads to the highest purity for RQ4's dataset; however, it does not lead to the best coverage (only 91%, ranked 7-th). Since in safety-critical systems one would prioritize the discovery of all failure scenarios, Pipeline 26 should be a better option than Pipeline 8; indeed, Pipeline 26 achieves top coverage while having a very high purity (87% VS 92% of Pipeline 8). Further, for pre-existing failure scenarios, Pipeline 26 is the only pipeline with a purity rank up to 4 being among the best 10 pipelines for coverage.

Pipeline 26 and Pipeline 80 are the only two pipelines being among the best ten for both purity and coverage, with pre-existing failure scenarios. Also, they are among the ten best pipelines for all the other datasets (i.e., RQ1, RQ2, and R3). More in general, the Pipelines 26, 44, 62, and 80, which are all the pipelines relying on transfer learning, DBSCAN, and not using fine-tuning, lead to top ranked results. However, only Pipeline 26 achieves the highest rank for more than one dataset, thus confirming it is a preferable choice as we suggested in Section 4.6.3.

Interestingly, four of the ten best-ranked pipelines for coverage with pre-existing failure scenarios include fine-tuning; however, they poorly perform in terms of purity. Based on our discussion in Section 4.4.3, it is predictable that fine-tuning performs better with pre-existing failure scenarios; indeed, the failure-inducing images do not differ from the ones considered for fine-tuning (i.e., fine-tuning captures features that are present in the failure-inducing test set). However, the reason fine-tuning did not help achieve clusters with high purity is its reliance on a dataset with different scenarios occurring according to very different frequencies. Indeed, fine-tuning may overfit the features belonging to the most frequent scenarios, consequently the fine-tuned autoencoder may not extract relevant features for infrequent scenarios. To conclude, fine-tuning seems not to be advisable because (1) failure scenarios, as shown in our experiment, are unlikely to include the same proportion of images, (2) it is not realistic to expect engineers to construct datasets with the same proportion of images for all failure scenarios, and (3) failure scenarios may largely differ from the images observed in the training set, which led to poor performance for fine-tuned pipelines in Section 4.4.3.

4.8 Discussion

The results of RQ1 and RQ2 show that there is a family of pipelines leading to higher purity (i.e., they simplify the identification of root causes) and coverage (i.e., they enable the identification of all root causes). Such pipelines rely on transfer learning, UMAP for dimensionality reduction,

Table 12. Purity and Coverage of the best ten pipelines on datasets with injected and pre-existing failure scenarios; the values between parentheses indicate the rank of a pipeline for a RQ's dataset. We selected the top-ten ranked pipelines based on each of the datasets considered for RQ1, RQ2, RQ3, and RQ4.

#	Best Pipelines				Purity			Coverage		
	FE	FT	CA	DR	RQ1	RQ3	RQ4	RQ2	RQ3	RQ4
5	HUDD	NoFT	Dbscan	PCA	55 (60)	57 (65)	85 (10)	25 (54)	32 (60)	86 (9)
8	HUDD	NoFT	Dbscan	UMAP	99 (5)	85 (13)	92 (1)	96 (9)	68 (14)	91 (7)
17	LRP	NoFT	Dbscan	UMAP	99 (5)	84 (15)	86 (8)	92 (10)	64 (21)	70 (15)
19	VGG16	NoFT	K-means	None	94 (14)	91 (6)	85 (10)	79 (15)	79 (8)	51 (52)
20	VGG16	NoFT	Dbscan	None	78 (24)	76 (26)	87 (3)	58 (22)	56 (25)	62 (25)
22	VGG16	NoFT	K-means	PCA	84 (19)	81 (19)	87 (4)	54 (34)	57 (23)	47 (61)
25	VGG16	NoFT	K-means	UMAP	99 (5)	90 (8)	87 (4)	100 (1)	76 (11)	51 (52)
26	VGG16	NoFT	Dbscan	UMAP	100 (1)	97 (3)	87 (4)	100 (1)	93 (4)	100 (1)
27	VGG16	NoFT	HDBSCAN	PCA	61 (50)	81 (19)	90 (2)	37 (47)	70 (13)	34 (94)
35	VGG16	FT	Dbscan	UMAP	74 (30)	72 (29)	53 (67)	42 (43)	10 (85)	77 (10)
43	ResNet50	NoFT	K-means	UMAP	97 (10)	86 (12)	62 (49)	92 (10)	71 (12)	47 (61)
44	ResNet50	NoFT	Dbscan	UMAP	100 (1)	97 (3)	87 (4)	100 (1)	94 (3)	95 (3)
53	ResNet50	FT	Dbscan	UMAP	76 (26)	66 (48)	52 (68)	42 (43)	32 (60)	86 (9)
55	InceptionV3	NoFT	K-means	None	99 (5)	91 (6)	65 (39)	100 (1)	81 (6)	47 (61)
61	InceptionV3	NoFT	K-means	UMAP	100 (1)	93 (5)	65 (39)	100 (1)	85 (5)	47 (61)
62	InceptionV3	NoFT	Dbscan	UMAP	100 (1)	98 (1)	84 (12)	100 (1)	97 (1)	94 (6)
71	InceptionV3	FT	Dbscan	UMAP	70 (37)	70 (37)	55 (63)	42 (43)	38 (57)	95 (3)
73	Xception	NoFT	K-means	None	97 (10)	89 (9)	66 (35)	100 (1)	77 (10)	51 (52)
75	Xception	NoFT	HDBSCAN	None	96 (12)	88 (11)	63 (46)	92 (10)	81 (6)	55 (48)
79	Xception	NoFT	K-means	UMAP	89 (17)	89 (9)	64 (42)	92 (10)	79 (8)	43 (72)
80	Xception	NoFT	Dbscan	UMAP	99 (5)	98 (1)	86 (8)	100 (1)	95 (2)	97 (2)
89	Xception	FT	Dbscan	UMAP	55 (60)	65 (50)	55 (63)	8 (76)	46 (43)	95 (3)
98	AE	NoFT	Dbscan	UMAP	91 (16)	77 (24)	83 (13)	67 (18)	54 (27)	76 (12)

DBSCAN for clustering, and are not using fine-tuning. Among such pipelines, considering that it is reasonable to expect unsafe scenarios to be infrequent, based on the results of RQ3, we suggest to use the pipeline relying on VGG16 (Pipeline 26) as transfer learning model. Pipeline 26 also leads to the best results when applied to pre-existing failure scenarios (RQ4), probably due to infrequent pre-existing failure scenarios.

In our study, we focused on effectiveness, not cost; indeed, our main purpose is to identify the pipeline that generates clusters that do not confuse the end-user (i.e., they are pure) and is likely to help identify all the root causes of failures (i.e., they have high coverage). In contrast, cost is related to the number of clusters being inspected. To discuss such cost, we report in Figure 14 a boxplot with the size of the clusters generated for RQ1, RQ2, RQ3, and RQ4 by Pipeline 26. As shown in Figure 14, across all our experiments, the number of images per cluster ranges from 2 to 76, with 75% of our clusters including at most 13 images (third quartile in Figure 14). Based on such numbers, we can conclude that the effort required to inspect a cluster is limited (i.e., at most 13 images to be visualized for 75% of our clusters); further, we have previously demonstrated through a user study that the inspection of five images per cluster is sufficient for a correct identification of the root cause of a DNN failure [4]. Finally, our root cause analysis toolset [25] includes the generation of animated gifs, one for each cluster, thus enabling the quick visualization of all the images in a cluster. In conclusion, either with animated gifs, or when cluster images are inspected in sequence, we conjecture that the number of clusters' images does not strongly impact cost since clusters are typically small and small subsets of larger clusters are sufficient for a correct identification of failure root causes.

What is important, instead, is the purity of clusters as low purity makes it difficult for the end-user to determine commonalities among images.




Fig. 14. Box plot capturing the distribution of the size of the clusters generated by the best pipeline (VGG16/Dbscan/UMAP/NoFT)

Nevertheless, to further discuss cost, we measure the number of clusters to be inspected for each pipeline considering the dataset used for RQ1 and RQ2. We count only clusters capturing the injected failure scenarios. A lower number of clusters should indicate lower cost and, since a number of clusters higher than the number of failure scenarios to be discovered implies the presence of redundant clusters, we compute the degree of redundancy as:

$$\text{redundancy ratio} = \frac{\text{number of clusters}}{\text{covered failure scenarios}}$$

Finally, to discuss how well each pipeline improves current practice in industry, we estimate the degree of savings with respect to the such practice, which entails the visual inspection of all images. To do so, we assume that inspecting a single cluster, especially when using animated gifs, is as inexpensive as visualizing one single image. Indeed, though clusters involve several images, through animation, they actually make it easier to quickly identify commonalities rather than guessing root causes from a single image. Figure 15 shows four example clusters where all the images present a commonality (i.e., the root cause of the DNN failure) that is easy to determine when visualizing all the images in a sequence. Therefore, we estimate savings as:

$$\text{savings} = 1 - \frac{\text{number of clusters}}{\text{number of images}}$$

Table 13 shows our results; it reports the number of RCCs generated for each case study DNN and across all of them. Further, it reports the *percentage and number of failure scenarios* covered by each pipeline (used to compute redundancy and providing information about the effectiveness of a pipeline), along with *redundancy ratio* and *savings*. We report only the results for the best pipelines identified when addressing RQ1 to RQ2 because there is no reason to select pipelines that do not achieve high purity and coverage.

The number of clusters generated by the selected pipelines ranges between 18 and 284. The pipelines leading to the lowest number of clusters are the ones including K-means: *ResNet50/K-means/UMAP/NoFT* (18), *VGG16/K-means/None/NoFT* (19), and *VGG16/K-means/UMAP/NoFT* (24). Pipelines with DBSCAN and HDBSCAN lead to a much higher number of clusters. To discuss the practical impact of such a high number of clusters, we focus on the *redundancy ratio*, which ranges between 1.12 and 11.8; the redundancy ratio indicates that the pipeline with the highest number of

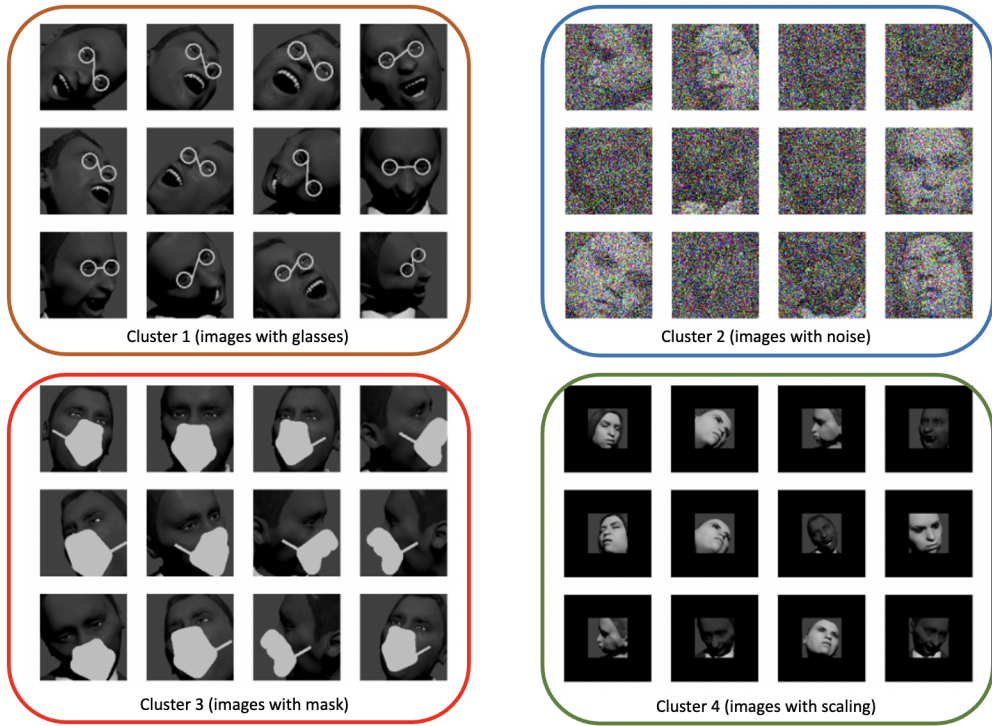


Fig. 15. Examples of clusters generated by pipeline 26 for the HPD case study subject.

Table 13. The number of redundant clusters generated by the best pipelines for each case study subject and across all of them. The last columns represent the number and the percentage of failure scenarios covered by the pipelines, the redundancy ratio, and the savings.

Pipelines	Number of generated clusters							CFS	RR	S
	GD	HPD	OC	SVIRO	CPD	SAP	TOTAL			
VGG16/K-means/None/NoFT	3	5	3	3	3	2	19	17 (59%)	1,12	0,99
VGG16/K-means/UMAP/NoFT	4	8	2	4	3	3	24	20 (69%)	1,20	0,99
ResNet50/K-means/UMAP/NoFT	3	4	3	2	2	4	18	15 (52%)	1,20	0,99
VGG16/Dbscan/UMAP/NoFT	26	77	13	8	13	37	174	28 (97%)	6,21	0,91
ResNet50/Dbscan/UMAP/NoFT	42	51	5	10	27	44	179	27 (93%)	6,63	0,91
Inception_V3/Dbscan/UMAP/NoFT	28	60	7	9	15	42	161	29 (100%)	5,55	0,92
Xception/Dbscan/UMAP/NoFT	33	30	9	2	9	14	97	25 (86%)	3,88	0,95
ResNet50/HDBSCAN/None/NoFT	14	171	15	7	74	3	284	24 (83%)	11,83	0,86

Legend: CFS: Covered failure scenarios (percentage %); RR: Redundancy Ratio; S: Savings.

clusters (i.e., *ResNet50/HDBSCAN/None/NoFT*), on average, presents 11 redundant clusters for each identified failure scenario. Given that, in the presence of pure clusters, understanding the scenario captured by one pipeline is quick with animated gifs, we consider that inspecting 11 redundant clusters per fault has a limited impact on cost. Finally, if we focus on *savings*, we can observe that respect to current practice, all the pipelines except (*ResNet50/HDBSCAN/Only/NoFT*) lead to savings above 90%, thus showing that their adoption is highly beneficial.

Although the pipelines including K-means lead to the lowest cost, their coverage is particularly low for infrequent scenarios (see Table 10, with coverage below 35% for the range [0-5], and below 60% for the range [5-10]), which is bound to be a common situation in practice. Since pipelines leading to a small number of clusters can be highly ineffective in realistic safety-critical contexts (i.e., when some failure scenarios are infrequent), assuming that redundant clusters are easy to manage, we conclude that the best choice are the pipelines that maximize purity and coverage, as discussed above (i.e., Pipeline 26, *VGG16/DBSCAN/UMAP/NoFT*). A possible tradeoff is Pipeline 80 (*Xception/DBSCAN/UMAP/NoFT*), which is among the best performing for RQ3 (e.g., coverage above 40% for the range [0-5], and above 70% for the range [5-10]) and leads to 3.6 redundant clusters only, on average.

4.9 Threats to validity

We discuss internal, conclusion, construct, and external validity according to conventional practices [98].

4.9.1 Internal validity. Since 72 of our 99 pipelines use a Transfer Learning pre-trained model to extract features, a possible internal threat is that this model can negatively affect our results if inadequate. Indeed, clustering relies on the similarity computed on the extracted features. However, since every pre-trained model is integrated into at least one of the best pipelines identified in our experiments (see Table 12), we conclude that they are suitable. Also, to mitigate the risk that our purity metric might not reflect what is perceived by the end-user as a pure cluster, we relied on the same purity metric adopted in our previous work [4] to conduct an empirical study with human subjects, which demonstrated that end-users can understand the root causes of failures by looking at a small random subset of images in each cluster. Further, we visually inspected a random subset of our clusters to check their consistency. Such consistency suggests that the features extracted by the models contain enough information to cluster the images based on their similarity.

Another potential threat might be that the dataset (with the injected faults) was created with the proposed approach in mind. Therefore, there might be a risk of bias. To mitigate this risk, all the methods used in our pipelines (feature extraction methods, clustering algorithms, dimensionality reduction techniques) are independent of the data. These methods do not require any a priori knowledge on the data. We also publish our data to further mitigate this risk. All the experiments can be reproduced with any injected faulty scenario.

4.9.2 External validity. To alleviate the threats related to the choice of the case study DNNs, we use six well-studied datasets with diverse complexity. Four out of six subject DNNs implement tasks motivated by IEE business needs. These DNNs address problems that are quite common in the automotive industry. The other two DNNs are also related to the automotive industry and were used in many Kaggle challenges [65, 91].

Although our pipelines were only tested on case study DNNs related to the automotive industry, we believe they will perform well with other datasets. This is because the models used for the feature extraction were pre-trained on ImageNet, which means that the model can capture features related to 1,000 classes, including humans, animals, and objects. As for AE, it can learn the aspects of any dataset during training and provide high-quality clusters. Finally, for HUDD and LRP, the extraction of heatmap-based features is performed on well-known layer types that are part of any DNN model, regardless of the task at hand (i.e., they can be extended to DNNs that were not studied in this work).

4.9.3 Construct validity. The construct considered in our work is effectiveness. We measure the effectiveness through complementary indicators as follows:

For RQ1, we evaluate the effectiveness of our pipelines by computing the purity of the generated clusters. The purity of a cluster is measured as the maximum proportion of images representing one faulty scenario in this cluster.

For RQ2, we evaluate the effectiveness of our pipelines based on the coverage of the injected faulty scenarios by the root cause clusters. A faulty scenario is covered by a cluster if at least 90% of the images in this cluster represent such faulty scenario.

Finally, for RQ3, we consider both the purity and the coverage to measure the robustness of the top-performing pipelines to rare faulty scenarios.

4.9.4 Conclusion validity. Conclusion validity addresses threats that impact the ability to conclude appropriately. To mitigate such threats and to avoid violating parametric assumptions in our statistical analysis, we rely on a non-parametric test and effect size measure (i.e., Mann Whitney U-test and the Vargha and Delaney's \hat{A}_{12} statistics, respectively) to assess the statistical significance of differences in our results. Additionally, we applied the Fisher's exact test when comparing coverage results related to different distributions of faulty scenarios (i.e., RQ3), which is commonly used in similar contexts. All results were reported based on both purity and coverage parameters, and six datasets were analyzed during our experiments.

4.10 Data Availability

All our implementations, the failure-inducing sets, the generated root-cause clusters and the data generated to address our research questions are available online [5].

5 RELATED WORK

Our paper is related to the literature on DNN debugging and applications of transfer learning to perform root cause analysis [64, 87].

Heatmap-based approaches [15, 60, 69, 77, 81, 100, 102] explain the DNN's prediction of an image by highlighting which region of that image influenced the most the DNN output. For example, Grad-CAM generates a heatmap from the gradient flowing into the last layer. The heatmap is then superposed on the original image to highlight the regions of the image that activated the DNN and influenced the decision [77]. The main limitation of these approaches is that they require the inspection of all the heatmaps generated for the images processed by the DNN (e.g., error-inducing images) and, different from our pipelines, do not provide engineers with guidance for their inspection (i.e., one cluster for each failure root cause). SHAP (SHapley Additive exPlanation) [54] generates explanations by calculating the contribution of each feature to predictions, thus explaining what features are the most important for each prediction. In the case of an image CNN, SHAP considers a group of pixels as a feature and calculates their contribution to the decision made by the DNN. Like heatmap-based approaches, SHAP does not provide guidance for the investigation of multiple failure-inducing images.

DeepJanus [74] helps identify misbehaviors in a Deep Learning system by finding a set of pairs of inputs that are similar to each other and that trigger different behaviors of the Deep Learning system. This set of pairs represents the border between the input regions where the DNN behaves as expected or fails. Different from our work, DeepJanus characterizes the behaviour of a DNN that can be tested with a simulator but cannot provide explanations for failures observed with real-world images.

Some DNN testing approaches explain the input regions where DNN errors are observed by relying on decision trees constructed using the simulator parameters used to generate test input images [2, 37]. Although decision trees are an effective mean to provide explanations for failures detected during simulator-based testing, they cannot be applied to provide explanations for failures

observed with real-world images. To overcome such a limitation, we have recently developed SEDE [26], a technique that applies HUDD to failure-inducing real-world images to generate root cause clusters and then relies on evolutionary algorithms to drive the generation, for each RCC, of similar images using simulators. The simulator parameter values used to generate such images are then fed into PART [29], a tree-based rule learning algorithm to characterize each RCCs in terms of simulator parameters (i.e., it generates expressions that constrain simulator parameters). The work in this paper is complementary to SEDE since the latter can be applied to clusters generated with the best pipeline (i.e., Pipeline 26).

Pan et al. [64] combine Transfer Learning with clustering to find root causes of hardware failures. The proposed method uses different clustering algorithms (K-means [56], decision tree clustering [52], hierarchical clustering [49]) on hardware test data to cluster failures likely due to the same causes. Different from their work, we aim to explain failures in DNNs that process images (i.e., our feature space is much larger). Ter Burg et al. [87] explain DNNs based on a transfer learning model that has been fine-tuned to detect geometric shapes connecting face landmarks. Such shapes are treated as features and the contribution of each feature is computed by relying on SHAP. The output should help end-users determine what influenced the DNN output. Unfortunately, similar to heatmap-based approaches, this approach does not support the explanation of multiple failures but require engineers to process them one by one.

To conclude, our previous works (i.e., HUDD [24] and SAFE [4]) have been the first to apply clustering algorithms to white-box and black-box feature extraction approaches to explain failure causes in DNN-based systems. This study is the first to systematically assess and compare the performance of alternative white-box and black-box feature extraction approaches, dimensionality reduction techniques, and clustering algorithms using a wide variety of practical, realistic failure scenarios.

6 CONCLUSION

In this paper, we presented a large-scale empirical evaluation of 99 different pipelines for root cause analysis of DNN failures. Our pipelines receive as input a set of images leading to DNN failures and generate as output cluster of images sharing similar characteristics. As demonstrated by our previous work, by visualizing the images in each cluster, an engineer can notice commonalities across the images in each cluster; such commonalities represent the root causes of failures, help characterize failure scenarios and, thus, support engineers in improving the system (e.g., by selecting additional similar images to retrain the DNN or by introducing countermeasures in the system).

We considered 99 pipelines resulting from the combination of five methods for feature extraction, two techniques for dimensionality reduction and three clustering algorithms. Our methods for feature extraction include white-box (i.e., heatmap generation techniques) and black-box approaches (i.e., fine-tuned and non-finetuned transfer learning models). Additionally, we rely on PCA and UMAP for dimensionality reduction and K-means, DBSCAN, and HDBSCAN for clustering.

We evaluated our pipelines in terms of clusters' purity and coverage of failures based on a controlled set of failure scenarios artificially injected into our datasets and widely varying in terms of frequency, thus analyzing the impact of rare scenarios on our best pipelines. Further, we assess the performance of our clustering pipelines in identifying failure scenarios not artificially injected but naturally present in our datasets. Based on six case study subjects in the automotive domain, our results show that the best results are obtained with a pipeline relying on VGG-16 as transfer learning model, not using fine tuning, leveraging UMAP as a dimensionality reduction technique, and using DBSCAN as clustering algorithm. When the failure scenarios are equally distributed, the best pipeline achieved a purity of 94.3% (i.e., almost all the images in RCCs present the same failure scenario) and a coverage of 96.7%. The same pipeline also performs well with rare failure

scenarios; indeed, when images belonging to failure scenarios represent between 5 and 10% of the total number of images, it still can cover 90% of the failure scenarios with a cluster purity above 70%. Finally, we observed that the pipeline performing the best with injected failure scenarios also leads to the best results with failure scenarios already present in the datasets; specifically, it achieves 100% coverage and an average purity of 87% across clusters.

ACKNOWLEDGMENTS

This project has received funding from IEE Luxembourg, Luxembourg's National Research Fund (FNR) under grant BRIDGES2020/IS/14711346/FUNTASY, and NSERC of Canada under the Discovery and CRC programs. Authors would like to thank Thomas Stifter from IEE for his valuable support. The experiments presented in this paper were carried out using the HPC facilities of the University of Luxembourg (see <http://hpc.uni.lu>).

REFERENCES

- [1] Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/> Software available from tensorflow.org.
- [2] Raja Ben Abdesslem, Shiva Nejati, Lionel C Briand, and Thomas Stifter. 2018. Testing vision-based control systems using learnable evolutionary algorithms. In *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)*. IEEE, 1016–1026.
- [3] Mohammed Oualid Attaoui, Nassima Dif, Hanene Azzag, and Mustapha Lebbah. 2022. Regions of interest selection in histopathological images using subspace and multi-objective stream clustering. *The Visual Computer* (2022), 1–19.
- [4] Mohammed Oualid Attaoui, Hazem Fahmy, Fabrizio Pastore, and Lionel Briand. 2022. Black-Box Safety Analysis and Retraining of DNNs Based on Feature Extraction and Clustering. *ACM Trans. Softw. Eng. Methodol.* (jul 2022). <https://doi.org/10.1145/3550271>
- [5] Mohammed Oualid Attaoui, Hazem Fahmy, Fabrizio Pastore, and Lionel Briand. 2023. DNN Explanation for Safety Analysis: an Empirical Evaluation of Clustering-based Approaches - Replicability package. https://figshare.com/projects/DNN_Explanation_for_Safety_Analysis_an_Empirical_Evaluation_of_Clustering-based_Approaches/157973
- [6] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. 2016. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316* (2016).
- [7] Diogo V Carvalho, Eduardo M Pereira, and Jaime S Cardoso. 2019. Machine learning interpretability: A survey on methods and metrics. *Electronics* 8, 8 (2019), 832.
- [8] Sully Chen. 2016. *Autopilot-Tensorflow*. <https://github.com/SullyChen/Autopilot-TensorFlow>
- [9] Xinyang Chen, Sinan Wang, Bo Fu, Mingsheng Long, and Jianmin Wang. 2019. Catastrophic forgetting meets negative transfer: Batch spectral shrinkage for safe transfer learning. *Advances in Neural Information Processing Systems* 32 (2019).
- [10] Yukang Chen, Yanwei Li, Tao Kong, Lu Qi, Ruihang Chu, Lei Li, and Jiaya Jia. 2021. Scale-aware automatic augmentation for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9563–9572.
- [11] François Chollet. 2017. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1251–1258.
- [12] Francois Chollet et al. 2015. *Keras*. <https://github.com/fchollet/keras>
- [13] Alex Clark. 2015. Pillow (PIL Fork) Documentation. <https://buildmedia.readthedocs.org/media/pdf/pillow/latest/pillow.pdf>
- [14] Steve Dias Da Cruz, Oliver Wasenmüller, Hans-Peter Beise, Thomas Stifter, and Didier Stricker. 2020. SVIRO: Synthetic Vehicle Interior Rear Seat Occupancy Dataset and Benchmark. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 962–971. <https://doi.org/10.1109/WACV45572.2020.9093315>
- [15] Piotr Dabkowski and Yarín Gal. 2017. Real Time Image Saliency for Black Box Classifiers. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*. Curran Associates Inc., Red Hook, NY,

USA, 6970–6979.

- [16] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [17] Steve Dias Da Cruz, Bertram Taetz, Thomas Stifter, and Didier Stricker. 2022. Autoencoder Attractors for Uncertainty Estimation. In *Proceedings of the IEEE International Conference on Pattern Recognition (ICPR)*.
- [18] Nassima Dif, Mohammed Oualid Attaoui, Zakaria Elberrichi, Mustapha Lebbah, and Hanene Azzag. 2021. Transfer learning from synthetic labels for histopathological images classification. *Applied Intelligence* (2021), 1–20.
- [19] Raghav Bali Dipanjan Sarkar. 2021. *Transfer Learning in Action*. Manning Publications. 69–76 pages. <https://livebook.manning.com/book/transfer-learning-in-action/chapter-1/v-1/69> (visited 2022-02-19).
- [20] Shuyang Du, Haoli Guo, and Andrew Simpson. 2019. Self-driving car steering angle prediction based on image recognition. *arXiv preprint arXiv:1912.05440* (2019).
- [21] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. 2019. CenterNet: Keypoint Triplets for Object Detection. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE Computer Society, 6568–6577.
- [22] Ekin EKİNCİ, Hidayet TAKCI, and Sultan ALAGÖZ. 2022. Poet Classification Using ANN and DNN. *Electronic Letters on Science and Engineering* 18, 1 (2022), 10–20.
- [23] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise.. In *Kdd*, Vol. 96. 226–231.
- [24] Hazem Fahmy, Fabrizio Pastore, Mojtaba Bagherzadeh, and Lionel Briand. 2021. Supporting Deep Neural Network Safety Analysis and Retraining Through Heatmap-Based Unsupervised Learning. *IEEE Transactions on Reliability* (2021), 1–17. <https://doi.org/10.1109/TR.2021.3074750>
- [25] Hazem Fahmy, Fabrizio Pastore, and Lionel Briand. 2022. HUDD: A tool to debug DNNs for safety analysis. In *2022 IEEE/ACM 44th International Conference on Software Engineering: Companion Proceedings (ICSE-Companion)*. 100–104. <https://doi.org/10.1145/3510454.3516858>
- [26] Hazem Fahmy, Fabrizio Pastore, Lionel Briand, and Thomas Stifter. 2022. Simulator-Based Explanation and Debugging of Hazard-Triggering Events in DNN-Based Safety-Critical Systems. *ACM Trans. Softw. Eng. Methodol.* (oct 2022). <https://doi.org/10.1145/3569935>
- [27] Hazem Fahmy, Fabrizio Pastore, Lionel Briand, and Thomas Stifter. 2023. Simulator-based explanation and debugging of hazard-triggering events in DNN-based safety-critical systems. *ACM Transactions on Software Engineering and Methodology* 32, 4 (2023), 1–47.
- [28] Florent Forest, Mustapha Lebbah, Hanene Azzag, and Jérôme Lacaille. 2021. Deep embedded self-organizing maps for joint representation learning and topology-preserving clustering. *Neural Computing and Applications* 33, 24 (2021), 17439–17469.
- [29] Eibe Frank and Ian H. Witten. 1998. Generating Accurate Rule Sets Without Global Optimization. In *Fifteenth International Conference on Machine Learning*, J. Shavlik (Ed.). Morgan Kaufmann, 144–151.
- [30] Keinosuke Fukunaga and Larry Hostetler. 1975. The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Transactions on information theory* 21, 1 (1975), 32–40.
- [31] Rafael Garcia, Alexandru C. Telea, Bruno Castro da Silva, Jim Torresen, and Joao Luiz Dihl Comba. 2018. A task-and-technique centered survey on visual analytics for deep learning model engineering. *Computers and Graphics* 77 (2018), 30 – 49. <https://doi.org/10.1016/j.cag.2018.09.018>
- [32] Leilani H Gilpin, David Bau, Ben Z Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. 2018. Explaining explanations: An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*. IEEE, 80–89.
- [33] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- [34] Alexander N Gorban and Andrei Y Zinovyev. 2010. Principal graphs and manifolds. In *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*. IGI Global, 28–59.
- [35] David Gunning and David Aha. 2019. DARPA’s explainable artificial intelligence (XAI) program. *AI magazine* 40, 2 (2019), 44–58.
- [36] Jiawei Han, Micheline Kamber, and Jian Pei. 2011. *Data Mining: Concepts and Techniques* (3rd ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- [37] Fitash Ul Haq, Donghwan Shin, Lionel C. Briand, Thomas Stifter, and Jun Wang. 2021. Automatic Test Suite Generation for Key-Points Detection DNNs Using Many-Objective Search (Experience Paper). In *Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA 2021)*. Association for Computing Machinery, New York, NY, USA, 91–102. <https://doi.org/10.1145/3460319.3464802>
- [38] Fitash Ul Haq, Donghwan Shin, Shiva Nejati, and Lionel Briand. 2021. Can Offline Testing of Deep Neural Networks Replace Their Online Testing? A Case Study of Automated Driving Systems. *Empirical Softw. Engg.* 26, 5 (sep 2021),

30. <https://doi.org/10.1007/s10664-021-09982-4>
- [39] Samuel W Hasinoff, Dillon Sharlet, Ryan Geiss, Andrew Adams, Jonathan T Barron, Florian Kainz, Jiawen Chen, and Marc Levoy. 2016. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics (ToG)* 35, 6 (2016), 1–12.
- [40] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [41] Torsten Hothorn, Kurt Hornik, and Achim Zeileis. 2015. ctree: Conditional inference trees. *The comprehensive R archive network* 8 (2015).
- [42] Xinyu Huang, Peng Wang, Xinjing Cheng, Dingfu Zhou, Qichuan Geng, and Ruigang Yang. 2020. The ApolloScape Open Dataset for Autonomous Driving and Its Application. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42, 10 (oct 2020), 2702–2719. <https://doi.org/10.1109/tpami.2019.2926463>
- [43] Nargiz Humbatova, Gunel Jahangirova, and Paolo Tonella. 2021. Deepcrime: mutation testing of deep learning systems based on real faults. In *Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 67–78.
- [44] IEE. 2020. Vitasense radar-based child detection. <https://iee-sensing.com/automotive/safety-and-comfort/vitasense/>.
- [45] IEE. 2023. IEE Sensing solutions. www.iee.lu.
- [46] International Organization for Standardization. 2020. ISO, ISO-24765-2017, Systems and software engineering - Vocabulary.
- [47] International Organization for Standardization. 2020. ISO, ISO26262-1:2018, Road vehicles: Functional safety.
- [48] Barbara Kitchenham, Lech Madeyski, David Budgen, Jacky Keung, Pearl Brereton, Stuart Charters, Shirley Gibbs, and Amnart Pohthong. 2017. Robust Statistical Methods for Empirical Software Engineering. *Empirical Software Engineering* 22, 2 (2017), 579–630. <https://doi.org/10.1007/s10664-016-9437-5>
- [49] Hans-Friedrich Köhn and Lawrence J Hubert. 2014. Hierarchical cluster analysis. *Wiley StatsRef: Statistics Reference Online* (2014), 1–13.
- [50] Hans-Peter Kriegel, Peer Kröger, Jörg Sander, and Arthur Zimek. 2011. Density-based clustering. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 1, 3 (2011), 231–240.
- [51] Zhong Li, Minxue Pan, Tian Zhang, and Xuandong Li. 2021. Testing DNN-based Autonomous Driving Systems under Critical Environmental Conditions. In *International Conference on Machine Learning*. PMLR, 6471–6482.
- [52] Bing Liu, Yiyuan Xia, and Philip S Yu. 2000. Clustering through decision tree construction. In *Proceedings of the ninth international conference on Information and knowledge management*. 20–29.
- [53] J. Liu and Wang Y. 2020. Noise Reduction Algorithms for Digital Image Processing. *Journal of Physics: Conference Series* 1464(1) (2020).
- [54] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30 (2017).
- [55] J. MacQueen. 1967. Some methods for classification and analysis of multivariate observations. *Proceedings of BSMSPP* (1967), 281–297.
- [56] J. MacQueen. 1967. Some Methods for Classification and Analysis of Multivariate Observations. In *Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability - Vol. 1*, L. M. Le Cam and J. Neyman (Eds.). University of California Press, Berkeley, CA, USA, 281–297.
- [57] Leland McInnes, John Healy, and Steve Astels. 2017. hdbscan: Hierarchical density based clustering. *Journal of Open Source Software* 2, 11 (2017), 205.
- [58] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. 2018. UMAP: Uniform Manifold Approximation and Projection. *The Journal of Open Source Software* 3, 29 (2018), 861.
- [59] Gregoire Montavon. 2019. WIFS 2017 Tutorial on Methods for Understanding DNNs and their Predictions. <http://heatmapping.org/wifs2017/>
- [60] Grégoire Montavon, Alexander Binder, Sebastian Lapuschkin, Wojciech Samek, and Klaus Robert Müller. 2019. *Layer-Wise Relevance Propagation: An Overview*. Springer International Publishing, Cham, 193–209. https://doi.org/10.1007/978-3-030-28954-6_10
- [61] Gregoire Montavon, Sebastian Lapuschkin, Alexander Binder, Wojciech Samek, and Klaus-Robert Muller. 2017. Explaining nonlinear classification decisions with deep Taylor decomposition. *Pattern Recognition* 65 (2017), 211 – 222. <https://doi.org/10.1016/j.patcog.2016.11.008>
- [62] Fionn Murtagh and Pierre Legendre. 2014. Ward’s Hierarchical Agglomerative Clustering Method: Which Algorithms Implement Ward’s Criterion? *Journal of Classification* 31, 3 (01 Oct 2014), 274–295. <https://doi.org/10.1007/s00357-014-9161-z>
- [63] A Kousar Nikhath, MD Abdul Rab, N Venkata Bharadwaja, L Goutham Reddy, K Saicharan, C Venkat Manas Reddy, et al. 2022. An Intelligent College Enquiry Bot using NLP and Deep Learning based techniques. In *2022 International Conference for Advancement in Technology (ICONAT)*. IEEE, 1–6.

- [64] Renjian Pan, Xin Li, and Krishnendu Chakrabarty. 2021. Unsupervised Root-Cause Analysis with Transfer Learning for Integrated Systems. In *2021 IEEE 39th VLSI Test Symposium (VTS)*. IEEE, 1–6.
- [65] ADVAY PATIL. 2022. Car Object Detection. <https://www.kaggle.com/code/advaypatil/car-object-detection>
- [66] Karl Pearson. 1901. LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science* 2, 11 (1901), 559–572.
- [67] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [68] Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. 2017. Deepxplore: Automated whitebox testing of deep learning systems. In *proceedings of the 26th Symposium on Operating Systems Principles*. 1–18.
- [69] Vitali Petsiuk, Abir Das, and Kate Saenko. 2018. RISE: Randomized Input Sampling for Explanation of Black-box Models. In *Proceedings of the British Machine Vision Conference (BMVC)*.
- [70] Autumn pre-trained model. 2017. *Udacity Self-Driving Car Challenge 2*. <https://github.com/udacity/self-driving-car/tree/master/steering-models/community-models/autumn>
- [71] PyTorch. 2020. PyTorch DNN framework. <https://pytorch.org>
- [72] Sizhe Rao, Minghui Wang, Cuixia Tian, Xiangqiao Ao, et al. 2021. A Hierarchical Tree-Based Syslog Clustering Scheme for Network Diagnosis. In *2021 17th International Conference on Network and Service Management (CNSM)*. IEEE, 28–34.
- [73] Satti RG Reddy, GP Varma, and Rajya Lakshmi Davuluri. 2022. Deep Neural Network (DNN) Mechanism for Identification of Diseased and Healthy Plant Leaf Images Using Computer Vision. *Annals of Data Science* (2022), 1–30.
- [74] Vincenzo Riccio and Paolo Tonella. 2020. Model-based exploration of the frontier of behaviours for deep learning system testing. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 876–888.
- [75] Peter J Rousseeuw. 1987. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics* 20 (1987), 53–65.
- [76] V. Satopaa, J. Albrecht, D. Irwin, and B. Raghavan. 2011. Finding a “Kneedle” in a Haystack: Detecting Knee Points in System Behavior. In *2011 31st International Conference on Distributed Computing Systems Workshops*. 166–171.
- [77] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. 2017. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In *2017 IEEE International Conference on Computer Vision (ICCV)*. 618–626. <https://doi.org/10.1109/ICCV.2017.74>
- [78] Jonathon Shlens. 2014. A tutorial on principal component analysis. *arXiv preprint arXiv:1404.1100* (2014).
- [79] Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*. <https://arxiv.org/abs/1409.1556>
- [80] Xibin Song, Peng Wang, Dingfu Zhou, Rui Zhu, Chenye Guan, Yuchao Dai, Hao Su, Hongdong Li, and Ruigang Yang. 2019. Apollocar3d: A large 3d car instance understanding benchmark for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5452–5462.
- [81] J.T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. 2015. Striving for Simplicity: The All Convolutional Net. In *ICLR (workshop track)*.
- [82] Stanford Vision Lab. 2015. ImageNet, image database organized according to the WordNet hierarchy. <https://www.image-net.org>. Accessed: 2021-12-07.
- [83] Helmut Strasser and Christian Weber. 1999. *On the Asymptotic Theory of Permutation Statistics* (january 1999 ed.). WorkingPaper 27. SFB Adaptive Information Systems and Modelling in Economics and Management Science, WU Vienna University of Economics and Business.
- [84] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1–9.
- [85] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2818–2826.
- [86] Muhammed Talo. 2019. Automated classification of histopathology images using transfer learning. *Artificial Intelligence in Medicine* 101 (2019), 101743.
- [87] Kaya ter Burg and Heysem Kaya. 2022. Comparing Approaches for Explaining DNN-Based Facial Expression Classifications. *Algorithms* 15, 10 (2022), 367.
- [88] Robert L Thorndike. 1953. Who belongs in the family? *Psychometrika* 18, 4 (1953), 267–276. <https://doi.org/10.1007/BF02289263>
- [89] Yuchi Tian, Kexin Pei, Suman Jana, and Baishakhi Ray. 2018. DeepTest: Automated Testing of Deep-Neural-Network-Driven Autonomous Cars. In *Proceedings of the 40th International Conference on Software Engineering (ICSE ’18)*.

- Association for Computing Machinery, New York, NY, USA, 303–314. <https://doi.org/10.1145/3180155.3180220>
- [90] Udacity. 2017. *Udacity Self-driving Car Challenge 2*. Accessed: 2022-08-16.
 - [91] Peking University/Baidu. 2020. *Peking University/Baidu - Autonomous Driving*. Accessed: 2022-08-16.
 - [92] Graham JG Upton. 1992. Fisher’s exact test. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 155, 3 (1992), 395–402.
 - [93] Stéfan van der Walt, Johannes L. Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D. Warner, Neil Yager, Emmanuelle Gouillart, and Tony Yu. 2014. scikit-image: image processing in Python. *PeerJ* 2 (jun 2014), e453. <https://doi.org/10.7717/peerj.453>
 - [94] Qi Wang, Jian Chen, Jianqiang Deng, and Xinfang Zhang. 2021. 3D-CenterNet: 3D object detection network for point clouds with center estimation priority. *Pattern Recognition* 115 (2021), 107884.
 - [95] Jr. Ward, J. H. 1963. Hierarchical grouping to optimize an objective function. *American Statistical Association Journal* 58 (1963), 236–244.
 - [96] Eric W. Weisstein. 2022. Fisher’s Exact Test. <https://mathworld.wolfram.com/FishersExactTest.html>. Accessed: 2022-11-15.
 - [97] Katharina Weitz, Teena Hassan, Ute Schmid, and Jens Garbas. 2018. Towards explaining deep learning networks to distinguish facial expressions of pain and emotions. In *Forum Bildverarbeitung*. 197–208.
 - [98] Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, Björn Regnell, and Anders Wesslén. 2012. *Experimentation in software engineering*. Vol. 9783642290. 1–236 pages. <https://doi.org/10.1007/978-3-642-29044-2>
 - [99] S Paul Wright. 1992. Adjusted p-values for simultaneous inference. *Biometrics* (1992), 1005–1013.
 - [100] Matthew D. Zeiler and Rob Fergus. 2014. Visualizing and Understanding Convolutional Networks. In *Computer Vision – ECCV 2014*, David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.). Springer International Publishing, Cham, 818–833.
 - [101] Yang Zheng, Zhenye Feng, Zheng Hu, and Ke Pei. 2021. MindFI: A Fault Injection Tool for Reliability Assessment of MindSpore Applications. In *2021 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW)*. IEEE, 235–238.
 - [102] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. 2016. Learning Deep Features for Discriminative Localization. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2921–2929. <https://doi.org/10.1109/CVPR.2016.319>

A ADDITIONAL MATERIAL FOR RQ1

Table 14. Comparing the clusters generated by the different pipelines based on the average of the purity across root cause clusters. The last column represents the average of averages.

Pipelines					Case Study Subjects						
#	FE	FT	DR	CA	GD	OC	HPD	SVIRO	SAP	CPD	Avg.
1	HUDD	None	None	K-Means	51.3%	36.6%	40.7%	62.4%	78.9%	39.9%	51.6%
2	HUDD	None	None	DBSCAN	56.2%	53.4%	43.1%	53.0%	80.6%	63.7%	58.3%
3	HUDD	None	None	HDBSCAN	68.5%	61.7%	43.7%	45.4%	51.2%	27.4%	49.6%
4	HUDD	None	PCA	K-Means	49.1%	56.4%	40.8%	74.8%	79.7%	27.4%	54.7%
5	HUDD	None	PCA	DBSCAN	43.4%	54.6%	48.7%	48.3%	80.6%	27.4%	50.5%
6	HUDD	None	PCA	HDBSCAN	68.5%	61.7%	35.5%	45.4%	74.1%	26.9%	52.0%
7	HUDD	None	UMAP	K-Means	56.7%	47.6%	42.3%	54.3%	61.1%	32.2%	49.0%
8	HUDD	None	UMAP	DBSCAN	69.6%	58.7%	68.3%	59.3%	68.7%	53.9%	63.1%
9	HUDD	None	UMAP	HDBSCAN	68.5%	61.7%	33.3%	45.4%	74.1%	27.4%	51.7%
10	LRP	None	None	K-Means	42.9%	36.6%	56.5%	33.8%	82.8%	73.6%	54.4%
11	LRP	None	None	DBSCAN	39.5%	28.7%	69.8%	50.7%	96.5%	35.4%	53.4%
12	LRP	None	None	HDBSCAN	69.0%	58.3%	56.2%	47.8%	42.4%	27.3%	50.2%
13	LRP	None	PCA	K-Means	54.2%	56.4%	54.9%	35.7%	82.9%	73.6%	59.6%
14	LRP	None	PCA	DBSCAN	47.2%	20.6%	71.8%	48.9%	79.7%	35.4%	50.6%
15	LRP	None	PCA	HDBSCAN	52.5%	58.3%	25.5%	47.8%	46.8%	26.2%	42.8%
16	LRP	None	UMAP	K-Means	55.9%	47.6%	54.7%	31.8%	67.0%	32.2%	48.2%
17	LRP	None	UMAP	DBSCAN	67.0%	62.8%	68.9%	49.7%	80.3%	33.5%	60.4%
18	LRP	None	UMAP	HDBSCAN	69.0%	58.3%	56.2%	47.8%	51.6%	26.4%	51.5%
19	VGG-16	None	None	K-Means	91.7%	92.1%	95.5%	82.5%	97.3%	99.7%	93.2%
20	VGG-16	None	None	DBSCAN	87.0%	85.9%	96.7%	57.3%	98.0%	100.0%	87.5%
21	VGG-16	None	None	HDBSCAN	52.6%	99.0%	30.7%	73.4%	77.8%	54.5%	64.7%
22	VGG-16	None	PCA	K-Means	90.5%	87.6%	58.3%	87.7%	94.2%	92.2%	85.1%
23	VGG-16	None	PCA	DBSCAN	90.7%	94.9%	81.0%	71.6%	96.0%	91.8%	87.7%
24	VGG-16	None	PCA	HDBSCAN	45.6%	95.1%	56.2%	93.5%	100.0%	76.0%	77.7%
25	VGG-16	None	UMAP	K-Means	97.6%	84.4%	93.7%	82.4%	90.3%	97.8%	91.0%
26	VGG-16	None	UMAP	DBSCAN	99.0%	93.0%	99.6%	79.7%	98.1%	96.6%	94.3%
27	VGG-16	None	UMAP	HDBSCAN	78.0%	96.7%	56.2%	88.9%	44.1%	79.9%	74.0%
28	VGG-16	FT	None	K-Means	26.2%	33.9%	15.8%	24.3%	27.9%	25.5%	25.6%
29	VGG-16	FT	None	DBSCAN	26.4%	38.5%	18.2%	32.8%	29.6%	25.2%	28.4%
30	VGG-16	FT	None	HDBSCAN	23.4%	51.1%	14.3%	48.1%	25.7%	53.2%	36.0%
31	VGG-16	FT	PCA	K-Means	25.4%	37.7%	16.7%	24.5%	26.8%	25.8%	26.1%
32	VGG-16	FT	PCA	DBSCAN	29.2%	51.4%	23.4%	46.6%	32.3%	29.0%	35.3%
33	VGG-16	FT	PCA	HDBSCAN	22.7%	43.7%	13.8%	41.5%	25.3%	23.3%	28.4%
34	VGG-16	FT	UMAP	K-Means	26.3%	33.6%	18.0%	25.3%	26.2%	26.2%	25.9%
35	VGG-16	FT	UMAP	DBSCAN	45.4%	42.8%	27.0%	39.1%	43.8%	44.0%	40.4%
36	VGG-16	FT	UMAP	HDBSCAN	23.5%	40.4%	14.1%	36.6%	22.0%	24.2%	26.8%
37	ResNet-50	None	None	K-Means	84.2%	84.6%	74.0%	61.2%	86.0%	83.7%	78.9%
38	ResNet-50	None	None	DBSCAN	63.5%	84.6%	87.5%	72.6%	75.5%	72.0%	76.0%
39	ResNet-50	None	None	HDBSCAN	96.4%	100.0%	100.0%	78.8%	87.5%	100.0%	93.8%
40	ResNet-50	None	PCA	K-Means	67.4%	79.6%	61.3%	53.4%	85.8%	75.6%	70.5%
41	ResNet-50	None	PCA	DBSCAN	79.7%	72.9%	51.1%	45.0%	89.8%	80.3%	69.8%
42	ResNet-50	None	PCA	HDBSCAN	40.8%	79.6%	56.2%	32.0%	42.5%	49.7%	50.2%
43	ResNet-50	None	UMAP	K-Means	99.4%	93.0%	82.3%	79.6%	99.6%	97.7%	91.9%
44	ResNet-50	None	UMAP	DBSCAN	100.0%	95.8%	95.8%	79.0%	99.7%	99.3%	94.9%
45	ResNet-50	None	UMAP	HDBSCAN	82.6%	87.3%	38.8%	60.0%	30.5%	69.4%	61.4%
46	ResNet-50	FT	None	K-Means	26.7%	37.4%	19.3%	30.0%	26.2%	25.6%	27.5%
47	ResNet-50	FT	None	DBSCAN	47.2%	40.9%	32.1%	33.7%	35.2%	39.4%	38.1%

Table 14. Comparing the clusters generated by the different pipelines based on the average of the purity across root cause clusters. The last column represents the average of averages.

Pipelines					Case Study Subjects						
#	FE	FT	DR	CA	GD	OC	HPD	SVIRO	SAP	CPD	Avg.
48	ResNet-50	FT	None	HDBSCAN	55.0%	46.7%	15.4%	45.2%	26.4%	24.9%	35.6%
49	ResNet-50	FT	PCA	K-Means	29.5%	37.1%	17.8%	39.5%	26.6%	26.2%	29.5%
50	ResNet-50	FT	PCA	DBSCAN	40.1%	45.6%	23.8%	41.5%	39.4%	39.4%	38.3%
51	ResNet-50	FT	PCA	HDBSCAN	23.7%	50.7%	15.7%	48.2%	24.6%	23.3%	31.1%
52	ResNet-50	FT	UMAP	K-Means	25.5%	34.6%	17.3%	25.7%	27.5%	25.6%	26.0%
53	ResNet-50	FT	UMAP	DBSCAN	37.8%	54.8%	35.9%	48.4%	41.7%	50.6%	44.9%
54	ResNet-50	FT	UMAP	HDBSCAN	23.9%	44.5%	15.1%	23.3%	23.7%	24.2%	25.8%
55	Inception-V3	None	None	K-Means	84.6%	86.0%	95.1%	69.2%	91.2%	87.8%	85.7%
56	Inception-V3	None	None	DBSCAN	100.0%	63.2%	80.9%	17.9%	98.4%	60.4%	70.1%
57	Inception-V3	None	None	HDBSCAN	62.6%	96.0%	99.8%	77.9%	62.6%	95.6%	82.4%
58	Inception-V3	None	PCA	K-Means	66.5%	74.5%	80.9%	68.6%	88.9%	75.7%	75.8%
59	Inception-V3	None	PCA	DBSCAN	97.9%	83.8%	86.0%	96.5%	92.0%	82.2%	89.7%
60	Inception-V3	None	PCA	HDBSCAN	92.5%	86.1%	53.4%	70.8%	69.4%	34.3%	67.8%
61	Inception-V3	None	UMAP	K-Means	94.1%	87.4%	95.0%	67.0%	90.1%	71.3%	84.2%
62	Inception-V3	None	UMAP	DBSCAN	93.4%	95.2%	98.1%	76.6%	97.4%	83.1%	90.7%
63	Inception-V3	None	UMAP	HDBSCAN	74.0%	83.3%	55.9%	80.1%	65.8%	70.0%	71.5%
64	Inception-V3	FT	None	K-Means	26.6%	34.4%	15.6%	23.6%	26.7%	25.2%	25.4%
65	Inception-V3	FT	None	DBSCAN	50.0%	38.3%	24.7%	17.0%	51.0%	44.1%	37.5%
66	Inception-V3	FT	None	HDBSCAN	49.9%	48.2%	15.4%	44.4%	53.0%	53.8%	44.1%
67	Inception-V3	FT	PCA	K-Means	24.4%	31.5%	16.7%	25.6%	27.2%	25.2%	25.1%
68	Inception-V3	FT	PCA	DBSCAN	32.1%	50.6%	29.1%	38.5%	35.1%	38.6%	37.3%
69	Inception-V3	FT	PCA	HDBSCAN	46.7%	44.3%	15.1%	45.2%	25.0%	22.9%	33.2%
70	Inception-V3	FT	UMAP	K-Means	26.1%	31.0%	17.1%	25.8%	25.2%	27.2%	25.4%
71	Inception-V3	FT	UMAP	DBSCAN	45.7%	50.4%	23.4%	45.1%	44.4%	52.5%	43.6%
72	Inception-V3	FT	UMAP	HDBSCAN	47.2%	26.2%	15.3%	37.8%	24.1%	25.0%	29.4%
73	Xception	None	None	K-Means	85.2%	90.4%	88.8%	68.4%	95.2%	72.6%	83.4%
74	Xception	None	None	DBSCAN	60.3%	35.2%	82.6%	34.7%	73.6%	57.0%	57.2%
75	Xception	None	None	HDBSCAN	88.9%	91.7%	85.6%	71.5%	100.0%	92.2%	88.3%
76	Xception	None	PCA	K-Means	75.5%	73.3%	67.2%	57.6%	83.6%	44.4%	66.9%
77	Xception	None	PCA	DBSCAN	78.8%	87.7%	83.6%	71.2%	93.6%	35.2%	75.0%
78	Xception	None	PCA	HDBSCAN	75.8%	84.5%	47.5%	75.5%	62.0%	32.7%	63.0%
79	Xception	None	UMAP	K-Means	93.1%	90.2%	89.8%	66.3%	92.6%	66.7%	83.1%
80	Xception	None	UMAP	DBSCAN	94.7%	92.7%	98.7%	62.4%	99.0%	85.4%	88.8%
81	Xception	None	UMAP	HDBSCAN	77.9%	86.3%	36.7%	78.1%	99.4%	62.4%	73.5%
82	Xception	FT	None	K-Means	25.0%	35.6%	17.1%	27.3%	25.4%	26.2%	26.1%
83	Xception	FT	None	DBSCAN	32.3%	34.9%	22.7%	17.8%	20.2%	55.2%	30.5%
84	Xception	FT	None	HDBSCAN	34.7%	46.2%	15.9%	33.4%	25.2%	54.0%	34.9%
85	Xception	FT	PCA	K-Means	27.0%	35.6%	17.0%	26.2%	25.4%	36.8%	28.0%
86	Xception	FT	PCA	DBSCAN	44.3%	31.1%	21.4%	28.2%	35.4%	35.2%	32.6%
87	Xception	FT	PCA	HDBSCAN	46.8%	56.3%	14.8%	37.6%	24.9%	26.6%	34.5%
88	Xception	FT	UMAP	K-Means	26.4%	32.8%	17.1%	28.7%	27.0%	24.6%	26.1%
89	Xception	FT	UMAP	DBSCAN	36.9%	42.8%	26.7%	50.4%	46.2%	45.5%	41.4%
90	Xception	FT	UMAP	HDBSCAN	23.1%	26.2%	14.9%	40.9%	24.7%	26.4%	26.1%
91	AE	None	None	K-Means	40.9%	56.1%	47.0%	41.6%	72.5%	73.0%	55.2%
92	AE	None	None	DBSCAN	35.2%	60.8%	11.5%	16.9%	20.3%	21.1%	27.6%
93	AE	None	None	HDBSCAN	36.9%	67.8%	0.0%	67.5%	0.0%	63.0%	39.2%
94	AE	None	PCA	K-Means	62.5%	55.1%	51.7%	52.9%	60.9%	77.2%	60.0%
95	AE	None	PCA	DBSCAN	20.4%	73.3%	68.2%	63.6%	89.1%	76.7%	65.2%
96	AE	None	PCA	HDBSCAN	73.5%	60.8%	25.7%	68.1%	76.1%	27.8%	55.3%

Table 14. Comparing the clusters generated by the different pipelines based on the average of the purity across root cause clusters. The last column represents the average of averages.

Pipelines					Case Study Subjects						
#	FE	FT	DR	CA	GD	OC	HPD	SVIRO	SAP	CPD	Avg.
97	AE	None	UMAP	K-Means	43.7%	46.2%	36.4%	37.6%	69.2%	66.0%	49.9%
98	AE	None	UMAP	DBSCAN	65.3%	61.8%	59.0%	51.9%	69.1%	70.6%	62.9%
99	AE	None	UMAP	HDBSCAN	28.4%	60.5%	45.4%	47.2%	41.6%	41.7%	44.1%

B ADDITIONAL MATERIAL FOR RQ2

Table 15. Percentage of faulty scenarios covered by the root cause clusters generated for each pipeline. The last column represents the average of averages.

Pipelines					Case Study Subjects						
#	FE	FT	DR	CA	GD	OC	HPD	SVIRO	SAP	CPD	Avg.
1	HUDD	None	None	K-Means	0.0%	0.0%	0.0%	20.0%	25.0%	0.0%	7.5%
2	HUDD	None	None	DBSCAN	25.0%	25.0%	0.0%	80.0%	25.0%	25.0%	30.0%
3	HUDD	None	None	HDBSCAN	100.0%	25.0%	0.0%	0.0%	0.0%	0.0%	20.8%
4	HUDD	None	PCA	K-Means	0.0%	25.0%	0.0%	40.0%	50.0%	0.0%	19.2%
5	HUDD	None	PCA	DBSCAN	0.0%	25.0%	0.0%	20.0%	50.0%	0.0%	15.8%
6	HUDD	None	PCA	HDBSCAN	100.0%	25.0%	0.0%	0.0%	100.0%	0.0%	37.5%
7	HUDD	None	UMAP	K-Means	25.0%	0.0%	0.0%	0.0%	0.0%	0.0%	4.2%
8	HUDD	None	UMAP	DBSCAN	100.0%	50.0%	75.0%	60.0%	75.0%	100.0%	76.7%
9	HUDD	None	UMAP	HDBSCAN	100.0%	25.0%	0.0%	0.0%	100.0%	0.0%	37.5%
10	LRP	None	None	K-Means	0.0%	0.0%	12.5%	0.0%	50.0%	25.0%	14.6%
11	LRP	None	None	DBSCAN	0.0%	0.0%	37.5%	20.0%	50.0%	0.0%	17.9%
12	LRP	None	None	HDBSCAN	100.0%	25.0%	12.5%	20.0%	0.0%	0.0%	26.2%
13	LRP	None	PCA	K-Means	25.0%	25.0%	12.5%	0.0%	50.0%	25.0%	22.9%
14	LRP	None	PCA	DBSCAN	0.0%	0.0%	12.5%	20.0%	50.0%	0.0%	13.8%
15	LRP	None	PCA	HDBSCAN	0.0%	25.0%	0.0%	20.0%	0.0%	0.0%	7.5%
16	LRP	None	UMAP	K-Means	25.0%	0.0%	25.0%	0.0%	50.0%	0.0%	16.7%
17	LRP	None	UMAP	DBSCAN	75.0%	100.0%	75.0%	40.0%	100.0%	0.0%	65.0%
18	LRP	None	UMAP	HDBSCAN	100.0%	25.0%	12.5%	20.0%	0.0%	0.0%	26.2%
19	VGG-16	None	None	K-Means	50.0%	50.0%	87.5%	80.0%	100.0%	100.0%	77.9%
20	VGG-16	None	None	DBSCAN	50.0%	50.0%	75.0%	20.0%	100.0%	100.0%	65.8%
21	VGG-16	None	None	HDBSCAN	0.0%	75.0%	0.0%	60.0%	50.0%	0.0%	30.8%
22	VGG-16	None	PCA	K-Means	50.0%	50.0%	12.5%	60.0%	100.0%	75.0%	57.9%
23	VGG-16	None	PCA	DBSCAN	50.0%	50.0%	37.5%	40.0%	75.0%	75.0%	54.6%
24	VGG-16	None	PCA	HDBSCAN	0.0%	75.0%	12.5%	100.0%	75.0%	50.0%	52.1%
25	VGG-16	None	UMAP	K-Means	75.0%	75.0%	87.5%	80.0%	75.0%	100.0%	82.1%
26	VGG-16	None	UMAP	DBSCAN	100.0%	100.0%	100.0%	80.0%	100.0%	100.0%	96.7%
27	VGG-16	None	UMAP	HDBSCAN	50.0%	75.0%	12.5%	60.0%	0.0%	50.0%	41.2%
28	VGG-16	FT	None	K-Means	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
29	VGG-16	FT	None	DBSCAN	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
30	VGG-16	FT	None	HDBSCAN	0.0%	25.0%	0.0%	20.0%	0.0%	100.0%	24.2%
31	VGG-16	FT	PCA	K-Means	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
32	VGG-16	FT	PCA	DBSCAN	0.0%	25.0%	0.0%	20.0%	0.0%	0.0%	7.5%
33	VGG-16	FT	PCA	HDBSCAN	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
34	VGG-16	FT	UMAP	K-Means	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
35	VGG-16	FT	UMAP	DBSCAN	25.0%	0.0%	0.0%	0.0%	50.0%	25.0%	16.7%
36	VGG-16	FT	UMAP	HDBSCAN	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
37	ResNet-50	None	None	K-Means	50.0%	50.0%	25.0%	40.0%	50.0%	50.0%	44.2%
38	ResNet-50	None	None	DBSCAN	25.0%	50.0%	37.5%	40.0%	50.0%	25.0%	37.9%
39	ResNet-50	None	None	HDBSCAN	50.0%	100.0%	100.0%	60.0%	75.0%	100.0%	80.8%
40	ResNet-50	None	PCA	K-Means	25.0%	50.0%	12.5%	20.0%	50.0%	25.0%	30.4%
41	ResNet-50	None	PCA	DBSCAN	50.0%	50.0%	12.5%	0.0%	50.0%	25.0%	31.2%
42	ResNet-50	None	PCA	HDBSCAN	0.0%	100.0%	12.5%	0.0%	0.0%	0.0%	18.8%
43	ResNet-50	None	UMAP	K-Means	100.0%	75.0%	50.0%	40.0%	100.0%	100.0%	77.5%
44	ResNet-50	None	UMAP	DBSCAN	100.0%	100.0%	100.0%	60.0%	100.0%	100.0%	93.3%
45	ResNet-50	None	UMAP	HDBSCAN	50.0%	75.0%	0.0%	20.0%	0.0%	25.0%	28.3%
46	ResNet-50	FT	None	K-Means	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
47	ResNet-50	FT	None	DBSCAN	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

Table 15. Percentage of faulty scenarios covered by the root cause clusters generated for each pipeline. The last column represents the average of averages.

Pipelines					Case Study Subjects						
#	FE	FT	DR	CA	GD	OC	HPD	SVIRO	SAP	CPD	Avg.
48	ResNet-50	FT	None	HDBSCAN	50.0%	0.0%	0.0%	0.0%	0.0%	0.0%	8.3%
49	ResNet-50	FT	PCA	K-Means	0.0%	0.0%	0.0%	20.0%	0.0%	0.0%	3.3%
50	ResNet-50	FT	PCA	DBSCAN	0.0%	0.0%	0.0%	20.0%	0.0%	0.0%	3.3%
51	ResNet-50	FT	PCA	HDBSCAN	0.0%	0.0%	0.0%	40.0%	0.0%	0.0%	6.7%
52	ResNet-50	FT	UMAP	K-Means	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
53	ResNet-50	FT	UMAP	DBSCAN	0.0%	50.0%	0.0%	40.0%	25.0%	100.0%	35.8%
54	ResNet-50	FT	UMAP	HDBSCAN	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
55	Inception-V3	None	None	K-Means	75.0%	75.0%	87.5%	40.0%	100.0%	50.0%	71.2%
56	Inception-V3	None	None	DBSCAN	75.0%	25.0%	50.0%	0.0%	100.0%	25.0%	45.8%
57	Inception-V3	None	None	HDBSCAN	25.0%	100.0%	87.5%	100.0%	25.0%	100.0%	72.9%
58	Inception-V3	None	PCA	K-Means	50.0%	50.0%	37.5%	40.0%	75.0%	25.0%	46.2%
59	Inception-V3	None	PCA	DBSCAN	50.0%	75.0%	62.5%	40.0%	75.0%	25.0%	54.6%
60	Inception-V3	None	PCA	HDBSCAN	25.0%	75.0%	12.5%	60.0%	25.0%	0.0%	32.9%
61	Inception-V3	None	UMAP	K-Means	75.0%	75.0%	87.5%	40.0%	75.0%	50.0%	67.1%
62	Inception-V3	None	UMAP	DBSCAN	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
63	Inception-V3	None	UMAP	HDBSCAN	25.0%	100.0%	12.5%	100.0%	25.0%	25.0%	47.9%
64	Inception-V3	FT	None	K-Means	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
65	Inception-V3	FT	None	DBSCAN	0.0%	0.0%	0.0%	0.0%	25.0%	0.0%	4.2%
66	Inception-V3	FT	None	HDBSCAN	75.0%	25.0%	0.0%	40.0%	75.0%	100.0%	52.5%
67	Inception-V3	FT	PCA	K-Means	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
68	Inception-V3	FT	PCA	DBSCAN	0.0%	25.0%	0.0%	0.0%	0.0%	0.0%	4.2%
69	Inception-V3	FT	PCA	HDBSCAN	25.0%	0.0%	0.0%	20.0%	0.0%	0.0%	7.5%
70	Inception-V3	FT	UMAP	K-Means	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
71	Inception-V3	FT	UMAP	DBSCAN	25.0%	50.0%	0.0%	20.0%	25.0%	75.0%	32.5%
72	Inception-V3	FT	UMAP	HDBSCAN	50.0%	0.0%	0.0%	0.0%	0.0%	0.0%	8.3%
73	Xception	None	None	K-Means	50.0%	75.0%	87.5%	40.0%	100.0%	75.0%	71.2%
74	Xception	None	None	DBSCAN	25.0%	0.0%	37.5%	0.0%	25.0%	25.0%	18.8%
75	Xception	None	None	HDBSCAN	100.0%	100.0%	62.5%	60.0%	50.0%	100.0%	78.8%
76	Xception	None	PCA	K-Means	50.0%	50.0%	37.5%	0.0%	75.0%	0.0%	35.4%
77	Xception	None	PCA	DBSCAN	50.0%	75.0%	50.0%	0.0%	75.0%	0.0%	41.7%
78	Xception	None	PCA	HDBSCAN	100.0%	50.0%	0.0%	80.0%	25.0%	0.0%	42.5%
79	Xception	None	UMAP	K-Means	75.0%	75.0%	62.5%	40.0%	100.0%	50.0%	67.1%
80	Xception	None	UMAP	DBSCAN	100.0%	100.0%	100.0%	40.0%	100.0%	100.0%	90.0%
81	Xception	None	UMAP	HDBSCAN	50.0%	75.0%	0.0%	60.0%	100.0%	25.0%	51.7%
82	Xception	FT	None	K-Means	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
83	Xception	FT	None	DBSCAN	0.0%	0.0%	0.0%	0.0%	0.0%	25.0%	4.2%
84	Xception	FT	None	HDBSCAN	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%	16.7%
85	Xception	FT	PCA	K-Means	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
86	Xception	FT	PCA	DBSCAN	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
87	Xception	FT	PCA	HDBSCAN	75.0%	50.0%	0.0%	0.0%	0.0%	0.0%	20.8%
88	Xception	FT	UMAP	K-Means	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
89	Xception	FT	UMAP	DBSCAN	0.0%	0.0%	0.0%	60.0%	50.0%	0.0%	18.3%
90	Xception	FT	UMAP	HDBSCAN	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
91	AE	None	None	K-Means	0.0%	25.0%	12.5%	20.0%	50.0%	50.0%	26.2%
92	AE	None	None	DBSCAN	0.0%	25.0%	0.0%	0.0%	0.0%	0.0%	4.2%
93	AE	None	None	HDBSCAN	0.0%	25.0%	0.0%	40.0%	0.0%	25.0%	15.0%
94	AE	None	PCA	K-Means	0.0%	25.0%	12.5%	0.0%	50.0%	50.0%	22.9%
95	AE	None	PCA	DBSCAN	0.0%	25.0%	0.0%	20.0%	50.0%	50.0%	24.2%
96	AE	None	PCA	HDBSCAN	100.0%	50.0%	0.0%	80.0%	50.0%	0.0%	46.7%

Table 15. Percentage of faulty scenarios covered by the root cause clusters generated for each pipeline. The last column represents the average of averages.

Pipelines					Case Study Subjects						
#	FE	FT	DR	CA	GD	OC	HPD	SVIRO	SAP	CPD	Avg.
97	AE	None	UMAP	K-Means	0.0%	0.0%	0.0%	0.0%	25.0%	50.0%	12.5%
98	AE	None	UMAP	DBSCAN	50.0%	50.0%	37.5%	40.0%	50.0%	100.0%	54.6%
99	AE	None	UMAP	HDBSCAN	0.0%	25.0%	0.0%	20.0%	0.0%	0.0%	7.5%

C ADDITIONAL MATERIAL FOR RQ3

Table 16. Distribution of faults for the different failure inducing sets for each case study subject.

Case Study	Dataset	Injected Failure Scenarios									
		N	B	S	D	M	H	SG	EG	EO	NF
GD	GD_1	64	40	48	72	-	-	-	-	-	56
	GD_2	48	32	24	16	-	-	-	-	-	64
	GD_3	24	64	16	40	-	-	-	-	-	8
	GD_4	72	16	24	48	-	-	-	-	-	40
	GD_5	40	24	72	16	-	-	-	-	-	8
	GD_6	56	32	48	16	-	-	-	-	-	24
	GD_7	64	32	8	56	-	-	-	-	-	24
	GD_8	72	24	16	8	-	-	-	-	-	64
	GD_9	40	64	48	16	-	-	-	-	-	56
	GD_10	56	8	64	24	-	-	-	-	-	48
OC	OC_1	18	6	10	4	-	-	-	-	-	8
	OC_2	4	18	12	2	-	-	-	-	-	14
	OC_3	2	14	10	6	-	-	-	-	-	16
	OC_4	2	4	6	10	-	-	-	-	-	8
	OC_5	6	4	16	18	-	-	-	-	-	10
	OC_6	8	6	10	12	-	-	-	-	-	16
	OC_7	18	8	16	6	-	-	-	-	-	2
	OC_8	16	18	14	10	-	-	-	-	-	4
	OC_9	14	16	4	10	-	-	-	-	-	2
	OC_10	10	2	14	8	-	-	-	-	-	18
HPD	HPD_1	45	72	54	9	36	63	81	18	-	27
	HPD_2	27	81	45	18	54	63	72	36	-	9
	HPD_3	54	81	27	63	18	45	9	36	-	72
	HPD_4	36	18	63	72	9	81	54	27	-	45
	HPD_5	27	63	18	72	36	9	45	81	-	54
	HPD_6	45	36	54	63	81	9	72	27	-	18
	HPD_7	63	45	81	36	27	72	18	54	-	9
	HPD_8	72	9	63	27	36	18	81	54	-	45
	HPD_9	72	63	18	27	45	9	81	54	-	36
	HPD_10	54	81	63	27	45	72	18	9	-	36
SVIRO	SVIRO_1	6	12	18	21	-	-	-	-	24	15
	SVIRO_2	9	24	6	12	-	-	-	-	15	3
	SVIRO_3	15	18	21	3	-	-	-	-	27	9
	SVIRO_4	21	9	12	24	-	-	-	-	6	27
	SVIRO_5	27	21	3	9	-	-	-	-	18	24
	SVIRO_6	3	27	24	6	-	-	-	-	21	12
	SVIRO_7	24	6	15	18	-	-	-	-	3	21
	SVIRO_8	15	3	27	24	-	-	-	-	12	6
	SVIRO_9	12	15	3	27	-	-	-	-	9	18
	SVIRO_10	18	21	9	15	-	-	-	-	21	18
CPD	CPD_1	87	74	55	28	-	-	-	-	-	44
	CPD_2	43	56	27	22	-	-	-	-	-	74
	CPD_3	6	32	4	62	-	-	-	-	-	35
	CPD_4	49	22	88	34	-	-	-	-	-	5
	CPD_5	24	69	37	57	-	-	-	-	-	86
	CPD_6	13	69	58	54	-	-	-	-	-	25
	CPD_7	3	32	51	9	-	-	-	-	-	59
	CPD_8	77	62	12	53	-	-	-	-	-	4
	CPD_9	85	27	78	30	-	-	-	-	-	62
	CPD_10	65	46	66	89	-	-	-	-	-	40

Table 16. Distribution of faults for the different failure inducing sets for each case study subject.

Case Study	Dataset	Injected Failure Scenarios									
		N	B	S	D	M	H	SG	EG	EO	NF
SAP	SAP_1	22	33	54	48	-	-	-	-	-	72
	SAP_2	75	22	48	17	-	-	-	-	-	57
	SAP_3	22	4	57	42	-	-	-	-	-	81
	SAP_4	74	21	40	36	-	-	-	-	-	42
	SAP_5	15	14	86	74	-	-	-	-	-	51
	SAP_6	73	60	2	83	-	-	-	-	-	72
	SAP_7	58	57	47	83	-	-	-	-	-	43
	SAP_8	6	75	26	16	-	-	-	-	-	70
	SAP_9	89	86	66	32	-	-	-	-	-	68
	SAP_10	67	77	14	4	-	-	-	-	-	55