

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2023.0322000

Adaptive Visual Quality Inspection Based on Defect Prediction from Production Parameters

ZVEZDAN LONČAREVIĆ¹, SIMON REBERŠEK¹, SAMO ŠELA², JURE SKVARČ², ALEŠ UDE¹,
(Member, IEEE), and Andrej Gams¹, (Member, IEEE)

¹Department of Automatics, Biocybernetics and Robotics, Jožef Stefan Institute, Jamova cesta 39, 1000 Ljubljana, Slovenia (e-mail: author@ijs.si)

²SICK d.o.o, Cesta v Gorice 34b, 1000 Ljubljana, Slovenia (e-mail: author@sick.si)

Corresponding author: Zvezdan Lončarević (e-mail: zvezdan.loncarevic@ijs.si).

This work has received funding from the DIGITOP project funded by Ministry of Higher Education, Science and Innovation of Slovenia, Slovenian Research and Innovation Agency and European Union – NextGenerationEU, and from the program group P2-0076 Automation, robotics and biocybernetics supported by the Slovenian Research Agency.

ABSTRACT At the end of a production process, the manufactured products must usually be visually inspected to ensure their quality. Often, it is necessary to inspect the final product from several viewpoints. However, the inspection of all possible aspects might take too long and thus create a bottleneck in the production process. In this paper we propose and evaluate a methodology for adaptive, robot-aided visual quality inspection. With the proposed method, the most probable defects are first predicted based on the production process parameters. A suitable classifier for defect prediction is learnt in an unsupervised manner from a database that includes the produced parts and the associated parameters. A robot then steers the camera only towards viewpoints associated with predicted defects, which implies that the trajectories of robot motion for the inspection might be different for every product. To enable dynamic planning of camera trajectories, we describe a methodology for evaluation and selection of the most appropriate autonomous motion planner. The proposed defect prediction approach was compared to other methods and evaluated on the products from a real-world production line for injection moulding, which was implemented for a producer of parts in the automotive industry.

INDEX TERMS industrial informatics, injection moulding, machine learning, production parameters, quality inspection, robot motion planning

I. INTRODUCTION

THE required increase in the quality of products, services and processes of the Industry 4.0 paradigms [1] gave rise to the Zero Defects Manufacturing (ZDM) approach [2]. *Per se*, ZDM does not explicitly rely on defect and fault detection, but rather on defect and fault prediction and provision of suggestions on how those can be avoided [3]. However, quality inspection remains an integral step of production processes, sometimes in the form of selective inspections based on the analysis of impact of the action on the overall economic, production logistics and quality performance [4].

In a multi-stage production process, the quality of produced parts is often checked during or after every stage [5]. For example, in production of injection-moulded parts that have added motor windings, which is also the main use-case for this paper, the product gets inspected after the insertion of side windings, after injection moulding, after cooling, after additional machining of parts, etc. In this example, each stage of the production is followed by an intermediate quality in-

spection process, based on different physical quantities. However, even if the part is deemed good after each intermediate step, the final part, when inspected by a human worker, is sometimes still classified as bad, because defects such as, e.g., porous material, are difficult to detect [6]. Thus, one of the solutions to prevent faulty products, besides sampling of products and inspecting them manually, is to visually inspect the parts also in the final stage of the production [7].

Visual inspection of all possible aspects of a complex product, however, requires the observation of the part from many different viewpoints. As stated in [8], this process can be automated in two different manners. The first option is to install several cameras so that all the required aspects of the object are observable at the same time. However, the cameras, especially if there are a lot of them, could get in the way of the production process or other means of quality control. In our practical example with an automotive parts producer, inspection with dial indicators was already integrated into the production line and prevented installation of additional

cameras.

The second option is to move the camera around the object with a robot using an in-hand camera to sequentially collect the required images [9]. The caveat of the solution with the moving camera is that the process might take more time and thus generate a bottleneck in the production process. A conceptually similar solution of moving the object around the camera is also a possibility [8], but often not a viable one because the object is partially occluded when grasped by the robot.

Machine learning approaches are increasingly being applied to improve the reliability and robustness of quality control and even production processes themselves [10], including injection moulding [11]. They can detect anomalies that are difficult or even impossible to identify for a production-line engineer. In this paper, we advance the application of machine learning approaches to deal with such challenges.

A. PROBLEM FORMULATION

The problem can be formulated as follows. A visual quality inspection with an in-hand camera needs to be added to an existing production line, and it should not create a bottleneck. This implies that only the most likely defective aspects of the object should be inspected. However, such aspects need to be first determined, and later the robot needs to move the camera so that they can be observed. This further implies that the motion of the robot might be different for every workpiece, and thus the motion of the camera needs to be generated dynamically as a part of the inspection cycle.

B. CONTRIBUTION

The contribution of this paper is a methodology for adaptive robot-aided visual quality inspection, where not all aspects of the product but only the most likely predicted areas with defects are inspected. To achieve this goal, we predict the possible defects using a database of the available production parameters and perform visual quality inspection only from viewpoints that provide clear images of the possible defects.

The methodology for the prediction of defects is based on a database of production parameters. Starting with the analysis of several classification methods, we propose an approach based on clustering of production parameters projected into a latent space of a deep autoencoder, which is trained with the parameters of all available good and faulty parts. The latent space projections of faulty parts are clustered based on different types of defects, which can be checked from specific viewpoints. For each new part, the system checks the posterior probability of its latent space projection belonging to the clusters, and the robot then steers the camera to the appropriate viewpoint with sufficiently large posterior probabilities in descending order to check the area where the predicted defect is located.

Because the camera path is not determined in advance but dynamically defined by the sequence of potential viewpoints, the robot motion cannot be pre-programmed. We therefore rely on a motion planner to generate the required motion tra-

jectories. While many robot motion planners exist, we outline and evaluate a methodology to test and select the best motion planner using the planner arena motion planning evaluation tool [12]. In our experiments, motion planners available in the Open Motion Planning Library (OMPL) [13] were evaluated.

All aspects of the proposed adaptive visual quality approach, the prediction of defects and the planning of robot camera steering trajectories were evaluated on data from a real-world injection-moulding production of a parts producer in the automotive industry. A simulated image of the experimental environment and its physical twin are depicted in Fig. 1.

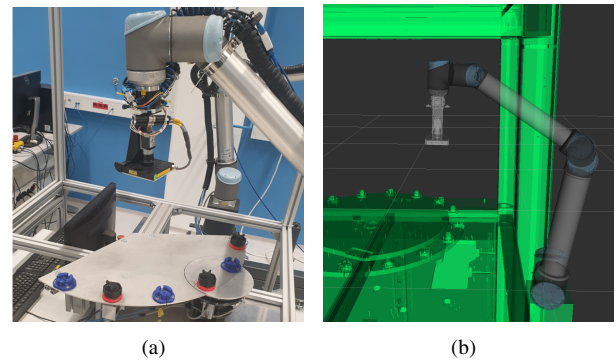


FIGURE 1. Physical implementation of the visual quality inspection cell (a) and its visualization in RViz (b)

II. RELATED WORK

This paper deals with the application of machine learning methods for quality prediction, robot-aided visual quality inspection, and robot motion planning. The related works in each of these areas are discussed separately.

A. MACHINE LEARNING METHODS FOR QUALITY PREDICTION

Machine learning methods have been applied in different manners for quality prediction [10]. In [14], the authors show that four examined machine learning (ML) algorithms adequately predict quality in injection moulding even with very little training data. Pressure data from the mould cavity was used for quality prediction. A similar analysis was performed in [15], where tree-based algorithms, regression based algorithms, and autoencoder were compared. Using a plethora of collected data, the paper shows that autoencoder outperforms other examined tree-based machine learning algorithms in accuracy, precision, recall, and F1-score. Artificial neural networks (ANN) were compared to decision trees in [16] for prediction of parts quality in thermoplastics injection moulding, showing over 99% accuracy rates. On the other hand, Silva et al. [17] show that only a combination of ANN and support vector machines (SVM) allowed them to reach accuracy above 99%. That work was extended to identify the relevant parameters [18]. Since many injection moulding machines are not prepared to access their data on a real-time basis, digitalization of such devices was explored in [19].

Given that in the real-world the datasets are often uneven or imbalanced, different methods of dealing with such datasets have been proposed. [20]. Supervised methods are known to need large quantities of data [21]. Undersampling of the majority class [22] or oversampling of the minority class with repeated or synthetic data are often used [23]. Both can lead to overfitting or loss of characteristics. On the other hand, unsupervised classification methods can be applied for uneven databases [24]. These include, among others, clustering and autoencoders. The latter can be used to detect difference through the reconstruction error [25]. Clustering and deep autoencoders can be combined with deep embedded clustering (DEC) [26].

In this paper we compare several supervised classifiers and evaluate our proposed unsupervised approach based on clustering in the latent space of a deep AE.

B. ROBOT-AIDED VISUAL QUALITY INSPECTION

Robot-aided visual quality inspection has been applied in different manners, either by moving the camera around the object, or by moving the object to be inspected in front of the camera [8]. In both cases, achieving a proper spatial relation between the camera and the object is critical for obtaining high quality images for subsequent processing and detection of defects, for example by using advanced deep learning methods [27], [28]. Thus, camera location, [29], [30] including a robot-supported autofocus mechanism, [31] and camera motion [9] are often the subject of optimization for quality inspection. Even so, the space in production lines is usually rather confined due to the production and inspection equipment and the robot may not be able to achieve optimal viewpoints for all the aspects of the object that need to be inspected. Thus, robot in-hand camera motion is sometimes combined with a rotation of the object to be inspected [8]. In this paper we use an in-hand camera that is moved around the product, but the motion of the robot is not predefined – it is planned on-line using a motion planner.

C. ROBOT MOTION PLANNING

In recent years, considerable progress was made to realize fast and robust algorithms for robot motion planning [32]. Their goal is to enable robots to automatically compute their motions from descriptors of tasks and models acquired from sensors [33]. As per the definition, the task of a path-planner is to “find a collision-free motion between an initial (start) and a final configuration (goal) within a specified environment” [34]. Various software packages offer readily available planning algorithms, one notable example being MoveIt! [35]. There exist several planning libraries that offer different types of planners. For instance, the Open Motion Planning Library (OMPL) [13] provides geometric and control-based planners and also includes algorithms for constrained planning [36]. In addition, alternatives such as Stochastic Trajectory Optimization for Motion Planning (STOMP) [37], Gradient Optimization Techniques for Efficient Motion Planning (CHOMP)

[38], and Search-based Planning Library (SBPL) [39] are available.

Tools are available within the Robot Operating System (ROS) framework to evaluate the performance of different motion planners. This includes the aforementioned MoveIt!, RViz for visualization, and Planner arena [12] for benchmarking of different planners. Open Motion Planning Library (OMPL) [13] is an integral part of available motion planning tools. We used it in the context of visual quality inspection to find planners that can generate collision-free motion trajectories in a fast and reliable manner when steering the camera to the desired viewpoints.

III. METHODOLOGY

In this section we provide the methodology of the proposed defect prediction algorithm.

The proposed defect prediction approach is based on the assumptions that 1. similar production parameters produce a similar output and 2. there is a subset of most common defects, which cover the large majority of all cases. Therefore, we use the production parameters from previously produced parts to predict the outcome of the production process for each new produced part, i.e., we compute the most probable defects based on its production parameters.

As discussed in Section II, there exist different classification approaches. The most important factors in deciding which of these classifiers can be applied effectively are the type and amount of production parameters gathered in the training database. In this section, we outline our unsupervised approach that relies on clustering of faulty parts in the latent space of a deep autoencoder (AE), which is trained on all available production data. In Section IV we show why other methods, such as for example a random forest classifier [40] or a deep neural network classifier with or without a softmax layer, cannot be effectively applied in our case.

In the proposed approach, we first train a deep AE using the complete database of production parameters. We then project the production parameters of parts onto a low dimensional latent space of a deep AE. The autoencoder passes its input data to the output so that the output matches the input with the highest possible precision. The data are pushed through differently sized layers, including through the lowest-dimensional layer, known as bottleneck or latent space. The projected data in the latent space now encodes the data with a small number of parameters. To discover the set of most common defects, the production parameters of available faulty parts are projected onto this latent space and clustered using GMM clustering. It is important to note that the real-world data is used to match the clusters and the defects. This process is illustrated in Fig. 2. Each defect type is associated with a camera viewpoint or a set of viewpoints from where its occurrence can be inspected.

At production time, the production parameters of each new part are projected onto the AE latent space. The posterior probabilities of Gaussian mixture components are then computed to predict the most probable defects. They determine

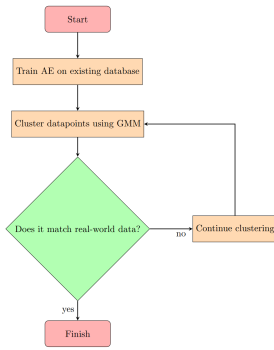


FIGURE 2. Diagram of data clustering.

the sequence of viewpoints, where the robot should place the camera to check for the most probable defects, as detailed in Algorithm 1. Note that only potential defects where the posterior probability is larger than a cut-off posterior probability p_c , determined empirically based on the producer's requirements, are checked. Thus, time is saved by performing less quality checks, testing only the defects that are likely to occur.

Algorithm 1 Procedure for inspecting the viewpoints of one product based on the process parameters

```

procedure Check viewpoints
  project new part  $\theta$  (Table 1) into the latent space  $\theta^{AE}$ 
    of a previously trained AE;
  compute posterior probabilities  $\mathbf{P}$  of the established  $c$ 
    GMM clusters associated to  $c$  most common defects;
  arrange  $\mathbf{P}$  in descending order;
  set  $i = 1$ 
  for probabilities  $p \in \mathbf{P}$ :
    if  $p > p_c$ :
      choose viewpoint(s)  $\mathbf{V}_i$  corresponding to  $p_i$ ;
      plan trajectory  $\tau$  between current robot posture
        and  $\mathbf{V}_i$ ;
      execute trajectory  $\tau$  and set  $i = i + 1$ ;
      if vision check (Fig. 11) returns error:
        discard product;
        break
  end
  
```

IV. EVALUATION OF CLASSIFICATION APPROACHES

In the following section, we introduce a functional database containing production parameters derived from the injection-moulding process used in manufacturing car parts. We then discuss and analyze design choices and the effectiveness of the proposed approach by utilizing the data obtained from this production process.

A. PRODUCTION PARAMETERS IN INJECTION MOULDING

During the production process, hundreds of production parameters are stored for every product (part). However, a large portion of the parameters is not relevant for the process

itself (server name, user name, folder, ...). In collaboration with a production engineer working at the manufacturer, we identified 22 relevant production parameters. They are listed in Table 1. Some can easily be understood, while the others, such as PD_DIA average and minimum values, were proposed by the engineer. Normalized data from the production production process are, without identifiers, made available at https://github.com/abr-ijis/production_parameters.

TABLE 1. Injection Moulding Parameters θ Used for Prediction of Defects

cycle time	switch after volume
injection time	machine cycle counter
max injection pressure	feeder housing temperature
switch after pressure	tool 1 heating circuit
dozing time	tool 2 heating circuit
material cushion	tool 4 heating circuit
tool 7 heating circuit	blasting nest
tool 10 heating circuit	nest insertion
tool 11 heating circuit	nest withdraw
cooling nest	PD_DIA minimum value
marking nest	PD_DIA average value

B. DATABASE

For evaluation we gathered a database that consists of production parameters for 4436 good (OK) parts and 45 faulty (NOK) parts. The production parameters for good parts were collected in 5 days. We also had 32 faulty physical parts with associated parameters available. The faulty parts were collected by workers at the production line by manually inspecting samples of products over the course of three months. We refer to this complete database as DB1.

The main reason for many more entries for good parts is that a typical industrial production process produces significantly more good than bad parts. In addition, the parts that reach the end of the production line have already passed intermediate checks, and thus the number of bad parts is relatively small [10]. In the use-case of this paper, the parts that are deemed faulty at the intermediate checks are discarded and their production data is not stored or is incomplete. In any case, the data about faulty parts identified at the end of the production line can initially only be collected by assigning a human worker with sufficient knowledge [41] to manually inspect the parts. Thus in general, the data for faulty parts is difficult to obtain.

To obtain the production parameters for each part, their QR code identifier is scanned and used to extract the data from the complete production line database. In our practical example, the discrepancy between the number of faulty physical parts and the number of data for faulty parts in database DB1 is due to some parts having the (small) QR code identifier not readable.

C. DEFECT PROBABILITY PREDICTION USING DEEP AE AND CLUSTERING

As explained in Algorithm 1, we pass all the parameters gathered in DB1 through a deep autoencoder (AE). The projected data is then normalized for each parameter. The autoencoder




FIGURE 3. Mean Squared Error (MSE) of samples after pushing the data through AEs with different latent space dimensions.

was designed empirically, with a 22 dimensional input and output layers and 5 hidden layers of the size 15, 10, x_{LS} , 10, 15, respectively. The hyperbolic tangent sigmoid activation function was used for the neurons and mean squared error (MSE) for the cost function. We first checked the size of the latent space that still preserves most of the information. As shown in Fig. 3, the mean squared error remains consistent after $x_{LS}=3$. An illustrative example projections of parameters of parts onto the latent space of a deep AE for DB1 is depicted in Fig. 4.

From the manufacturer we know that there are 5 typical errors that occur with the considered parts. These include porous material, faulty/broken sides, incomplete bottom edge, fault/broken pins and faulty windings as a result of the injection moulding process. The defects are depicted in Fig. 5. We therefore partitioned the data into $c = 5$ clusters. In order to exclude the effect of the selected clustering method on the result of clustering, we tested two methods; k-means and Gaussian Mixture Model (GMM) clustering, with 5 target clusters for each. The former, which is the most commonly used clustering approach, forms spheres around the centers of the clusters and performs hard classification, while the latter can handle non-round shapes of clusters and perform soft classification [42]. To check if the clustering depends on the autoencoder, we also trained the deep AE with different random initial parameters of the deep AE network. The results of clustering using both methods in latent spaces of 6 different AEs are for illustration shown in Fig. 7. Clusters are marked




FIGURE 4. Database DB1 datapoints projected into the 3-dimensional latent space of a deep AE with red dots marking datapoints for faulty parts and smaller black dots for good parts.




FIGURE 5. Five typical defects: (a) porous material – e1, (b) faulty/broken sides – e5, (c) incomplete bottom edge –e2, (d) faulty/broken pins – e3, and (e) faulty windings as a result of the injection moulding process –e4.

with different colors. Except for one object, both algorithms clustered the data in the same manner in all shown cases. This shows that training of deep AE with different initial parameters will change the latent space but not necessarily the clustering. Due to its soft classification outcome, we utilized the GMM clustering method in our methodology.

Thus, for any new data point, we check to which cluster it belongs using the posterior probability for all GMM components in the latent space. Results for the posterior probability of all DB1 data points are shown in Fig. 6.

Note, that clustering starts with random initial seeds. As depicted in Fig. 2, to use the clusters for classification of real data, they have to match. By associating the latent space data points with the physical parts with known defects, shown in Fig. 8, we can see that only one one part fits into two clusters (green and purple, shown in Fig. 8 in the 5-th row left). This result confirms that the clustering procedure is appropriate.




FIGURE 6. Probability of defect determined by posterior distribution of GMM components. The colors of the highest defect probability are associated with defects in Fig. 5.

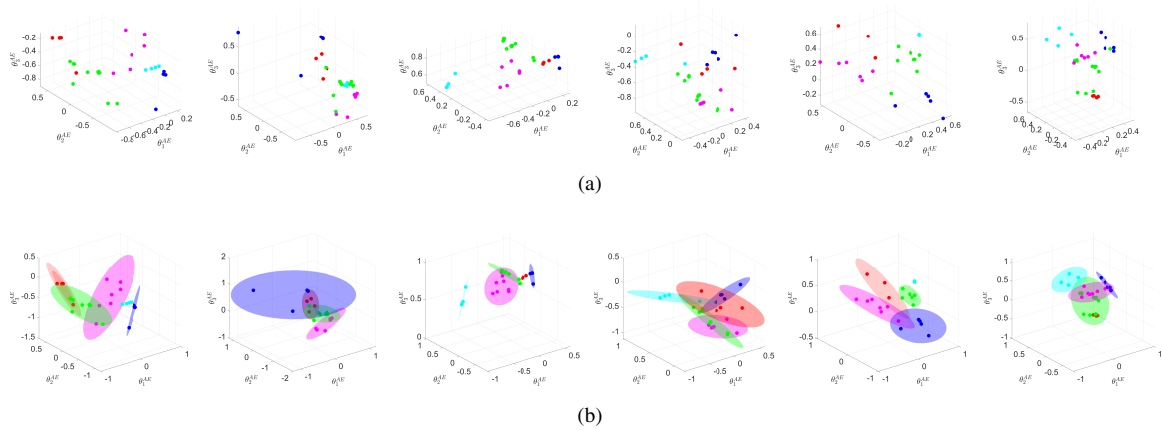


FIGURE 7. Clustering of faulty datapoints in the latent space of six autoencoders trained with different randomly selected initial autoencoder parameters. (a): k-means, (b): GMM clustering. In all cases, the clustering and color coding is the same as in Fig. 5.



FIGURE 8. The defects on physical parts coincide with the clustering of the associated production parameters in the latent space of the autoencoder, highlighted with the same colors in Fig. 7.

D. COMPARISON OF CLASSIFICATION WITH DIFFERENT CLASSIFIERS

As stated in Section II, there exist various classifiers that can be used to classify the data. In this section we compare the classification results of our proposed approach based on clustering in the latent space of the AE with three different supervised classifier methods on the parts with defects from

DB1: simple neural network classification, softmax neural network classification and random forest classification. For each of these three methods, we used 70% of the data points of parts with defects for training, 15% for validation and 15% for testing. We additionally tested another unsupervised method, i.e., classification based on AE reconstruction error.

Deep neural network classification (DNN): We defined a deep neural network with the input and output normalized in range $[-1,1]$, with the sizes of hidden layers $[15\ 10\ 6\ 3]$. When performing the forward pass, the denormalized output was rounded to the closest number in range $[1, 5]$, with the rounded output corresponding to the identity of the defect.

Softmax deep neural network classification (DNN+Softmax): We used the same deep neural network but added a softmax layer at the end.

Random Forest Classification (RF): We used 50 estimators that describe how many trees are trained. The minimal number of leaf node observations was set to 1.

Autoencoder Reconstruction Error (AE-re): The error of AE reconstruction can be used, e.g., see [25]. For this case, we trained the autoencoder on the good parts and then checked the reconstruction of good and bad parts, when passed through the AE. We clustered the reconstruction error to check if it can match with the real data.

GMM clustering in the latent space of an AE (AE LS GMM) We clustered the projection in the latent space to match with the real data We also tested whether clustering the data points in the latent space of an AE is consistent. To do this, we first obtained the clusters that match the real-world data and then re-clustered the data by changing the initial center of the clusters.

Note, as mentioned above, the choice of classifier highly depends on the available database and potentially, with a more even database, supervised classifier would do much better. Please see also the discussion in Section VII.

Results in Table 2 show that, as expected, the three supervised classifiers do not perform well with a low number of data samples because each data point that is left out for vali-

TABLE 2. Average classification and clustering results for 100 attempts using different classifiers and the proposed AE LS GMM clustering algorithm.

Classifier	DNN	DNN+Softmax	RF
Accuracy (%)	0.29167	0.353	0.579
Classifier	AE-re	AE LS GMM	
Accuracy (%)	0.5545	0.969	

dation makes a large difference. Using all the data, however, would lead to over-fitting. Unsupervised anomaly detection performed better. AR-re could reliably predict that there was a defect for 78.44% cases, but could not be used to classify the defect, as the defect classification was only 55.45% accurate. GMM clustering in the latent space of an AE (AE LS GMM), on the other hand, led to better classification rates, where 31 out of 32 items are correctly classified (96.9%), and was more consistent with the real-world data, as changing the cluster centers still matched the data in 85% of the time.

Our approach proceeds with checking the parts in the descending order of probabilities of possible defects. This means that the motion of the robot is potentially different for each part that needs to be inspected. We therefore generate the robot motions that place the camera at the required viewpoints online using a suitable motion planner.

V. ADAPTIVE VISUAL QUALITY INSPECTION

The production example presented in the introduction calls for the inspection of parts from 10 different viewpoints, even though there are only 5 most common defects. This is because some of the defects have to be inspected from more than one viewpoint (e.g. pins, porous material). In addition, there is the starting point of the robot. Thus, the robot potentially needs to move between 11 different postures. The physical implementation and simulation of the real visual quality inspection cell are depicted in Fig. 1.

A. INSPECTION VIEWPOINTS

The task of the production engineer is to associate the potential defects with the predefined camera viewpoints (V). The viewpoints must be selected in such a way that the critical surface area of the part is visible and that the acquired images are sharp. Additionally, the engineer should make sure that the associated robot postures are reachable in the confined space of the workcell. Some of the robot postures that provide good viewpoints of the observed part are displayed in Fig. 9. Because the part is small, some of the postures are quite close together. In our practical example, the engineer selected two viewpoints from the side of the object, one for the sides and windings and one for the bottom of the product. The other defects were observed from the top.

The issue with adaptive visual quality inspection is that potentially there are a lot of different motion trajectories between viewpoints. If we consider the trajectories from one viewpoint to another and back as two different motion trajectories and there are m different camera viewpoints (plus

the initial robot posture), then there are $(m + 1) \times m$ possible motion trajectories between different viewpoints. In our practical example, this results in 110 possible trajectories. They can be either generated in advance and stored or they can be generated online.

B. MOTION PLANNERS

Many different motion planners exist, as presented in the related work section. We used MoveIt! [35] software package for motion planning and Planner arena [12] for benchmarking of different planners for steering the camera in adaptive visual quality inspection. We tested 21 different motion planners available in Open Motion Planning Library (OMPL) [13]. The planners are listed in Table 3 and the number associated with each planner is used to present the results of our evaluations.

TABLE 3. List of tested OMPL planning algorithms

Pl. #	Planner name
1	BiEST (Bidirectional version of EST)
2	ProjEST (Projection-based version of EST)
3	LazyPRM
4	LazyPRM*
5	SPARS (SPARse Roadmap Spanner algorithm)
6	SPARS2
7	SBL (Single-query, Bi-directional and Lazy Collision Checking)
8	EST (Expansive Space Trees)
9	LBKPIECE (Lazy BKPIECE)
10	BKPIECE (Bidirectional KPIECE)
11	KPIECE (Kinematic Planning by Interior-Exterior Cell Exploration)
12	RRT (Rapidly-exploring Random Trees)
13	RRT Connect
14	RRT*
15	T-RRT (Transition-based RRT)
16	PRM (Probabilistic roadmap)
17	PRM*
18	FMT (Fast Marching Tree algorithm)
19	PDST (Path-Directed Subdivision Trees)
20	STRIDE (Search Tree with Resolution Independent Density Estimation)
21	BiTRRT (Bi-directional Transition-based RRT)

C. EXPERIMENTAL PROTOCOL AND METRICS

Each planner was used to generate trajectories between all possible pairs of 11 robot postures (10 viewpoints + the starting point), 250 times for each pair. Several executions were executed because the planners use random sampling when generating the trajectory points. In general, there is no motion planner that is optimal for every given use-case [12]. Thus, we checked various parameters to determine which planner performed best in our practical example. The calculations were done on a PC with AMD Ryzen Threadripper PRO 5975WX processor and 512GB of RAM.

First, one needs to check if the robot motion generated by the selected planner reaches the desired final pose. If the planner does not converge, it is immediately ruled out.

Next, since short cycle times are needed to prevent bottlenecks in the production, the average time required for




FIGURE 9. Seven out of eleven postures for visual quality inspection shown in RViz visualization (a) and in real world (b).

planning was checked.

The length of motion, as the next evaluation metrics, was computed as the sum of joint angle differences travelled by all robot joints starting at the initial and final robot configuration \mathbf{q}_0 and \mathbf{q}_n , respectively

$$L = \sum_{i=1}^n \|\mathbf{q}_i - \mathbf{q}_{i-1}\|. \quad (1)$$

Finally, smoothness S of the Cartesian space path, as introduced in [43], was evaluated. It was obtained by averaging the squared angles between consecutive trajectory segments

$$S = \frac{1}{n} \sum_{i=1}^{n-1} \phi_i^2. \quad (2)$$

To compute the angles, we sampled a sequence of $n + 1$ positions $\mathbf{r}_0, \mathbf{r}_1, \dots, \mathbf{r}_n \in \mathbb{R}^3$ on the Cartesian space path, distributed from the beginning to the end of the path. The angle ϕ_i between each triple of consecutive points $\{\mathbf{r}_{i-1}, \mathbf{r}_i, \mathbf{r}_{i+1}\}$ can be calculated using

$$\phi_i = \pi - \arccos \left(\frac{(\mathbf{r}_{i-1} - \mathbf{r}_i)^\top (\mathbf{r}_{i+1} - \mathbf{r}_i)}{\|\mathbf{r}_{i-1} - \mathbf{r}_i\| \|\mathbf{r}_{i+1} - \mathbf{r}_i\|} \right). \quad (3)$$

VI. MOTION PLANNER EVALUATION

We evaluated the planners listed in Table 3. The results of convergence to the desired posture are shown in Table 4. It can be seen that some planners did not converge to the solution in all 250 planning sessions. These planners were immediately ruled out.

TABLE 4. Success rate of planners for 250 planning attempts and 110 movements between the specified camera postures

#	1	2	3	4	5	6	7	8	9	10	11
%	1	1	1	1	0.98	0.66	1	1	1	1	1
#	12	13	14	15	16	17	18	19	20	21	
%	0.99	1	0.99	0.98	1	1	1	1	1	1	

As shown in Fig. 10a, some of the planners took an order of magnitude longer than the others and reached the cut-off

time of 2 seconds. These planners were also excluded from further consideration.

Length of motion metrics results, as defined in (1), are presented in Fig. 10b. These were combined with Cartesian path smoothness. Considering the results in Table 4 and Fig. 10, the best result was obtained by planner #21, Bi-directional Transition-based Rapidly-exploring Random Trees (BiRRT). This planner combines the exploratory strength of RRTs with the efficiency of stochastic optimization methods (e.g., Monte Carlo optimization) [44]. The BiRRT method converged to a solution all the time, was very fast, and the resulting trajectories were smooth and short. Variations of RRT, such as T-RRT performed on the same level if not better, but did not converge in all planning attempts. Planners #1, 2, 19 and 20 were excluded because of longer and less smooth trajectories, even though the results were close to the chosen planner.

VII. DISCUSSION

It is important to note that parts for the automotive industry are produced in large series. Any possible defects are problematic for the overall production process because they might cause delays in the time of delivery of the final product or even recall of the vehicles that are already being used in traffic. Since such recalls can be expensive for the manufacturer and cause the loss of trust from the customers, it is important for the manufacturers to invest and implement better methods for visual quality inspection. It is desirable to reduce the number of undetected faulty parts to zero. Note that technological limitations still cannot provide 100% optical control in real-world production environments.

The defect prediction approach presented in this paper mostly depends on the available data. Additional data-points can always be added and the clustering re-trained. Thus, with an increasing amount of data points for good and faulty parts, it is possible to raise the trust in the prediction of possible defects. Manufacturing companies, however, tend to produce orders of magnitude more of good parts than faulty parts. Thus, collecting a large database, especially at the end of the production process where the parts that reached this point

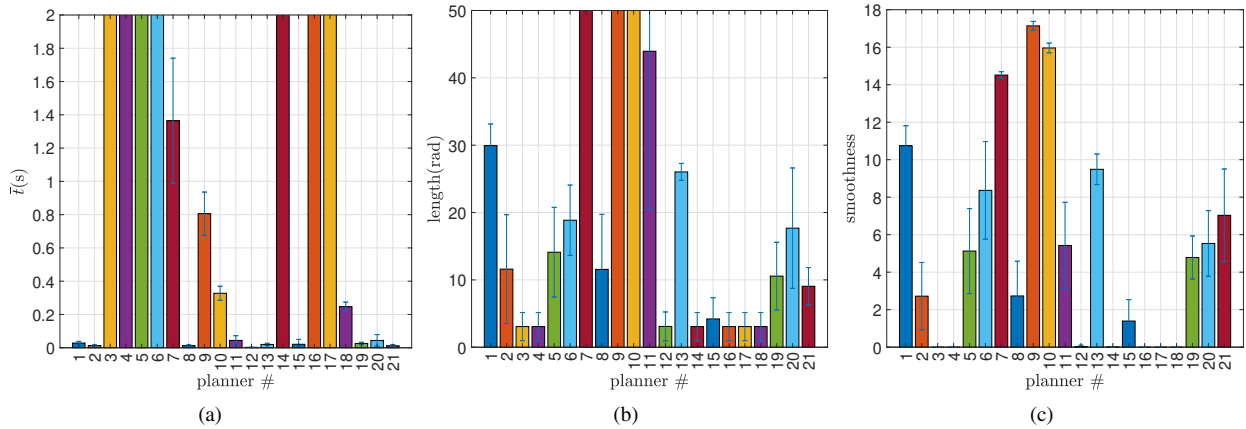


FIGURE 10. Planner evaluation results. (a) planning times. (b) length as defined in (1), cut at 50 rad. (c) smoothness as defined in (2).

have already passed the intermediate quality checks, is not easy. As can be seen in the practical production example considered in this paper, the database of good and faulty parts (DB1) can be extremely lopsided.

With more data available, the choice of classification algorithms may also change. Our proposed approach shows good classification results even though the amount of relevant data was low. Furthermore, the comparisons show that some approaches are really not appropriate in this case, i.e, supervised learning classifiers are known to require a lot of data [21]. The proposed approach combines both AEs and clustering. Using them separately, for clustering in the data space, or for detecting defects through reconstruction error has not shown good results. Conceptually similar to our method is also deep embedded clustering (DEC), where the clustering is implemented as an additional layer of the encoder. Since it can only be trained on the outcome of already implemented clustering, it offers only an additional implementation which will take the data as the input and provide directly the possibility of belonging to a cluster, without the need to check separately. Thus, the results will be the same.

Our approach can be used in two ways: either to i) check all the parts after the final production stage, but using the prediction model to change the order of inspections, or ii): to only checks the most likely defects. Under i), all defects will be found, but it will take the longest, as shown in Table 5, which shows the average times for 100 iterations of the inspections for all the defect combinations. With enough confidence in the prediction model, we can save close to 80% of the time when checking for only one defect, or around 55% of the time when checking for two, as listed in Table 5. Checking for 3 defects saved roughly 36% of the time, while checking for 4 out of the five most likely defects on average saved around 23% of the inspection time.

In our approach, the robot must frequently move between different pairs of viewpoints no matter if we check for all or only some of the defects. We therefore tested which of the available motion planners is best for our practical production process. As stated in [12], there is no motion planner that

TABLE 5. Time required for inspecting different combinations of defects using the selected BiRRT motion planner. The Table shows the mean and standard deviation of computational time needed for trajectory planning and execution for different defect combinations as well as the percent of the saved time compared to the inspection from all viewpoints.

Inspected defect	μ (s)	σ (s)	saved (%)
All	42.21	0.51	0.00 ± 1.21
bottom	2.66	0.40	93.69 ± 1.02
pins	24.92	0.50	40.96 ± 1.91
porous	11.06	0.44	73.80 ± 1.36
windings	2.68	0.38	93.66 ± 0.98
sides	2.67	0.38	93.67 ± 0.98
Average 1	8.80	0.42	79.16 ± 1.25
bottom + pins	29.12	0.85	31.01 ± 2.84
bottom + porous	20.23	1.46	52.07 ± 4.04
bottom + windings	8.34	0.44	80.24 ± 1.29
bottom + sides	8.73	0.22	79.32 ± 0.77
pins + porous	28.08	0.25	33.47 ± 1.39
pins + windings	29.20	0.37	30.82 ± 1.72
pins + sides	29.22	0.38	30.77 ± 1.73
porous + windings	18.13	0.17	57.05 ± 0.93
porous + sides	18.42	0.38	56.37 ± 1.44
windings + sides	2.71	0.38	93.59 ± 0.98
Average 2	19.22	0.49	54.47 ± 1.71
bottom + pins + porous	32.07	0.47	24.03 ± 2.04
bottom + pins + windings	32.97	0.47	21.88 ± 2.06
bottom + pins + sides	32.93	0.43	22.00 ± 1.96
bottom + porous + windings	24.90	2.22	41.02 ± 5.97
bottom + porous + sides	24.79	1.80	41.26 ± 4.97
bottom + windings + sides	8.71	0.19	79.37 ± 0.70
pins + porous + windings	31.91	0.44	24.40 ± 1.95
pins + porous + sides	32.07	0.38	24.03 ± 1.83
pins + windings + sides	29.10	0.47	31.05 ± 1.94
porous + windings + sides	18.18	0.35	56.94 ± 1.35
Average 3	26.76	0.72	36.6 ± 2.48
bottom + pins + porous + windings	35.82	0.32	15.15 ± 1.78
bottom + pins + porous + sides	35.79	0.48	15.22 ± 2.17
bottom + pins + windings + sides	32.83	0.41	22.23 ± 1.92
bottom + porous + windings + sides	24.77	2.55	41.32 ± 6.76
Average 4	32.27	0.84	23.56 ± 2.91

works best in every situation. For the considered production process, the extensive statistical evaluation has shown that several planners might be appropriate. The best results were obtained with Bi-directional Transition-based Rapidly-exploring Random Trees (BiTRRT) method. In this example, learning 110 trajectories and then recalling them from the

database is still viable. Nonetheless, it is easy to imagine visual inspection processes where many more motion trajectories would need to be generated. This is, however, not necessary because the available motion planners can generate trajectories for visual quality inspection in a fast and reliable manner.

For completeness of the adaptive visual quality inspection approach, we implemented a deep neural network classifier that takes the image as the input and classifies the part as good or faulty at the output. The results of this classification for the incomplete bottom edge defect are shown in Fig. 11.




FIGURE 11. Results of final visual quality inspection for the "incomplete bottom edge" (e2) defect of several parts. Faulty parts are marked with a red edge.

VIII. CONCLUSION

The presented approach, demonstrated on real-world data and objects, takes the production parameters of an injection moulding line to predict potential defects of the manufactured parts. We have shown that both supervised and unsupervised classification methods can struggle with the problem. As is common with learning approaches, a small database prevents the use of supervised learning classifiers, as they either overfit or achieve poor classification results. Our proposed, unsupervised method shows how to use autoencoders and clustering of production parameters to predict what kind of defect might be present.

The evaluation has shown that the clustering and classification are robust against the training of AE networks. The classification results and the comparison to the real data show, that the proposed method can be effectively applied for default prediction. The use of the proposed method can, without having to check the most common defects, reduce their number without introducing a delay in the production. It should be noted, that the proposed approach provides the initial, starting point in the production. Once added to the production line, the database can be potentially increased. Thus, either our approach or even supervised deep learning classification approaches could, at a later stage, be applied more effectively.

We have also analyzed different motion planners to generate motions that bring visual inspection cameras to the required viewpoints in a collision-free manner. The timing results show that considerable time can be saved when only

some of the most common potential defects are inspected. However, it is the prediction which potential defects need to be checked for, that makes the approach more reliable.

While in our approach we close the feedback loop between the prediction of the defect and the camera viewpoint, for a true ZDM approach the closed loop should consider changing the production parameters based on the prediction of feedback. This remains an open research question and future work. With a larger database, one could improve the confidence in the predictions from model. Additionally, a larger database would enable one to evaluate the effect of specific models on the outcome of the production, for example with Shapely values [45], and thus influence the production process in a closed-loop manner.

REFERENCES

- [1] Harald Foidl and Michael Felderer. Research challenges of industry 4.0 for quality management. In Michael Felderer, Felix Piazzolo, Wolfgang Ortner, Lars Brehm, and Hans-Joachim Hof, editors, *Innovations in Enterprise Information Systems Management and Engineering*, pages 121–137, Cham, 2016. Springer International Publishing.
- [2] Foivos Psarommatis, Gökan May, Paul-Arthur Dreyfus, and Dimitris Kiritsis. Zero defect manufacturing: state-of-the-art review, shortcomings and future directions in research. *International Journal of Production Research*, 58(1):1–17, 2020.
- [3] Artem A. Nazarenko, Joao Sarraipa, Luis M. Camarinha-Matos, Christian Grunewald, Marc Dorchain, and Ricardo Jardim-Goncalves. Analysis of relevant standards for industrial systems to support zero defects manufacturing process. *Journal of Industrial Information Integration*, 23:100214, 2021.
- [4] Florian Eger, Daniel Coupek, Davide Caputo, Marcello Colledani, Mariluz Penalva, Jon Ander Ortiz, Hermann Freiberger, and Gernot Kollegger. Zero defect manufacturing strategies for reduction of scrap and inspection effort in multi-stage production systems. *Procedia CIRP*, 67:368–373, 2018. 11th CIRP Conference on Intelligent Computation in Manufacturing Engineering, 19-21 July 2017, Gulf of Naples, Italy.
- [5] Mohamed Ismail, Noha A Mostafa, and Ahmed El-Assal. Quality monitoring in multistage manufacturing systems by using machine learning techniques. *Journal of Intelligent Manufacturing*, 33(8):2471–2486, 2022.
- [6] Ming Gan and Jianhua Wang. Applications of image processing technique in porous material characterization. In Dimitrios Venzas, editor, *Advanced Image Acquisition, Processing Techniques and Applications*, chapter 6. IntechOpen, Rijeka, 2012.
- [7] Jinjiang Wang, Peilun Fu, and Robert X Gao. Machine vision intelligence for product defect inspection based on deep learning and hough transform. *Journal of Manufacturing Systems*, 51:52–60, 2019.
- [8] Niko Herakovic. *Robot vision in industrial assembly and quality control processes*. INTECH Open Access Publisher London, UK, 2010.
- [9] Zvezdan Lončarević, Andrej Gams, Simon Reberšek, Bojan Nemeč, Jure Škrabar, Jure Skvarč, and Aleš Ude. Specifying and optimizing robotic motion for visual quality inspection. *Robotics and Computer-Integrated Manufacturing*, 72:102200–1–102200–14, 2021.
- [10] Justus Zipfel, Felix Verworner, Marco Fischer, Uwe Wieland, Mathias Kraus, and Patrick Zschech. Anomaly detection for industrial quality assurance: A comparative evaluation of unsupervised deep learning models. *Computers & Industrial Engineering*, 177:109045, 2023.
- [11] Olga Ogorodnyk, Ole Vidar Lyngstad, Mats Larsen, and Kristian Martinsen. Prediction of width and thickness of injection molded parts using machine learning methods. *EcoDesign and Sustainability I: Products, Services, and Business Models*, pages 455–469, 2021.
- [12] Mark Moll, Ioan A. Sucan, and Lydia E. Kavraki. Benchmarking motion planning algorithms: An extensible infrastructure for analysis and visualization. *IEEE Robotics & Automation Magazine*, 22(3):96–102, 2015.
- [13] Ioan A. Sucan, Mark Moll, and Lydia E. Kavraki. The Open Motion Planning Library. *IEEE Rob. & Aut. Magazine*, 19(4):72–82, 2012.
- [14] Richárd Dominik Párizs, Dániel Török, Tatyana Ageyeva, and József Gábor Kovács. Machine learning in injection molding: An industry 4.0 method of quality prediction. *Sensors*, 22(7), 2022.

- [15] Hail Jung, Jinsu Jeon, Dahui Choi, and Jung-Ywn Park. Application of machine learning techniques in injection molding quality prediction: Implications on sustainable manufacturing industry. *Sustainability*, 13(8), 2021.
- [16] Olga Ogorodnyk, Ole Vidar Lyngstad, Mats Larsen, Kesheng Wang, and Kristian Martinsen. Application of machine learning methods for prediction of parts quality in thermoplastics injection molding. In *Advanced Manufacturing and Automation VIII 8*, pages 237–244. Springer, 2019.
- [17] Bruno Silva, João Sousa, and Guillem Alenya. Machine learning methods for quality prediction in thermoplastics injection molding. In *2021 Int. Conf. on Electrical, Computer and Energy Technologies (ICECET)*, 2021.
- [18] Bruno Silva, Ruben Marques, Tiago Santos, João Sousa, and Guillem Alenya. Relevant parameters identification in traditional & stretch and blow thermoplastics injection molding. In *2022 IEEE 9th International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications (CIVEMSA)*, 2022.
- [19] Bruno Silva, João Sousa, and Guillem Alenya. Data acquisition and monitoring system for legacy injection machines. In *2021 IEEE International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications (CIVEMSA)*, 2021.
- [20] Le Wang, Meng Han, Xiaojuan Li, Ni Zhang, and Haodong Cheng. Review of classification methods on unbalanced data sets. *IEEE Access*, 9:64606–64628, 2021.
- [21] Sotiris B Kotsiantis, Ioannis Zaharakis, P Pintelas, et al. Supervised machine learning: A review of classification techniques. *Emerging artificial intelligence applications in computer engineering*, 160(1):3–24, 2007.
- [22] Debashree Devi, Saroj K. Biswas, and Biswajit Purkayastha. A review on solution to class imbalance problem: Undersampling approaches. In *2020 International Conference on Computational Performance Evaluation (ComPE)*, pages 626–631, 2020.
- [23] Anuraganand Sharma, Prabhat Kumar Singh, and Rohitash Chandra. Smotified-gan for class imbalanced pattern classification problems. *IEEE Access*, 10:30655–30665, 2022.
- [24] Mohamed Alloghani, Dhiya Al-Jumeily, Jamila Mustafina, Abir Hussain, and Ahmed J. Aljaaf. *A Systematic Review on Supervised and Unsupervised Machine Learning Algorithms for Data Science*, pages 3–21. Springer International Publishing, Cham, 2020.
- [25] Yibo Zhou. Rethinking reconstruction autoencoder-based out-of-distribution detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7379–7387, June 2022.
- [26] Xuan Huang, Zhenlong Hu, and Lin Lin. Deep clustering based on embedded auto-encoder. *Soft Computing*, 27(2):1075–1090, 2023.
- [27] Javier Villalba-Diez, Daniel Schmidt, Roman Gevers, Joaquín Ordieres-Meré, Martin Buchwitz, and Wanja Wellbrock. Deep Learning for Industrial Computer Vision Quality Control in the Printing Industry 4.0. *Sensors (Basel, Switzerland)*, 19(18), 2019.
- [28] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition*, 112:107706, 2021.
- [29] B. Triggs and C. Laugier. Automatic camera placement for robot vision tasks. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1732–1737, Nagoya, Japan, 1995.
- [30] C. Zang and K. Hashimoto. A flexible camera positioning strategy for robot-based visual inspection applications. In *IEEE Int. Conf. on Mechatronics and Automation*, pages 527–532, Chengdu, China, 2012.
- [31] Robert Bevec, Timotej Gašpar, and Aleš Ude. Robot-driven autofocus control mechanism for an in-hand fixed focus camera. In Nikos A. Aspragathos, Panagiotis N. Koustoumpardis, and Vassilis C. Moulitanitis, editors, *Advances in Service and Industrial Robotics*, pages 551–559, Cham, 2019. Springer International Publishing.
- [32] Pablo Malvido Fresnillo, Saigopal Vasudevan, Wael M. Mohammed, Jose L. Martinez Lastra, and Jose A. Perez Garcia. Extending the motion planning framework–MoveIt with advanced manipulation functions for industrial applications. *Robotics and Computer-Integrated Manufacturing*, 83:102559–1–102559–18, 2023.
- [33] Howie Choset, Kevin M. Lynch, Seth Hutchinson, George Kantor, Wolfram Burgard, Lydia Kavraki, and Sebastian Thrun. *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press, 5 2005.
- [34] Alessandro Gasparetto, Paolo Boscaroli, Albano Lanzutti, and Renato Vidoni. Path planning and trajectory planning algorithms: A general overview. *Mechanisms and Machine Science*, 2015.
- [35] David Coleman, Ioan Sucan, Sachin Chitta, and Nikolaus Correll. Reducing the barrier to entry of complex robotic software: a moveit! case study, 2014.
- [36] Zachary Kingston, Mark Moll, and Lydia E. Kavraki. Exploring implicit spaces for constrained sampling-based planning. *Intl. J. of Robotics Research*, 38(10–11):1151–1178, September 2019.
- [37] Mrinal Kalakrishnan, Sachin Chitta, Evangelos Theodorou, Peter Pastor, and Stefan Schaal. Stomp: Stochastic trajectory optimization for motion planning. In *IEEE Int. Conf. on Robotics and Automation*, pages 4569–4574, 2011.
- [38] Nathan Ratliff, Matt Zucker, J. Andrew Bagnell, and Siddhartha Srinivasa. Chomp: Gradient optimization techniques for efficient motion planning. In *IEEE International Conference on Robotics and Automation*, pages 489–494, 2009.
- [39] Maxim Likhachev. Search-based planning library (sbpl), 2016.
- [40] L Breiman. Random forests. *Machine Learning*, 45:5–32, 10 2001.
- [41] Bekhzod Alisher ugli Olimov, Kalyana C. Veluvolu, Anand Paul, and Jeonghong Kim. Uzadl: Anomaly detection and localization using graph laplacian matrix-based unsupervised learning method. *Computers & Industrial Engineering*, 171:108313, 2022.
- [42] Zhiqiang Wang, Catherine Da Cunha, Mathieu Ritou, and Benoît Furet. Comparison of k-means and gmm methods for contextual clustering in hsm. *Procedia Manufacturing*, 28:154–159, 2019. 7th International conference on Changeable, Agile, Reconfigurable and Virtual Production (CARV2018).
- [43] Benjamin Cohen, Ioan A. Şucan, and Sachin Chitta. A generic infrastructure for benchmarking motion planners. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 589–595, 2012.
- [44] Didier Devaurs, Thierry Siméon, and Juan Cortés. Enhancing the transition-based rrt to deal with complex cost spaces. In *2013 IEEE International Conference on Robotics and Automation*, pages 4120–4125, 2013.
- [45] Scott Lundberg and Su-In Lee. A unified approach to interpreting model predictions, 2017.



ZVEZDAN LONČAREVIĆ received his B.S. degree in Mechatronics Engineering from the Faculty of Technical Sciences, University of Novi Sad, in 2017. He obtained his M.S. and Ph.D. degrees in Information and Communication Technologies from Jožef Stefan International Postgraduate School, Ljubljana, in 2019 and 2023, respectively.

Since 2019, he has been a Research Assistant at Jožef Stefan Institute - Department of Automatics, Biocybernetics, and Robotics (E1), 1000 Ljubljana, Slovenia. His research interests include robotics, reinforcement learning, deep learning, and transfer learning. He has worked on the QU4LITY (Digital Reality in Zero Defect Manufacturing) project, where he applied optimization techniques for zero-defect manufacturing and used deep learning for adaptive visual quality inspection. Additionally, he contributed to the GOSTOP (The Building Blocks, Tools, and Systems for the Factories of the Future) project, focusing on learning and optimization for industrial robotics.



SIMON REBERŠEK received the M.S. degree in electrical engineering from Faculty of electrical engineering University of Ljubljana.

He is currently working as research assistant at Jožef Stefan Institute in the field of robotics.



SAMO ŠELA received a B.S. degree from the Faculty of Computer and Information Science at the University of Ljubljana in 2016. He has actively complemented his computer vision and artificial intelligence knowledge by taking courses in the increasingly widespread mass open educational platforms such as Coursera and Udacity.

From 2012 to 2022, he held various roles in the company Kolektor in machine vision and deep learning. He contributed to the realization of demanding machine vision solutions in industrial environments and made vital contributions to developing the KiS, Kolektor Imaging Software. His drive and thinking about the future took him on the road of research and development for the factories of the future at Kolektor Group in 2016. In 2018, he began his journey to obtain a doctoral degree. Since 2022, he has been the head software engineer at SICK Slovenia.



ANDREJ GAMS received the Diploma degree in electrical engineering in 2004, and the Ph.D. degree in robotics from the University of Ljubljana, Slovenia, in 2009.

He is currently the Head of Humanoid and Cognitive Robotics Laboratory at the Department of Automatics, Biocybernetics, and Robotics, Jožef Stefan Institute, Ljubljana. He was a Postdoctoral Researcher with the Biorobotics Laboratory, Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland, during 2012–2013. His research interests include robot and imitation learning, applied for industrial and household robotics and handling of deformable materials with robots.

...



JURE SKVARČ received his bachelor's degree in 1992 and his Ph.D. in 1998 from the Faculty of Electrical Engineering at the University of Ljubljana. His research covers image processing, machine vision, machine learning and digital signal processing.

Since 1998, his primary focus has been on the implementation of complex image-processing systems in industrial environments. Over the last 25 years, he has been involved in more than a thousand projects for users in various industries. In 1998, he started his non-academic career as an independent researcher and became Technical Director at FDS Research. From 2012 to 2022, he was Managing Director at Kolektor, BU VISION. Since 2022, he has been responsible for the R&D team at SICK Slovenia. He leads the interdisciplinary team specialising in the research, development and implementation of vision systems. His current research focuses on integrating new concepts into the value chain structure, focusing on machine vision and artificial intelligence as key technologies of Industry 4.0. He is active in international organisations, particularly EMVA and VDMA, where he contributed to the first version of the international standard OPC UA Vision CS.



ALEŠ UDE received the diploma degree in applied mathematics from the University of Ljubljana, Slovenia, in 1990, and the doctoral degree in engineering sciences from the University of Karlsruhe, Germany, in 1995.

He is currently the Head of the Department of Automatics, Biocybernetics, and Robotics, Jožef Stefan Institute, Ljubljana. His research interests include robot learning, imitation learning, reconfigurable robotic systems, and humanoid robotics. He has been a coordinator and/or a principal investigator of numerous national and international projects in these areas.