An Experimental Comparison of Anomaly Detection Methods for Collaborative Robot Manipulators

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY-NC-SA 4.0

SUBMISSION DATE / POSTED DATE

24-01-2022 / 31-01-2022

CITATION

Graabæk, Søren Goddixen; Ancker, Emil Vincent; Christensen, Anders Lyhne; Fugl, Andreas Rune (2022): An Experimental Comparison of Anomaly Detection Methods for Collaborative Robot Manipulators. TechRxiv. Preprint. <https://doi.org/10.36227/techrxiv.19006643.v1>

DOI

[10.36227/techrxiv.19006643.v1](https://doi.org/10.36227/techrxiv.19006643.v1)

An Experimental Comparison of Anomaly Detection Methods for Collaborative Robot Manipulators

Søren G. Graabæk*, Emil Vincent Ancker*,
Andreas Rune Fugl, and Anders Lyhne Christensen, *Senior Member, IEEE*

Abstract—A key advantage of collaborative robots is their ability to operate in non-shielded environments alongside humans. However, the lack of a highly controlled environment makes robots prone to external perturbations that can compromise their efficiency if not detected in a timely manner. In this paper, we perform an experimental comparison of methods for detecting exogenous anomalies. We design a typical pick-and-place application for a collaborative manipulator with six degrees of freedom and we identify eight common anomalies. In an extensive experimental campaign comprising 600 runs under normal operation, and 80 runs where exogenous perturbations are present, we collect data on the robot’s joints, force-torque readings, and the time spent in each program node. We then train a set of common anomaly detection methods whose computational complexity is low enough to run on resource-constrained robot hardware. We find that several of the evaluated methods can achieve a high anomaly detection performance. We show that exploiting knowledge of the robot’s program tree can increase the performance for some types of anomalies. We also observe that performance, in general, deteriorates when the application includes: (i) rarely-visited program branches, (ii) physical contact with the environment, and (iii) stochastic trajectories.

Note to Practitioners—There exist a large number of methods that can be used for anomaly detection/fault detection in collaborative robots. However, studies on these methods tend to only focus on a single or a couple of such methods, which can make it challenging to gauge their relative merits in specific robot scenarios. In this paper, we conduct a comprehensive comparison of 15 methods for anomaly detection, including methods based on principle component analysis, local outlier factor, and autoencoders. The methods are assessed in a typical pick-and-place application with respect to their capacity to detect a broad range of exogenous anomalies. The results of the study show that several methods perform well, but that their performance profiles differ across the studied anomalies. The results also give an indication of the application characteristics that have the potential to make anomaly detection challenging.

Index Terms—Anomaly detection, collaborative robots, machine learning, neural networks, semi-supervised learning.

I. INTRODUCTION

WITH the current rate of technological improvement, collaborative robots (cobots) can be deployed in an increasing broader set of scenarios [1], including in environments with less structure than found in traditional robot manufacturing. Ideally, cobots can easily be reprogrammed, e.g. via drag-and-drop programming, to take on new tasks and collaborate seamlessly with humans [2]. However, when the operating environment is not highly controlled and frequently undergoes modifications, robot applications become

vulnerable to exogenous anomalies. Exogenous anomalies can negatively impact a robot’s efficiency, and severe anomalies can even cause robot breakdown. By detecting anomalies in a timely and reliable manner, predictive maintenance on the robot or its surroundings can be performed [3], which in turn can lead to improved overall system performance.

Anomaly detection refers to the problem of finding patterns in data that do not conform to expected behavior [4]. Since anomalies may be symptoms of a fault, anomaly detection has been widely studied as a means for fault detection in manufacturing robots, such as robot manipulators [5]. In most studies, however, only a single method or class of methods for anomaly detection are evaluated. In this study, we evaluate several reconstruction-based methods and benchmark them against two relatively simple instance-based methods. We collect data in a typical pick-and-place application for a general collaborative manipulator, a UR5e from Universal Robots, which has six Degree of Freedom (DoF) and the ability to lift a total payload of five kilograms. The robot has a set of sensors that enable measurements of the robot’s joint configurations, joint velocities, joint currents, and six-dimensional force-torque at the Tool Center Point (TCP). We use these features, combined with a feature describing the time spent in the robot’s current program node, and train a total of 15 different anomaly detection models. The discriminatory power of the trained models is then evaluated on eight common exogenous anomalies in collaborative robots applications in terms of the Precision Recall Area Under Curve (PR-AUC) and Receiver Operating Characteristic Area Under Curve (ROC-AUC) metrics [6].

We assess two distance-based methods, k -Nearest Neighbors (k -NN) and Local Outlier Factor (LOF), and find that they perform well in our pick-and-place application for the anomalies considered. We also experiment with different reconstruction-based methods, including methods based on autoencoders, and find that utilizing statistical representations (variational inference) in the latent space, can lead to improved performance. We evaluate the use of contextual information in the form of the program node ID currently being executed, and find that leveraging such information can significantly improve discriminatory power. In this study, we also analyze which aspects of the robot’s task make anomaly detection challenging in general, and identify three elements that negatively impact performance, namely rarely-visited program branches, variations in the task, and physical contact with the environment.

The remainder of the paper is organized as follows: in Section II, we discuss previous studies on anomaly detec-

*Authors contributed equally to the project.

tion and fault detection applied to manufacturing robots. In Section III, we present the hardware and data collection procedure used in our study. In Section IV, we present the anomaly detection models that we evaluate and compare. In Section V, we detail our experimental setup and the exogenous anomalies considered. In Section VI, we present and compare the anomaly detection performance achieved by the respective methods studied. In Section VII, we discuss the results and avenues for future work, and in Section VIII, we provide concluding remarks.

II. RELATED WORK

Fault detection and diagnosis approaches are typically divided into three categories [7]: data-driven, model-based, and knowledge-based. Model-based approaches use mathematical models or logic to describe the system and its components. De Luca and Mattone [8] use a model of the robot's dynamic and a linear state observer to detect unknown actuator faults based on the residual between the physical system and the model. They achieve a good sensitivity, with no false positives, using a low-pass filter to remove data noise and a low detection threshold. Knowledge-based approaches typically associate recognized behaviors with predefined known faults and can thus be used for fault diagnoses. Fantuzzi et al. [9] use an expert system to detect and isolate a well-known set of faults using simple logic on the sensor data. A major issue with knowledge-based systems and model-based approaches is that they require a priori knowledge or a reasonably accurate analytical model, which may not be feasible to obtain. Data-driven approaches, on the other hand, are model-free and thus have the potential to detect unknown faults since a system model is learned and adapted through training. Data-driven methods are often computationally expensive to train, but due to a reduced cost of computational power, they are becoming increasingly popular [7]. Below, we review key data-driven methods and past studies on their application in robotics.

Anomaly detection can be used to pinpoint data instances potentially caused by faults in a robotic system. The data type and the nature of the given data determine which detection methods are appropriate and practical to use [4]; e.g., Mahalanobis distances could be used to detect anomalies in datasets composed of normally distributed real numbers, but may not be applicable in other cases. Sathish et al. [10] study the effect of the data source and type of training data on the detection performance for endogenous faults (gearbox, motor, brake, and computer unit) in industrial robots using Principal Component Analysis (PCA). The authors recorded time series containing mechanical measurements, e.g., TCP speed and joint torque, and non-mechanical measurements, e.g., CPU temperature and fan speed. The data was recorded on ten robots performing different tasks in the same application domain and the data was divided into datasets, each consisting of 20 days of recording. The authors' initial assumption was that raw data features are more representative of the current task, whereas data features based on absolute difference capture variation in data and are

therefore more representative of the condition of the robot (i.e. *normal* or *faulty*). Their analysis showed that it is indeed possible to use a single reference robot to generate training data for testing multiple scenarios provided that the PCA model is trained on absolute difference data features.

Goncharov et al. [11] present a simple statistical approach for detecting exogenous anomalies in a pick-and-place application on a single robot. Specifically, the authors' aim was to detect a change in the robot's workload using the actuator currents as the only data feature. They calculated a detection threshold based on the summary statistics (minimum, maximum, and mean) for each robot operation. A limitation of the approach is that it can only detect external influences that excite the actuator currents above the extreme values.

Borgi et al. [5] proposes to build a linear regression model of the relation between features extracted from actuator current and the robot's accuracy, measured by external laser sensors. The authors recorded the actuator energy consumption from a robot moving between two configurations and calculated summary statistics (mean, standard deviation, skewness, and kurtosis). They then analyzed the Pearson correlation between the data features and the robot accuracy and found that the actuator energy and the mean and standard deviation of the actuator currents are strongly anti-correlated with the robot accuracy; A regression model was therefore trained on these features. The authors did not introduce any faults in their setup, but the predicted accuracy was highly correlated with the actual robot accuracy.

The actuator current can be preprocessed in many ways to train models or not processed at all. Sabry et al. [12] used Bode Equation Vector Fitting (BEVF) for fitting the model to the raw power measurements of a six DoF ABB IRB industrial manipulator as a function of the task, speed, load, temperature, and time: $P_t = f(task, speed, load, temp, time)$. The authors collected a time series dataset of the power consumption during two types of tasks using a sampling rate of up to 123 Hz. The first task was a pick-and-place task where the robot lifted a 200 g payload, and the second task was a polishing task. To detect anomalies, an analytical mathematical power model was calibrated to the individual joints of the robot. After the joint calibration, a BEVF was fitted for each of the tasks. More specifically, the power pattern was divided into intervals according to any change in the excitation of the joints and a detection threshold was set to $\pm 10\%$ of the reference/model values. The authors introduced three types of fault: (i) a spring was attached between the two links in the elbow joint, (ii) the electrical signal was removed from joint 2 and the drive belt was removed from joint 5, and (iii) an encoder fault on joint 2 was additionally introduced to (ii). The presented method achieved a good fault detection performance on the studied faults. However, the approach is not directly applicable to dynamic task environments since it relies on a fixed and accurate power consumption pattern.

In [13], Wu et al. propose a generative Bayesian non-parametric Hidden Markov Model (HMM) for encoding spatial and temporal dependencies in robot data in a low-dimensional hidden-state space. The non-parametric method enables the automatic inference of the number of latent states

from recorded data [14]. The authors evaluated the approach on the right end-effector of a Baxter humanoid robot in a human-robot collaboration task. The human was tasked with feeding six common household objects to the robot consisting of box-like shapes and bottles. During the task, seven types of exogenous anomalies were introduced: human collision, tool collision, object slip, human collision with object, wall collision, no object, and others. The anomalies were detected using a six DoF force-torque data feature, a six DoF Cartesian velocity data feature, and two 56 DoF tactile data features, one for each of the left and right tactile sensors. To make the data features more generic, they were normalized within their category, i.e., the three DoF forces were normalized, the three DoF torques were normalized, and so on. Lastly, the normalized data features and the raw data features (except the raw tactile data features) were combined into a single data feature used for training the non-parametric HMM. The authors achieved high sensitivity and good detection performance across varying demonstrations and tasks by adjusting the anomaly detection threshold during the task. The dynamic threshold was a function of the maximum and minimum of the log-likelihood of observations for latent states.

Using statistical methods to obtain a representation of the latent space has also provided good anomaly detection results in other domains, e.g., aviation [15], and the methods are also commonly used for feature engineering in computer vision applications [16]. Recent papers [17, 18] report on an attempt to incorporate the statistical latent space methods into reconstruction-based anomaly detection methods. Reconstruction-based methods assume that anomalous data instances lose more information than normal data instances when they are projected to a lower dimension space by an encoder, and hence that anomalous data instances are not properly reconstructed by the decoder. Reconstruction-based methods are often based on Autoencoders (AEs) [16].

Park et al. [17] used a PR2 humanoid robot for a robot-assisted feeding task. The authors developed a real-time anomaly detector to detect 12 types of exogenous anomalies: touch by user, aggressive eating, utensil collision by user, sound from user, face occlusion, utensil miss by user, unreachable location, environmental collision, environmental noise, utensil miss by system fault, utensil collision by system fault, and system freeze. The data consisted of 17 raw sensor features recorded from the robot: joint encoders, current sensors, force sensors, a microphone, and a camera. The presented AE method used *variational inference* [19, 20] to model the underlying latent space distribution of robot data and a *Long Short-Term Memory (LSTM)* structure [21] to model temporal dependencies. The AE structure was trained on the 17 raw sensor data using a denoising autoencoder criterion [22, 23] and it outperformed a likelihood-based classifier (Hidden Markov Model) trained on four hand-engineered data features developed by the same authors.

In [18], Chen et al. present a recent anomaly detection study on a KUKA KR6R 900 SIXX doing a pick-and-place application. The approach embeds variational inference and a Convolutional Neural Network (CNN) [24] into the AE model

Fig. 1. A screenshot of the *teach pendant*, with a visualization of the program nodes and program node IDs.

instead of LSTM neurons, because the authors claim that the CNN is better at capturing data dependencies. The proposed method achieved promising results but had to be trained offline due to a very large number of trainable parameters. The method was trained on 12 standardized sensor signals sampled at 33.3 Hz. The data features consisted of the six joint positions and currents.

In the following section, we review the theory behind some of the key methods used in the studies discussed above along with anomaly detection methods commonly used in the broader anomaly detection literature [4, 25, 26, 6].

III. ROBOTIC HARDWARE AND DATA COLLECTION

We use a UR5e [27] manipulator from Universal Robots in this study. The UR5e has six joints with a wide working range of $\pm 360^\circ$ and a maximum joint speed of $180^\circ/\text{s}$ for all joints. The robot's 850 mm reach and ± 0.03 mm pose repeatability makes it suitable for many high-precision robot tasks. The *teach pendant* provides an drag-and-drop programming interface that creates a program tree consisting of multiple robot operations. The teach pendant and a snippet from a program tree with the program nodes and program node IDs is shown in Fig. 1.

To collect data from the robot, we connected a computer to the robot's Real-Time Data Exchange (RTDE) interface [28]. The RTDE interface can sampled sensor data and information about the robot's state at a maximum frequency of 500 Hz. The sampled features are listed in Table I and they consist of 25 behavioral features and two contextual features; one being a custom designed feature, *Time_in_node*, derived from from the *timestamp* and the *Program node ID* feature. We standardize the features by computing a shared mean and variance for each category, e.g., all current measurements share a common mean and variance:

$$x'_j = \frac{x_j - \hat{\mu}_{c(j)}}{\hat{\sigma}_{c(j)}}, \quad (1)$$

where $\hat{\mu}_{c(j)}$ is the shared mean for all features belonging to category $c(j)$ and $\hat{\sigma}_{c(j)}$ is the shared standard deviation for all features belonging to category $c(j)$.

TABLE I
SUMMARY OF USED DATA FEATURE. B: BEHAVIORAL. C: CONTEXTUAL.

Feature	Size	B	C	Description
Joint configuration	6	✓		Relates to robot movements.
Joint velocity	6	✓		Relates to robot movements.
Joint current	6	✓		Relates to external wrench and robot movements.
TCP F/T sensor	6	✓		Relates to external wrench.
Timestamp	1		✓	Establishes temporal context.
Program node ID	1		✓	Establishes robot program context.
Time_in_node	1	✓		Provides temporal behavioral information.

IV. ANOMALY DETECTION METHODS COMPARED

In this study, we implement 15 different methods for anomaly detection and all the presented methods are trained semi-supervised. The training data consist of the 25 behavioral features listed in Table I recorded on a robot during normal operation; and for some of the methods, the two contextual features are also used (see details below). The methods will be evaluated based on their discriminatory power.

A. Instance-based Methods

Instance-based models are popular in the outlier analysis domain because of their simplicity, effectiveness, and intuitive nature [6, 29, 30]. The methods are trained by organizing the training data in memory, and the outlier score is calculated by computing the distance between the query data and the data stored in memory, based on a prespecified distance metric. The methods can be updated online by adding new data instances to the stored data. The outlier score is valid under the assumption that anomalous data instances have a significantly larger distance to their neighbors than normal data instances, measured under the specified distance metric. Thus, the performance of the method depends on the chosen distance metric. Basic instance-based methods have a space and computational complexity of $\mathcal{O}(N)$ [6], where N is the number of training instances. In our implementations, we use a k -dimensional index trees to lower the query and training computational complexity to $\mathcal{O}(\log N)$ [31, 32].

In this study, we evaluate a k -NN [6] and LOF [33] algorithm as performance baselines due to their relative simplicity, high performance, and their ability to detect anomalies in local neighborhoods. Ensemble structures have been shown to increase performance and decrease parameter sensitivity [34, 35]. Thus, based on preliminary experiments, we implemented both methods using a parametric ensemble consisting of ten ensemble components randomly initialized with $k \in [5, 50]$. The remaining parameters for the methods were also tuned based on preliminary experiments. In the k -NN ensemble, we compute the average of all ensemble components' k th neighbor distances.

For the LOF ensemble, a data instance is considered an outlier when its score deviates significantly from 1 [36]. Hence, we calculate the outlier score as the average of the ensemble components' absolute difference from 1.

B. Explicit Generalization Models

Explicit generalization models [6] are an alternative to instance-based methods that capture the normal behavior in

a fixed size model, e.g., linear regression models, Boltzmann Machines, and Neural Networks. Generalization models often benefit from a low memory requirement and many possible extensions of the basic methods.

1) *Principal Component Analysis*: PCA can either be used as a standalone anomaly detection method or to extract data features for another method as part of a hybrid method [25]. In both cases, PCA generate a k -dimensional uncorrelated latent space (with $k < d$ principal components) using a linear combination of the original d data features. The k latent space features are then either fed to another detection method, or a reconstruction error is computed by re-projecting the features to the original data space:

$$s = \left\| \mathbf{x} - \sum_{j=1}^k \mathbf{e}_j^T \mathbf{x} \mathbf{e}_j \right\|, \quad (2)$$

where s is the reconstruction error, $\mathbf{e}_j$ is the j th eigenvector, and $\mathbf{x}$ is the d dimensional input feature. PCA can only describe linear relations between the features which motivates the usage of nonlinear explicit generalization methods, e.g., kernel-PCA and neural network-based AE structures. In our PCA implementation, we use the randomized Singular Value Decomposition (SVD) algorithm [37] for obtaining the principal components since the standard SVD has a quadratic memory $\mathcal{O}(N^2)$ requirement.

We evaluate three PCA variants: (i) the standard PCA as described above, (ii) the Probabilistic Principal Component Analysis (PPCA) [38], and (iii) the Incremental Principal Component Analysis (IPCA) [39]. PCA and IPCA differ only by how the model is fitted; IPCA is trained in small batches, whereas PCA is deterministically found using all data instances at once. PPCA is trained like PCA, but it returns the probability of the query belonging to the learned model, instead of the Euclidean distance. We initialize each method as a parametric ensemble consisting of five ensemble components with $k \in [4, 12]$, we compute the outlier score by averaging the reconstruction error of each ensemble component. When training the IPCA, we use a batch size of 128. These parameters were found in preliminary experiments.

We also implemented a hybrid model consisting of an IPCA model ($k = 6$) for feature extraction and a One-Class Support Vector Machine (OCSVM) as the detector. We initialize it as an ensemble consisting of 30 OCSVM components using the Radial Basis Functions (RBF) as the kernels. We use a six-dimensional latent space because it yielded the best results in preliminary experiments, and we observed no gain in performance with more than 30 ensemble components. Additionally, to mitigate the long training time resulting from OCSVM's computational complexity of $\mathcal{O}(N^2)$, we train it using 1,000 randomly sampled data instances, as recommended in [6]. We combine the ensemble outputs by averaging the scores.

2) *Autoencoder Structures*: In an AE, the Neural Network (NN) can be thought of as two sub-networks: an encoder mapping the data features from the original space $\mathbf{x} \in \mathbb{R}^{d_x}$ to a latent space $\mathbf{z} \in \mathbb{R}^{d_z}$, $\mathbf{z} = e(\mathbf{x})$, and a decoder reconstructing the data features to the original space, $\hat{\mathbf{x}} = d(\mathbf{z})$; similar

to PCA. To avoid the network simply passing the input to the output, the AE structure must be regularized e.g., by using an undercomplete network, or the loss function must be regularizing [22], e.g. using the L2 norm. In the following subsections, we introduce the networks compared in this paper.

The performance of NN methods is highly sensitive to the choice of hyperparameters [40]. Therefore, an appropriate strategy for choosing hyperparameters must be considered when designing the autoencoders. We use Hyperband [41] from the keras-tuner module [42] to tune appropriate network parameters, e.g., layer size, activation function, and learning rate. We use Hyperband because of its focus on rapidly testing of configurations and fast elimination of non-promising configurations. We configure Hyperband to evaluate the models using an unsupervised loss function, i.e., reconstruction error. To prevent Hyperband from exploring very deep networks, we fix the structure of all autoencoders to consist of: (i) one encoding layer, (ii) a latent space layer, and (iii) one decoding layer. Hyperband tends to prematurely terminate models with a slow convergence rate, thus prioritizing models with a fast convergence rate even though they might converge to a suboptimal solution. We mitigated this issue by pinpointing multiple smaller search spaces using a random search tuner and then running Hyperband on the subspaces.

We use dense layers for all network layers and train all autoencoder variants using the Adam optimizer [43] for 30 epochs at a batch size of 128, unless stated otherwise. The parameters explicitly mentioned in the text are found using Hyperband, the remaining parameters are set to their default value given by scikit-learn (v. 0.23.2) [44] or TensorFlow (v. 2.4.0) [45].

a) Undercomplete Autoencoder: The Undercomplete Autoencoder (UAE) is one of the simplest AE structures and it regulates the network using an hourglass structure $d_x > d_z$, where d_x is the dimension of the input and d_z is the dimension of the latent space. We train the UAE using Mean Squared Error (MSE) as the loss function and a learning rate of $\alpha = 0.001$. The network details are shown in Table II.

TABLE II
THE IMPLEMENTED UAE NETWORK STRUCTURE .

Layer	Layer size	Activation function
x	25	
$h_1 = \text{Dense}(x)$	22	tanh
$z = \text{Dense}(h_1)$	14	tanh
$h_2 = \text{Dense}(z)$	22	tanh
$\hat{x} = \text{Dense}(h_2)$	25	linear

b) Denoising Autoencoder: In the Denoising Autoencoder (DAE) [22], the input data is intentionally corrupted during training, e.g., by adding isotropic Gaussian noise to all data instances, to make the network robust to noise:

$$\tilde{x} = x + \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \epsilon \mathbf{I}). \quad (3)$$

The loss is computed by calculating the MSE loss between the noisy data propagated through the network and the original noise-free data:

$$\mathcal{L}(x, d(e(\tilde{x}))). \quad (4)$$

In our implementation, we reuse the UAE network and add Gaussian noise, $\epsilon = 0.1$, to all data features except the six TCP force-torque features. We do not add noise to the TCP force-torque features as they are subject to sensor noise and the artificial noise reduced the performance of the method.

c) Autoencoder Ensemble: Chen et al. [35] introduced the Randomized Neural Network (RandNet) for Outlier Detection. RandNet is an autoencoder ensemble that has shown promising anomaly detection performance on some datasets from the UCI Machine Learning Repository¹. The networks are randomized in the sense that some of the layer connections are randomly discarded to create diversity between the ensemble components.

We implement a RandNet consisting of five ensemble components instead of the proposed 100 ensemble components used by the authors because our network is shallower and, thus, the additional components did not improve the performance. We randomly remove connections as proposed by the authors. The network details are shown in Table II.

TABLE III
THE IMPLEMENTED RANDNET NETWORK STRUCTURE.

Layer	Layer size	Activation function
x	25	
$h_1 = \text{Layer}(x)$	18	tanh
$z = \text{Layer}(h_1)$	12	tanh
$h_2 = \text{Layer}(z)$	18	tanh
$\hat{x} = \text{Layer}(h_2)$	25	linear

Fig. 2. Visualization of the VAE structure.

d) Variational Autoencoder: The Variational Autoencoder (VAE) is a generative model based on variational inference [19], and the latent space is thus a posterior distribution $q(z|x)$ instead of a feature vector. The encoder network is responsible for approximating the true posterior distribution of the latent features, $p(z|x)$. We assume that the latent features are distributed according to an isotropic Gaussian:

$$z \sim q(z|x) = \mathcal{N}(\mu_z, \sigma_z^2 \mathbf{I}). \quad (5)$$

We then use the reparameterization trick [19] to sample from the learned latent space posterior and feed it through

¹<https://archive.ics.uci.edu/ml/index.php>

the decoder network used to reconstruct the data instance into the original feature space:

$$\hat{x} = \mu_z + \sigma_z \cdot \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (6)$$

When training the VAE we minimize the reconstruction error and regularize the loss function using the Kullback-Leibler (KL) divergence between an isotropic Gaussian prior with unit variances, $p(z)$, and the approximated latent space posterior distribution, $q(z|x)$:

$$\mathcal{L}_{\text{VAE}} = \mathbb{E}_{z \sim q}[\log p(x|z)] - D_{KL}(q(z|x)||p(z)). \quad (7)$$

When using MSE as reconstruction error, it is equal to assuming that the original feature space posterior, $p(x|z)$, has unit variance [46]. Due to our standardization method presented in Section III, we know this assumption does not hold. Thus, we expand the decoder network to learn an isotropic multivariate Gaussian as the output, as visualized in Fig. 2. With an appropriate choice of batch size (greater than 100), we can directly use the computed log-likelihood as the expected log-likelihood in the VAE's loss function [19]. For a detailed explanation of auto-encoding variational Bayes, we refer the reader to [19, 47].

We train the network using a learning rate of $\alpha = 0.0005$, and Table V summarizes the VAE network.

TABLE IV
THE IMPLEMENTED VAE NETWORK STRUCTURE.

Layer	Layer size	Activation function
x	25	
$h_1 = \text{Dense}(x)$	26	ReLU
$\mu_z = \text{Dense}(h_1)$	24	linear
$\log(\sigma_z) = \text{Dense}(h_1)$	24	linear
$z = \text{Sample}(\mathcal{N}(\mu_z, \sigma_z^2 \mathbf{I}))$	24	-
$h_2 = \text{Dense}(z)$	26	ReLU
$\mu_x = \text{Dense}(h_2)$	25	linear
$\sigma_x = \text{Dense}(h_2)$	25	softplus

e) Clustering the Latent Space of the Variational Autoencoder: In preliminary experiments with the VAE structure, we observed that instances sampled within the same program nodes from the teach pendant (see Fig. 1) tend to cluster in the latent space. Fig. 3 illustrated this phenomenon with a 2-dimensional t-SNE plot [48]. We therefore construct a deep hybrid model using a Gaussian Mixture Model (GMM) [49] to cluster the latent space of the VAE.

In our implementation, we reuse the VAE network structure from Table V and we initialize the GMM with the same number of modes as the number of program nodes in the robot's program tree.

f) Conditional Variational Autoencoder: The Conditional Variational Autoencoder (CVAE) structure is an extension to the VAE structure in which a set of conditional (contextual) features are both concatenated with the input data to encoder and decoder network [46]. Providing contextual information to the network can potentially reduce the complexity of the feature mapping the method needs to learn.

We use the *program node ID* as categorical data and provide it as a one-hot encoded feature for the CVAE, since this transformation is commonly used when passing categorical

Fig. 3. The 2D t-SNE representation of the latent space feature encoded by the VAE structure used on the behavioral features for data collected during the normal behavior of a robot program. We apply a color to each data instance based on the program node it is collected under; data instances collected under the same program nodes get assigned the same color.

data to neural networks [50]. The network is trained using a learning rate of $\alpha = 0.0005$ and the network details are listed in Table V.

TABLE V
THE IMPLEMENTED CVAE NETWORK STRUCTURE.

Layer	Layer size	Activation function
x	$25 + \text{len}(\text{onehot})$	
$h_1 = \text{Dense}(x)$	26	ReLU
$\mu_z = \text{Dense}(h_1)$	8	linear
$\log(\sigma_z) = \text{Dense}(h_1)$	8	linear
$z = \text{Sample}(\mathcal{N}(\mu_z, \sigma_z^2 \mathbf{I}))$	$8 + \text{len}(\text{onehot})$	-
$h_2 = \text{Dense}(z)$	26	ReLU
$\mu_x = \text{Dense}(h_2)$	25	linear
$\sigma_x = \text{Dense}(h_2)$	25	softplus

3) Time Series Networks: Using time series models, such as LSTM networks or CNN networks, for anomaly detection in time series data has yielded promising results in previous studies [17, 51, 18]. However, the structures can rapidly accumulate many trainable parameters and often have a high tuning complexity to achieve good performance (see [18] for an example).

We briefly explore some time series methods to inspect if the extra model complexity introduced by analyzing sequences of data instances leads to a considerable increase in performance. The methods are implemented as estimation model-based methods as defined in [26]; we use the data sequence of length, t_{len} , consisting of the past and present data points, $\{x_{t-t_{\text{len}}}, \dots, x_t\}$, to calculate the expected value x_t .

a) Convolutional Neural Network: Munir et al. [52] designed a general CNN for anomaly detection called DeepAnT using 1D convolutional layers. We implement the network as it is presented in the paper except for the sequence length, which we set to ten during training. The network structure is shown in Table VI.

b) Long Short-Term Memory Autoencoder: We implement the LSTM with the network structure listed in Table VII, and we train the network using a learning rate of $\alpha = 0.001$. The network structure and sequence length were found using Hyperband.

c) Bilinear Autoencoder: Tran et al. [53, 54] recently presented a neural network based on Bilinear Layers (BL)

TABLE VI
THE IMPLEMENTED DEEPANT [52] NETWORK STRUCTURE.

Layer	Layer size	Kernel size	Filters	Activation function
x	(10, 25)			
$h = \text{Conv1D}(x)$	(32, 23)	3	32	ReLU
$h = \text{MaxPool1D}(h)$	(32, 11)	2	-	-
$h = \text{Conv1D}(h)$	(32, 9)	3	32	ReLU
$h = \text{MaxPool1D}(h)$	(32, 4)	2	-	-
$h = \text{Flatten}(h)$	128			-
$h = \text{Dense}(h)$	40			ReLU
$\hat{x} = \text{Dense}(h)$	25			linear

TABLE VII
THE IMPLEMENTED LSTM NETWORK STRUCTURE.

Layer	Layer size	Activation function
x	(10, 25)	
$h_1 = \text{LSTM}(x)$	(10, 24)	tanh
$h_2 = \text{LSTM}(h_1)$	6	tanh
$z = \text{RepeatVector}(h_2)$	(10, 6)	-
$h_3 = \text{LSTM}(z)$	(10, 6)	tanh
$h_4 = \text{LSTM}(h_3)$	24	tanh
$\hat{x} = \text{Dense}(h_4)$	25	linear

and found that their bilinear network could outperform a CNN and a LSTM, both in terms of computational complexity and precision, in stock price prediction. Furthermore, the concept of bilinear mapping has recently yielded promising results for time series anomaly detection in dynamic maps [55] among others, and we, therefore, evaluate the model in this study.

The bilinear map is a function that combines elements of two vector spaces into a third vector space [53]. In our case, the function maps a multivariate time series, $X \in \mathcal{R}^{d_x \times T}$, into a latent space using $W_1 \in \mathcal{R}^{d_z \times d_x}$ and $W_2 \in \mathcal{R}^{T \times d_t}$:

$$z = F(X) = W_1 X W_2, \quad (8)$$

where W_1 can be interpreted as mapping the d_x -dimensional representation of each time series to a d_z -dimensional subspace, while W_2 maps the contribution of the T times series into d_t times series using a weighted average approach.

We stack the Bilinear layer and train it as an autoencoder network. We use the MSE reconstruction error as the loss function and we regularize the network with a dropout rate of 0.1 and L_2 max-norm of 1.5 as recommended in [54]. We structure the network as shown in Table VIII. The structure is determined using the Hyperband tuner. However, we found that the network generally is hard to train in an unsupervised setting because the network tends to converge to local minima. Additionally, a network with low reconstruction error does not always achieve better performance than a network with a high reconstruction error, and it is thus difficult to tune the network properly using unsupervised methods. We, therefore, resorted to manual tuning.

C. Evaluation Metrics

The objective of the evaluation is to describe the method's discriminatory power, i.e., the ability to distinguish between

TABLE VIII
THE IMPLEMENTED BL AE NETWORK STRUCTURE.

Layer	Layer size	Activation function
x	(10, 24)	
$h = \text{BL}(x)$	(2, 45)	ReLU
$\hat{x} = \text{BL}(h)$	25	linear

normal data instances and anomalies. An often used metric within classification is accuracy:

$$\text{accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \quad (9)$$

where TP is the number of true positives, TN is the number of true negatives, FP is the number of false positives, and FN is the number of false negatives. However, as can be seen in eq. (9), accuracy depends on TN , and since anomaly detection tends to be an imbalanced problem where most data instances typically are normal, a high score is not necessarily indicative of high anomaly detection performance: simply classifying all data instances as normal would yield a high accuracy score.

Two often used metrics within anomaly detection for methods returning a continuous output score are the Receiver Operating Characteristic (ROC) curve [6, 17, 46, 35] and the Precision-Recall (PR) curve [6, 13, 18]. The ROC curve denotes the trade-off between the *true positive rate* (TPR), which is the percentage of anomalies correctly identified as anomalous, and the *false positive rate* (FPR), which is the percentage of normal instances wrongly classified as anomalies, by plotting the TPR against the FPR .

$$TPR = \frac{TP}{TP + FN}, \quad (10)$$

$$FPR = \frac{FP}{FP + TN}. \quad (11)$$

The ROC curve has the benefit of its Area Under Curve (AUC) being interpretable as the detector's discriminatory power [56] since ROC-AUC is interpretable as the probability of a detector assigning a higher score, e.g., greater distance or high reconstruction error, to an anomaly than a normal data instance [6]. However, ROC presents the performance as overly optimistic for imbalanced datasets, since a large change in FP can lead to a small change in FPR, when having a high number of normal instances and thereby possibly $TN \gg FP$ [57].

The *PR curve* denotes the trade-off between the *precision* and the *recall* of a detector, defined as follows:

$$\text{recall} = \frac{TP}{TP + FN}, \quad (12)$$

$$\text{precision} = \frac{TP}{FP + TP}. \quad (13)$$

The recall metric is the same as TPR , while precision is the number of identified anomalies actually being anomalous. Precision is not impacted by imbalanced data in the same sense as FPR , since it relates TP to FP instead of TN [57].

A ROC curve dominates another ROC curve, i.e., all parts of the other curve are beneath or equal, if and only if the corresponding PR curve for the method dominates as well [57].

Fig. 4. The robot workcell.

In [57], it is furthermore shown that optimizing the ROC-AUC does not necessarily optimize the PR-AUC. Based on these findings, it is appropriate to choose methods optimizing the PR-AUC, while taking into account the trade-offs reflected by the ROC and PR curves.

V. EXPERIMENTAL SETUP

To evaluate the presented methods, we design a pick-and-place application for a UR5e, because collaborative robots are commonly used for such applications [58]. The robot program consists of the main program and a rarely-visited branch case. The robot uses contact detection, i.e. moving until the TCP measure a contact, to locate the objects and to determine when to release the object during place operations. Finally, the program also includes a non-deterministic place location to represent applications in which the exact location of the objects is not known in advance.

Fig. 4 shows the robot in its home configuration and the workcell. The robot's task is to move the cylinder objects between the red squares marked on the table with the use of a Robotiq 2F-85 gripper. The four cylinders, from left to right, have a mass of 0.5 kg, 1.0 kg, 2.0 kg, and 2.5 kg, respectively. The cylinder of mass 2.0 kg is used as the default mass in our experiments.

The robot first picks up the cylinder at Square 1, and then it either moves to Square 3 or enters the rarely-visited program branch with a likelihood of 0.167, and moves to Square 2 before moving to Square 3. At Square 2, the cylinder is placed at the center of the marked square. The robot then waits for a random amount of time uniformly drawn from the interval four to ten seconds, before it picks up the cylinder. Afterwards, the robot moves to Square 3. The place location at Square 3 is randomly generated to be within the $20\text{ cm} \times 20\text{ cm}$ red

square sampled from a uniform distribution. The robot then moves the cylinder back to Square 1. A video of the robotic experiment can be found at the following link: https://youtu.be/a63tmL3yC_E.

Aside from collecting data during normal operation, we conduct experiments in which we introduce eight different anomalies, which can occur in common collaborative robot applications:

- 1) *Push*: we push the robot at either joint 3 or joint 5 (see Fig. 5a) with a force less than 50 N.
- 2) *TCP wrench*: we apply a wrench to the robot's TCP (Fig. 5c).
- 3) *Raised object*: we introduce a foreign object at the place location which makes the robot place the cylinder at a location higher than in the normal behavior (Fig. 5b).
- 4) *Outside square*: we sample a location outside the usual $20\text{ cm} \times 20\text{ cm}$ square (Fig. 5d).
- 5) *Change weight*: we change the default 2.0 kg cylinder with either the 0.5 kg, 1.0 kg, or 2.5 kg cylinder.
- 6) *Drop object*: the robot losses the cylinder (Fig. 6).
- 7) *Speed scale*: we reduce the robot's movement speed by 50%.
- 8) *Long wait*: we extend the four to ten seconds wait to either 12 s, 15 s, or 20 s.

We collected multiple datasets from the application: a normal dataset representing the normal behavior sampled from 600 program runs, and eight datasets with anomalies present – one for each of the introduced anomalies sampled over ten programs runs per anomaly, i.e., a total of 80 program runs with anomalies present. The datasets are available at [59].

The data is sampled at a frequency of 100 Hz, instead of the maximum 500 Hz, as it is sufficient to detect the classes of anomalies that we focus on in this study. Table IX summarizes the datasets.

TABLE IX
SUMMARY OF DATA SETS COLLECTED DURING NORMAL OPERATION AND WITH EXTERNAL ANOMALIES.

Dataset name	Total samples	Anomalous samples	Contamination [%]
Normal	1530921	0	0.00
Push	31117	2687	8.64
TCP wrench	31062	4660	15.00
Raised object	28438	3992	14.04
Outside box	23499	7539	32.08
Change weight	28978	9656	33.32
Drop object	22274	5105	22.92
Speed scale	32804	7449	22.71
Long wait	47979	5600	11.67

VI. EXPERIMENTAL RESULTS

Table X summarizes the discriminatory power of the 15 methods evaluated on the eight anomaly datasets individually and the three best performing methods are highlighted in bold for each anomaly. We compare the performance using the non-parametric Mann-Whitney U Test [60] by determining the degree of statistical significance of differences. We use Mann-Whitney since D'Agostino's test for normality [61] showed that some results are not normally distributed, i.e. $p < 0.5$.

Fig. 5. Illustration of the introduced anomalies.

Fig. 6. Illustration of the Drop object anomaly.

From the table we see that instance-based methods (k -NN and LOF) generally achieve a good performance despite their relative simplicity, while PCA-based reconstruction models (PCA, PPCA, and IPCA) generally have low performance; implementing PCA as a hybrid method did not increase the performance on our datasets. Additionally, we see that non-linear models outperform PCA models in nearly all aspect.

From the table we also see that implementing hybrid method (IPCA+OCSVM and VAE+GMM) lead to a slightly higher performance than computing the score based on the MSE.

By comparing the UAE and DAE, we see that the incorporating the denoising criterion into the training phase increases the performance of the UAE; with the largest increase observed for the *Push* anomaly. Additionally, by comparing UAE with the statistical AE methods (VAE, CVAE, and VAE+GMM), we see that using variational inference increase the performance for the challenging anomalies: *Raised object*, *Outside box*, and *TCP wrench*.

By comparing the time series methods, we see that the LSTM structure receives the highest score, and it ranks top three on half of the evaluated anomalies. Both DeepAnT and the bilinear network perform below average.

In order to gain insight into which aspects of robot applications complicate anomaly detection, we identified three hypotheses described below. To test the three hypotheses, we slightly modify the original pick-and-place task and collect a new training and evaluation dataset for each modification. In the new datasets, we only introduce the *Push* and *TCP wrench* anomaly and we only evaluated three anomaly detectors: LOF, VAE+GMM, and CVAE.

A. Rarely-visited Program Branches

Hypothesis 1: Distinguishing between normal and anomalous data instances in the rarely-visited program branch is significantly harder than in the rest of the program.

To test Hypothesis 1, we remove the program branch that makes the robot move to Square 2, see Section V, and we sample a new dataset, which we denote *no_branch*. The collected data is summarized in Table XI and Table XII lists the evaluation results for the methods.

We formulate the null hypothesis that the mean performance of the methods trained on the original data, μ_{orig} , is equal to the mean performance of the methods trained on the *no_branch* data, μ_{no_branch} :

$$H_0^{(1)} : \mu_{orig}^{LOF} = \mu_{no_branch}^{LOF} , \quad (14)$$

$$H_0^{(2)} : \mu_{orig}^{VAE+GMM} = \mu_{no_branch}^{VAE+GMM} , \quad (15)$$

$$H_0^{(3)} : \mu_{orig}^{CVAE} = \mu_{no_branch}^{CVAE} . \quad (16)$$

We use the Mann-Whitney U Test to test for equality of means, which we reject for all three methods, i.e., $p < 0.05$. We thus reject the null hypothesis.

We then formulate an alternative hypothesis that μ_{no_branch} is greater than μ_{orig} and test it using the one-sided Mann-Whitney U Test:

$$H_A^{(1)} : \mu_{orig}^{LOF} < \mu_{no_branch}^{LOF} , \quad (17)$$

$$H_A^{(2)} : \mu_{orig}^{VAE+GMM} < \mu_{no_branch}^{VAE+GMM} , \quad (18)$$

$$H_A^{(3)} : \mu_{orig}^{CVAE} < \mu_{no_branch}^{CVAE} . \quad (19)$$

For all methods, we also reject the alternative hypothesis, i.e., $p < 0.05$. Therefore, we conclude that the selected methods experience a significantly higher anomaly detection performance when scoring data instances in programs without rarely-visited program branches, which confirms Hypothesis 1.

B. Non-deterministic Trajectories

Hypothesis 2: Distinguishing between normal and anomalous data instances in deterministic trajectories is easier than in non-deterministic trajectories.

To test Hypothesis 2, we remove the program branch that makes the robot move to Square 2, we replace the semi-randomly generated place location a Square 3 with a fixed location, and we sample a new dataset, which we denote *deterministic*. The collected data is summarized in Table XIII and Table XIV lists the evaluation results of the anomaly detection methods.

We formulate the null hypothesis that the mean performance of the methods trained on the *no_branch* data, μ_{no_branch} ,

TABLE X

THE RESULTS ARE REPORTED AS MEAN ($\pm$ STANDARD DEVIATION) OBTAINED BY EVALUATING ALL METHODS 30 TIMES ON THE TEST DATASET (FITTING TO TRAINING DATA REPRESENTING THE ROBOT APPLICATION'S NORMAL BEHAVIOR AND EVALUATING ON THE DATASETS WITH ANOMALIES 30 TIMES). ALL STANDARD DEVIATIONS REPORTED AS ($\pm$ 0.0000) INDICATE THAT THE STANDARD DEVIATION IS BELOW 10^{-4} . THE THREE BEST PERFORMING METHODS ARE HIGHLIGHTED FOR EACH DATASET.

		All anomalies	Push	Raised object	Outside box	TCP wrench	Long wait	Drop object	Change weight	Speed scale
KNN	PR-AUC	0.7926 (± 0.0033)	0.6248 (± 0.0049)	0.3054 (± 0.0062)	0.6972 (± 0.0011)	0.7821 (± 0.0046)	0.9737 (± 0.0003)	0.9526 (± 0.0026)	0.6894 (± 0.0012)	0.8421 (± 0.0041)
	ROC-AUC	0.9428 (± 0.0014)	0.9429 (± 0.0010)	0.8029 (± 0.0061)	0.9213 (± 0.0005)	0.8930 (± 0.0020)	0.9948 (± 0.0000)	0.9860 (± 0.0006)	0.8805 (± 0.0007)	0.9586 (± 0.0011)
LOF	PR-AUC	0.8255 (± 0.0012)	0.5049 (± 0.0004)	0.6698 (± 0.0089)	0.7559 (± 0.0016)	0.8141 (± 0.0004)	0.9955 (± 0.0004)	0.9748 (± 0.0003)	0.7215 (± 0.0015)	0.7981 (± 0.0006)
	ROC-AUC	0.9531 (± 0.0002)	0.9322 (± 0.0001)	0.3422 (± 0.0010)	0.9300 (± 0.0004)	0.9518 (± 0.0001)	0.9991 (± 0.0000)	0.9891 (± 0.0001)	0.8788 (± 0.0002)	0.9413 (± 0.0004)
PCA	PR-AUC	0.3524 (± 0.0264)	0.3465 (± 0.9374)	0.1107 (± 0.0049)	0.4892 (± 0.0510)	0.4235 (± 0.0305)	0.6854 (± 0.2418)	0.3403 (± 0.0266)	0.4129 (± 0.0102)	0.2424 (± 0.0089)
	ROC-AUC	0.6470 (± 0.0146)	0.7952 (± 0.0215)	0.3422 (± 0.0140)	0.6971 (± 0.0674)	0.7397 (± 0.0251)	0.9354 (± 0.0584)	0.6595 (± 0.0322)	0.6040 (± 0.0238)	0.5565 (± 0.0186)
PPCA	PR-AUC	0.3882 (± 0.0112)	0.3319 (± 0.0282)	0.1144 (± 0.0035)	0.4838 (± 0.0591)	0.4203 (± 0.0255)	0.9806 (± 0.0113)	0.3134 (± 0.0203)	0.4320 (± 0.0077)	0.2275 (± 0.0061)
	ROC-AUC	0.6489 (± 0.0172)	0.8004 (± 0.0154)	0.3422 (± 0.0087)	0.6724 (± 0.0822)	0.7845 (± 0.0123)	0.9980 (± 0.0012)	0.6259 (± 0.0261)	0.6254 (± 0.0089)	0.5200 (± 0.0165)
IPCA	PR-AUC	0.3448 (± 0.0382)	0.2920 (± 0.0371)	0.1144 (± 0.0103)	0.5046 (± 0.0504)	0.3503 (± 0.0241)	0.7243 (± 0.2831)	0.3361 (± 0.0406)	0.3906 (± 0.0189)	0.2500 (± 0.0121)
	ROC-AUC	0.6395 (± 0.0209)	0.7630 (± 0.0278)	0.3698 (± 0.0226)	0.7361 (± 0.0690)	0.7712 (± 0.0351)	0.9445 (± 0.0599)	0.6425 (± 0.0463)	0.5647 (± 0.0270)	0.5468 (± 0.0210)
IPCA+OCSVM	PR-AUC	0.4803 (± 0.0023)	0.2528 (± 0.0024)	0.1307 (± 0.0011)	0.2279 (± 0.0012)	0.4300 (± 0.0051)	0.9984 (± 0.0004)	0.6593 (± 0.0051)	0.6415 (± 0.0031)	0.3384 (± 0.0040)
	ROC-AUC	0.6153 (± 0.0017)	0.6060 (± 0.0039)	0.3966 (± 0.0043)	0.2188 (± 0.0034)	0.7240 (± 0.0033)	0.9998 (± 0.0000)	0.7463 (± 0.0045)	0.7764 (± 0.0029)	0.4966 (± 0.0051)
UAE	PR-AUC	0.6489 (± 0.0226)	0.3607 (± 0.0278)	0.1443 (± 0.0248)	0.7640 (± 0.0607)	0.7505 (± 0.0102)	0.9901 (± 0.0105)	0.7176 (± 0.0352)	0.6613 (± 0.0635)	0.4929 (± 0.0370)
	ROC-AUC	0.8083 (± 0.0117)	0.7996 (± 0.0161)	0.5011 (± 0.0730)	0.9367 (± 0.0154)	0.8301 (± 0.0130)	0.9985 (± 0.0016)	0.8328 (± 0.0217)	0.7735 (± 0.0274)	0.7179 (± 0.0183)
DAE	PR-AUC	0.7014 (± 0.0101)	0.5246 (± 0.0362)	0.1745 (± 0.0115)	0.7169 (± 0.0121)	0.7718 (± 0.0225)	0.9995 (± 0.0004)	0.8169 (± 0.0160)	0.7646 (± 0.0193)	0.5275 (± 0.0237)
	ROC-AUC	0.8500 (± 0.0067)	0.9022 (± 0.0092)	0.5583 (± 0.0310)	0.9286 (± 0.0024)	0.8373 (± 0.0200)	0.9999 (± 0.0000)	0.9055 (± 0.0140)	0.8386 (± 0.0139)	0.7695 (± 0.0156)
VAE	PR-AUC	0.7018 (± 0.0306)	0.3732 (± 0.0456)	0.3333 (± 0.0981)	0.8910 (± 0.0571)	0.8787 (± 0.0384)	0.7143 (± 0.1755)	0.7922 (± 0.0198)	0.7697 (± 0.0375)	0.5586 (± 0.0392)
	ROC-AUC	0.8572 (± 0.0170)	0.8293 (± 0.0125)	0.6625 (± 0.0521)	0.9648 (± 0.0173)	0.9465 (± 0.0259)	0.8454 (± 0.1182)	0.8855 (± 0.0098)	0.8658 (± 0.0158)	0.8147 (± 0.0161)
VAE+GMM	PR-AUC	0.7925 (± 0.0212)	0.4160 (± 0.0619)	0.6860 (± 0.0761)	0.7704 (± 0.0522)	0.8879 (± 0.0469)	0.9626 (± 0.0768)	0.8746 (± 0.0236)	0.6836 (± 0.0436)	0.7542 (± 0.0305)
	ROC-AUC	0.9275 (± 0.0078)	0.8461 (± 0.0258)	0.9277 (± 0.0200)	0.9388 (± 0.0103)	0.9526 (± 0.0368)	0.9841 (± 0.0363)	0.9479 (± 0.0126)	0.8422 (± 0.0194)	0.9009 (± 0.0124)
CVAE	PR-AUC	0.7208 (± 0.0364)	0.4022 (± 0.0316)	0.8454 (± 0.0423)	0.8505 (± 0.0379)	0.9437 (± 0.0168)	0.6318 (± 0.2390)	0.8686 (± 0.0275)	0.6099 (± 0.0253)	0.5725 (± 0.0586)
	ROC-AUC	0.8778 (± 0.0259)	0.8405 (± 0.0103)	0.9461 (± 0.0166)	0.9499 (± 0.0092)	0.9742 (± 0.0083)	0.7409 (± 0.1830)	0.9494 (± 0.0121)	0.8353 (± 0.0149)	0.8216 (± 0.0246)
DeepAnT	PR-AUC	0.4169 (± 0.0030)	0.2157 (± 0.0076)	0.1695 (± 0.0041)	0.5593 (± 0.0183)	0.4248 (± 0.0072)	0.9950 (± 0.0001)	0.4113 (± 0.0094)	0.4143 (± 0.0062)	0.3331 (± 0.0086)
	ROC-AUC	0.7200 (± 0.0059)	0.7870 (± 0.0065)	0.4761 (± 0.0127)	0.7978 (± 0.0236)	0.8647 (± 0.0052)	0.9998 (± 0.0000)	0.7674 (± 0.0085)	0.6625 (± 0.0079)	0.6820 (± 0.0111)
LSTM AE	PR-AUC	0.7860 (± 0.0108)	0.6689 (± 0.0238)	0.3328 (± 0.0359)	0.7186 (± 0.0225)	0.7409 (± 0.0205)	0.9785 (± 0.0073)	0.8992 (± 0.0213)	0.7503 (± 0.0342)	0.7981 (± 0.0158)
	ROC-AUC	0.9192 (± 0.0049)	0.9457 (± 0.0073)	0.8049 (± 0.0294)	0.9265 (± 0.0037)	0.8603 (± 0.0189)	0.9970 (± 0.0016)	0.9524 (± 0.0111)	0.8697 (± 0.0105)	0.9163 (± 0.0066)
BL AE	PR-AUC	0.5827 (± 0.0451)	0.4674 (± 0.1135)	0.1564 (± 0.0256)	0.7391 (± 0.0430)	0.7218 (± 0.0258)	0.9538 (± 0.0301)	0.6273 (± 0.2179)	0.5701 (± 0.0806)	0.3587 (± 0.0925)
	ROC-AUC	0.7861 (± 0.0489)	0.8690 (± 0.0480)	0.4934 (± 0.1026)	0.9337 (± 0.0086)	0.8320 (± 0.0241)	0.9916 (± 0.0040)	0.8059 (± 0.1031)	0.7162 (± 0.0469)	0.6705 (± 0.0952)
RandNet	PR-AUC	0.3958 (± 0.0101)	0.3246 (± 0.0597)	0.1245 (± 0.0075)	0.6891 (± 0.0198)	0.4720 (± 0.0632)	0.9968 (± 0.0033)	0.3824 (± 0.0614)	0.4192 (± 0.0510)	0.2541 (± 0.0350)
	ROC-AUC	0.7681 (± 0.0024)	0.7602 (± 0.0411)	0.4200 (± 0.0279)	0.9057 (± 0.0139)	0.7787 (± 0.0262)	0.9996 (± 0.0004)	0.6866 (± 0.0504)	0.6164 (± 0.0466)	0.5867 (± 0.0480)

TABLE XI

NUMBER OF DATA INSTANCES SAMPLED FOR THE NORMAL CASE AND ALL ANOMALOUS PROGRAM RUNS WITHOUT THE PROGRAM BRANCH. THE NORMAL DATA IS GATHERED FROM 200 PROGRAM RUNS. THE *Push* ANOMALY IS INTRODUCED IN TEN PROGRAM RUNS AND THE *TCP wrench* ANOMALY IS INTRODUCED IN TEN PROGRAM RUNS.

Dataset name	Total samples	Anomalous samples	Contamination [%]
Normal	422510	0	0.00
Push	21233	1391	6.55
TCP wrench	21171	2803	13.23

TABLE XII

ORIGINAL RESULTS (WITH A RARELY-VISITED PROGRAM BRANCH) FROM TABLE X AND RESULTS OF THE NEW EXPERIMENTS WITHOUT THE RARELY-VISITED PROGRAM BRANCH. RESULTS ARE REPORTED AS MEAN ($\pm$ STANDARD DEVIATION) WHICH ARE OBTAINED FROM 30 EVALUATIONS OF EACH MODEL.

Original results				
		LOF	VAE+GMM	CVAE
Push	PR-AUC	0.5049 (± 0.0004)	0.4160 (± 0.0619)	0.4022 (± 0.0316)
	ROC-AUC	0.9322 (± 0.0001)	0.8461 (± 0.0258)	0.8405 (± 0.0103)
TCP wrench	PR-AUC	0.8141 (± 0.0004)	0.8879 (± 0.0469)	0.9437 (± 0.0168)
	ROC-AUC	0.9518 (± 0.0001)	0.9526 (± 0.0368)	0.9742 (± 0.0083)
Without rarely-visited program branch				
		LOF	VAE+GMM	CVAE
Push	PR-AUC	0.9366 (± 0.0007)	0.7350 (± 0.0857)	0.8067 (± 0.0228)
	ROC-AUC	0.9880 (± 0.0003)	0.9558 (± 0.0139)	0.9650 (± 0.0030)
TCP wrench	PR-AUC	0.9220 (± 0.0016)	0.9641 (± 0.0105)	0.9668 (± 0.0030)
	ROC-AUC	0.9910 (± 0.0001)	0.9899 (± 0.0018)	0.9905 (± 0.0010)

data is equal to the performance of the methods trained on the *deterministic* data, $\mu_{\text{deterministic}}$:

$$H_0^{(1)} : \mu_{\text{no_branch}}^{\text{LOF}} = \mu_{\text{deterministic}}^{\text{LOF}} \quad (20)$$

$$H_0^{(2)} : \mu_{\text{no_branch}}^{\text{VAE+GMM}} = \mu_{\text{deterministic}}^{\text{VAE+GMM}} \quad (21)$$

$$H_0^{(3)} : \mu_{\text{no_branch}}^{\text{CVAE}} = \mu_{\text{deterministic}}^{\text{CVAE}} \quad (22)$$

We use the Mann-Whitney U Test to test for equality of means, which we reject for all three methods, i.e., $p < 0.05$.

We then formulate an alternative hypothesis that $\mu_{\text{deterministic}}$ is greater than $\mu_{\text{no_branch}}$ and test it using the one-sided Mann-Whitney U Test:

$$H_A^{(1)} : \mu_{\text{no_branch}}^{\text{LOF}} < \mu_{\text{deterministic}}^{\text{LOF}} \quad (23)$$

$$H_A^{(2)} : \mu_{\text{no_branch}}^{\text{VAE+GMM}} < \mu_{\text{deterministic}}^{\text{VAE+GMM}} \quad (24)$$

$$H_A^{(3)} : \mu_{\text{no_branch}}^{\text{CVAE}} < \mu_{\text{deterministic}}^{\text{CVAE}} \quad (25)$$

We accept the alternative hypothesis for all methods and anomalies, i.e., $p < 0.05$, except for CVAE and the *Push* anomaly, i.e., $p > 0.05$. We conclude that most methods

significantly improve their anomaly detection performance when the robot program only consists of deterministic trajectories, which confirms that Hypothesis 2 holds for most of the methods evaluated.

TABLE XIII

NUMBER OF DATA INSTANCES SAMPLED FOR THE NORMAL CASE AND ALL ANOMALOUS PROGRAM RUNS WITHOUT THE PROGRAM BRANCH AND WITH A FIXED PLACE LOCATION. THE NORMAL DATA IS COLLECTED IN 200 PROGRAM RUNS. THE *Push* ANOMALY IS INTRODUCED IN TEN PROGRAM RUNS AND THE *TCP wrench* ANOMALY IS INTRODUCED IN TEN PROGRAM RUNS.

Dataset name	Total samples	Anomalous samples	Contamination [%]
Normal	420886	0	0.00
Push	21192	1401	6.61
TCP wrench	21152	2805	13.26

TABLE XIV

FIXED-PLACE-LOCATION RESULTS. RESULTS ARE REPORTED AS MEAN ($\pm$ STANDARD DEVIATION) WHICH ARE OBTAINED FROM 30 EVALUATIONS OF EACH MODEL.

		LOF	VAE+GMM	CVAE
Push	PR-AUC	0.9814 (± 0.0014)	0.8460 (± 0.0594)	0.7718 (± 0.0363)
	ROC-AUC	0.9962 (± 0.0004)	0.9716 (± 0.0080)	0.9598 (± 0.0096)
TCP wrench	PR-AUC	0.9861 (± 0.0010)	0.9847 (± 0.0045)	0.9964 (± 0.0009)
	ROC-AUC	0.9978 (± 0.0002)	0.9963 (± 0.0010)	0.9989 (± 0.0004)

C. Intentional Contact With the Environment

Hypothesis 3: Distinguishing between normal and anomalous data instances in robot programs with intentional contact with the environment is significantly harder than those without.

To test Hypothesis 3, we remove the program branch that makes the robot move to Square 2, we replace the contact-based pick-and-place operations with hard-coded locations, and we sample a new dataset, which we denote *no_contact*. The collected data is summarized in Table XV and Table XVI lists the evaluation results for the methods.

We formulate the null hypothesis that the mean performance of the methods trained on the *no_branch* data, $\mu_{\text{no_branch}}$ is equal to the performance of the methods trained on the *no_contact* data, $\mu_{\text{no_contact}}$:

$$H_0^{(1)} : \mu_{\text{no_branch}}^{\text{LOF}} = \mu_{\text{no_contact}}^{\text{LOF}} \quad (26)$$

$$H_0^{(2)} : \mu_{\text{no_branch}}^{\text{VAE+GMM}} = \mu_{\text{no_contact}}^{\text{VAE+GMM}} \quad (27)$$

$$H_0^{(3)} : \mu_{\text{no_branch}}^{\text{CVAE}} = \mu_{\text{no_contact}}^{\text{CVAE}} \quad (28)$$

We use the Mann-Whitney U Test to test for equality of means, which we reject for all three methods, i.e., $p < 0.05$.

We then formulate an alternative hypothesis that $\mu_{no_contact}$ is greater than μ_{no_branch} and test it using the one-sided Mann-Whitney U Test:

$$H_A^{(1)} : \mu_{no_branch}^{LOF} < \mu_{no_contact}^{LOF} \quad (29)$$

$$H_A^{(2)} : \mu_{no_branch}^{VAE+GMM} < \mu_{no_contact}^{VAE+GMM}, \quad (30)$$

$$H_A^{(3)} : \mu_{no_branch}^{CVAE} < \mu_{no_contact}^{CVAE} \quad (31)$$

We accept the alternative hypothesis for all three methods, i.e., $p < 0.05$. Therefore, we conclude that all methods yield a significantly higher anomaly detection performance when scoring data instances in programs without intentional contact with the environment, which confirms Hypothesis 3.

TABLE XV

NUMBER OF DATA INSTANCES SAMPLED FOR THE NORMAL CASE AND ALL ANOMALOUS PROGRAM RUNS WITHOUT THE PROGRAM BRANCH AND WITHOUT INTENTIONAL CONTACT WITH THE ENVIRONMENT. THE NORMAL DATA IS GATHERED FROM 200 PROGRAM RUNS. THE *Push* ANOMALY IS INTRODUCED IN TEN PROGRAM RUNS AND THE *TCP wrench* ANOMALY IS INTRODUCED IN TEN PROGRAM RUNS.

Dataset name	Total samples	Anomalous samples	Contamination [%]
Normal	307782	0	0.00
Push	15391	1488	9.67
TCP wrench	15383	1647	10.71

TABLE XVI

WITHOUT-INTENTIONAL-CONTACT RESULTS. THE RESULTS ARE REPORTED AS MEAN ($\pm$ STANDARD DEVIATION) WHICH ARE OBTAINED FROM 30 EVALUATIONS OF THE MODELS.

		LOF	VAE+GMM	CVAE
Push	PR-AUC	0.9575 (± 0.0024)	0.8951 (± 0.0280)	0.8825 (± 0.0183)
	ROC-AUC	0.9895 (± 0.0005)	0.9767 (± 0.0052)	0.9715 (± 0.0038)
TCP wrench	PR-AUC	0.9401 (± 0.0026)	0.9732 (± 0.0116)	0.9839 (± 0.0030)
	ROC-AUC	0.9932 (± 0.0002)	0.9936 (± 0.0018)	0.9959 (± 0.0010)

VII. DISCUSSION

In our experiments, we found that the instance-based methods k -NN and LOF achieve a remarkably high detection performance for the exogenous anomalies introduced in the pick-and-place application, particularly considering their relative simplicity. For the reconstruction-based method, we found that PCA achieves poor performance, which indicates that our data features have nonlinear dependencies. Neural network-based autoencoders, the nonlinear equivalent to PCA, on the other hand, achieve high detection performance for some of the evaluated network structures: (i) DAE, (ii) VAE+GMM, (iii) CVAE, and (iv) LSTM. In particular, structures that use variational inference achieve good performance. This observation aligns with current research [18, 17] that indicates that combining variational inference with, e.g., times series methods, yields good performance. The nonlinear explicit

generalization methods have several interesting characteristics that could be exploited in robot applications: (i) learning can be done incrementally, (ii) samples can be weighted differently when updating the model, and (iii) fast prediction times when models are kept relatively small. The studied generalization models can also easily be modified and combined with other methods to further enhance their performance.

We found three cobot program elements that can significantly decrease the performance of anomaly detection methods: (i) when the robot program contains rarely-visited program branches, (ii) when the program consists of non-deterministic trajectories, and (iii) when the robot's TCP intentionally makes contact with the environment. It would be interesting to explore which detectors are the least vulnerable to these situations, e.g., it might be possible to mitigate the issue of rarely-visited program branches by considering the imbalanced representation of the program nodes when training the AE structures.

The results of our study indicate that exploiting the robot program context directly in the CVAE structure generally leads to significantly better performance than for the (non-conditional) VAE structure. However, indirect exploitation as in the VAE+GMM model also yielded very promising performance. It would be interesting to further explore the incorporation contextual information from the robot program when performing anomaly detection in robot manipulators.

VIII. CONCLUSIONS

We evaluated 16 anomaly detection methods in a common pick-and-place application for a robot manipulator. We trained models on data collected under normal operation and assessed their discriminatory power on eight different anomalies. The two instance-based methods, namely k -NN and LOF, achieved good performance despite their relative simplicity compared to nonlinear explicit generalization models based on autoencoders. However, the performance profile of the evaluated methods varied significantly across the different types of anomalies. To achieve the best performance, the choice of method should thus depend on the expected likelihood and severity of different types of anomalies in a given environment. Alternatively, two (e.g. k -NN and VAE+GMM) or more methods could potentially be combined in an ensemble to achieve better overall performance, but at the cost of increased complexity.

REFERENCES

- [1] Y. Cohen, S. Shoval, and M. Faccio, "Strategic view on cobot deployment in assembly 4.0 systems," *IFAC-PapersOnLine*, vol. 52, no. 13, pp. 1519–1524, 2019.
- [2] S. El Zaatari, M. Marei, W. Li, and Z. Usman, "Cobot programming for collaborative industrial tasks: An overview," *Robotics and Autonomous Systems*, vol. 116, pp. 162–180, 2019.
- [3] P. Zhao, M. Kurihara, J. Tanaka, T. Noda, S. Chikuma, and T. Suzuki, "Advanced correlation-based anomaly detection method for predictive maintenance," in *2017 IEEE International Conference on Prognostics and Health Management (ICPHM)*. IEEE, 2017, pp. 78–83.
- [4] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Computing Surveys (CSUR)*, vol. 41, no. 3, pp. 1–58, 2009.

- [5] T. Borgi, A. Hidri, B. Neef, and M. S. Naceur, "Data analytics for predictive maintenance of industrial robots," in *2017 International Conference on Advanced Systems and Electric Technologies (IC_ASET)*. IEEE, 2017, pp. 412–417.
- [6] C. C. Aggarwal, *Outlier analysis*. Cham: Springer International Publishing, 2017. [Online]. Available: <http://link.springer.com/10.1007/978-3-319-47578-3>
- [7] E. Khalastchi and M. Kalech, "On fault detection and diagnosis in robotic systems," *ACM Computing Surveys (CSUR)*, vol. 51, no. 1, pp. 1–24, 2018.
- [8] A. De Luca and R. Mattone, "An identification scheme for robot actuator faults," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2005, pp. 1127–1131.
- [9] C. Fantuzzi, C. Secchi, and A. Visioli, "On the fault detection and isolation of industrial robot manipulators," *IFAC Proceedings Volumes*, vol. 36, no. 17, pp. 399–404, 2003.
- [10] V. Sathish, S. Ramaswamy, and S. Butail, "Training data selection criteria for detecting failures in industrial robots," *IFAC-PapersOnLine*, vol. 49, no. 1, pp. 385–390, 2016.
- [11] A. Goncharov, A. Savelev, N. Krinitsyn, and S. Mikhalevich, "Automated anomalies detection in the work of industrial robots," in *IOP Conference Series: Materials Science and Engineering*, vol. 1019, no. 1. IOP Publishing, 2021, p. 012095.
- [12] A. H. Sabry, F. H. Nordin, A. H. Sabry, and M. Z. A. Ab Kadir, "Fault detection and diagnosis of industrial robot based on power consumption modeling," *IEEE Transactions on Industrial Electronics*, vol. 67, no. 9, pp. 7929–7940, 2019.
- [13] H. Wu, Y. Guan, and J. Rojas, "A latent state-based multimodal execution monitor with anomaly detection and classification for robot introspection," *Applied Sciences*, vol. 9, no. 6, p. 1072, 2019.
- [14] J. V. Gael and Z. Ghahramani, *Nonparametric hidden Markov models*. Cambridge University Press, 2011, p. 317–340.
- [15] L. Basora, X. Olive, and T. Dubot, "Recent advances in anomaly detection methods applied to aviation," *Aerospace*, vol. 6, no. 11, p. 117, 2019.
- [16] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1798–1828, 2013.
- [17] D. Park, Y. Hoshi, and C. C. Kemp, "A multimodal anomaly detector for robot-assisted feeding using an lstm-based variational autoencoder," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1544–1551, 2018.
- [18] T. Chen, X. Liu, B. Xia, W. Wang, and Y. Lai, "Unsupervised anomaly detection of industrial robots using sliding-window convolutional variational autoencoder," *IEEE Access*, vol. 8, pp. 47072–47081, 2020.
- [19] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [20] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, "Variational inference: A review for statisticians," *Journal of the American statistical Association*, vol. 112, no. 518, pp. 859–877, 2017.
- [21] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [22] Y. Zhou, D. Arpit, I. Nwogu, and V. Govindaraju, "Is joint training better for deep auto-encoders?" *arXiv preprint arXiv:1405.1380*, 2014.
- [23] D. Im Im, S. Ahn, R. Memisevic, and Y. Bengio, "Denoising criterion for variational auto-encoding framework," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, no. 1, 2017.
- [24] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [25] R. Chalapathy and S. Chawla, "Deep learning for anomaly detection: A survey," *arXiv preprint arXiv:1901.03407*, 2019.
- [26] A. Blazquez-Garcia, A. Conde, U. Mori, and J. A. Lozano, "A review on outlier/anomaly detection in time series data," *arXiv preprint arXiv:2002.04236*, 2020.
- [27] *UR5e technical details*, Universal Robots, 2018. [Online]. Available: https://www.universal-robots.com/media/1802778/ur5e-32528_ur_technical_details_.pdf
- [28] Universal Robots, "Real-time data exchange (rtde) guide," <https://www.universal-robots.com/articles/ur/real-time-data-exchange-rtde-guide/>, 2021, accessed: 09-09-2020. [Online]. Available: <https://www.universal-robots.com/articles/ur/real-time-data-exchange-rtde-guide/>
- [29] E. Khalastchi, M. Kalech, G. A. Kaminka, and R. Lin, "Online data-driven anomaly detection in autonomous robots," *Knowledge and Information Systems*, vol. 43, no. 3, pp. 657–688, 2015.
- [30] M. Goldstein and S. Uchida, "A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data," *PLoS one*, vol. 11, no. 4, p. e0152173, 2016.
- [31] J. L. Bentley, "Multidimensional binary search trees used for associative searching," *Communications of the ACM*, vol. 18, no. 9, pp. 509–517, 1975.
- [32] R. A. Brown, "Building a balanced kd tree in $O(kn \log n)$ time," *arXiv preprint arXiv:1410.5420*, 2014.
- [33] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "Lof: Identifying density-based local outliers," in *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, 2000, pp. 93–104.
- [34] C. C. Aggarwal and S. Sathe, *Outlier ensembles: An introduction*. Springer, 2017.
- [35] J. Chen, S. Sathe, C. Aggarwal, and D. Turaga, "Outlier detection with autoencoder ensembles," in *Proceedings of the 2017 SIAM International Conference on Data Mining*. SIAM, 2017, pp. 90–98.
- [36] H.-P. Kriegel, P. Kroger, E. Schubert, and A. Zimek, "Interpreting and unifying outlier scores," in *Proceedings of the 2011 SIAM International Conference on Data Mining*. SIAM, 2011, pp. 13–24.
- [37] N. Halko, P.-G. Martinsson, and J. A. Tropp, "Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions," *SIAM Review*, vol. 53, no. 2, pp. 217–288, 2011.
- [38] M. E. Tipping and C. M. Bishop, "Mixtures of probabilistic principal component analyzers," *Neural Computation*, vol. 11, no. 2, pp. 443–482, 1999.
- [39] D. A. Ross, J. Lim, R.-S. Lin, and M.-H. Yang, "Incremental learning for robust visual tracking," *International Journal of Computer Vision*, vol. 77, no. 1-3, pp. 125–141, 2008.
- [40] M. Claesen and B. De Moor, "Hyperparameter search in machine learning," *arXiv preprint arXiv:1502.02127*, 2015.
- [41] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, "Hyperband: A novel bandit-based approach to type-r parameter optimization," *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 6765–6816, 2017.
- [42] T. O'Malley, E. Bursztin, J. Long, F. Chollet, H. Jin, L. Invernizzi et al., "Kerastuner," <https://github.com/keras-team/keras-tuner>, 2019.
- [43] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [44] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg et al., "Scikit-learn: Machine learning in python," *the Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [45] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard et al., "Tensorflow: A system for large-scale machine learning," in *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, 2016, pp. 265–283.
- [46] A. A. Pol, V. Berger, C. Germain, G. Cerminara, and M. Pierini, "Anomaly detection with conditional variational autoencoders," in *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*. IEEE, 2019, pp. 1651–

- 1657.
- [47] C. Doersch, "Tutorial on variational autoencoders," *arXiv preprint arXiv:1606.05908*, 2016.
 - [48] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of Machine Learning Research*, vol. 9, no. 11, 2008.
 - [49] C. M. Bishop, *Pattern recognition and machine learning*. Springer, 2006.
 - [50] J. T. Hancock and T. M. Khoshgoftaar, "Survey on categorical data for neural networks," *Journal of Big Data*, vol. 7, pp. 1–41, 2020.
 - [51] K. Hundman, V. Constantinou, C. Laporte, I. Colwell, and T. Soderstrom, "Detecting spacecraft anomalies using lstms and nonparametric dynamic thresholding," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018, pp. 387–395.
 - [52] M. Munir, S. A. Siddiqui, A. Dengel, and S. Ahmed, "Deepant: A deep learning approach for unsupervised anomaly detection in time series," *IEEE Access*, vol. 7, pp. 1991–2005, 2018.
 - [53] D. T. Tran, M. Magris, J. Kannianen, M. Gabbouj, and A. Iosifidis, "Tensor representation in high-frequency financial data for price change prediction," in *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 2017, pp. 1–7.
 - [54] D. T. Tran, A. Iosifidis, J. Kannianen, and M. Gabbouj, "Temporal attention-augmented bilinear network for financial time-series data analysis," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 5, pp. 1407–1418, 2018.
 - [55] X. Teng, Y.-R. Lin, and X. Wen, "Anomaly detection in dynamic networks using multi-view time-series hypersphere learning," in *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, 2017, pp. 827–836.
 - [56] J. Fan, S. Upadhye, and A. Worster, "Understanding receiver operating characteristic (roc) curves," *Canadian Journal of Emergency Medicine*, vol. 8, no. 1, pp. 19–20, 2006.
 - [57] J. Davis and M. Goadrich, "The relationship between precision-recall and roc curves," in *Proceedings of the 23rd International Conference on Machine Learning*, 2006, pp. 233–240.
 - [58] M. Stenmark, M. Haage, and E. A. Topp, "Simplified programming of re-usable skills on a safe industrial robot: Prototype and evaluation," in *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, 2017, pp. 463–472.
 - [59] S. Graabæk, E. V. Ancke, A. R. Fugl, and A. L. Christensen, "Dataset for: An experimental comparison of anomaly detection methods for collaborative robot manipulators," 2022, URL and DOI (10.5281/zenodo.5849300) will be available in the final version. [Online]. Available: IEEE
 - [60] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *The Annals of Mathematical Statistics*, pp. 50–60, 1947.
 - [61] R. B. d'Agostino, "An omnibus test of normality for moderate and large size samples," *Biometrika*, vol. 58, no. 2, pp. 341–348, 1971.