

Fast mesh denoising with data driven normal filtering using deep variational autoencoders

Stavros Nousias^{1,2}, Gerasimos Arvanitis², Aris S. Lalos¹, and Konstantinos Moustakas²

¹ *Industrial Systems Institute, Athena Research Center {nousias, lalos}@isi.gr*

² *Department of Electrical & Computer Engineering University of Patras, Greece {nousias, arvanitis, moustakas}@ece.upatras.gr*

Abstract—

Recent advances in 3D scanning technology have enabled the deployment of 3D models in various industrial applications like digital twins, remote inspection and reverse engineering. Despite their evolving performance, 3D scanners, still introduce noise and artifacts in the acquired dense models. In this work, we propose a fast and robust denoising method for dense 3D scanned industrial models. The proposed approach employs conditional variational autoencoders to effectively filter face normals. Training and inference are performed in a sliding patch setup reducing the size of the required training data and execution times. We conducted extensive evaluation studies using 3D scanned and CAD models. The results verify plausible denoising outcomes, demonstrating similar or higher reconstruction accuracy, compared to other state-of-the-art approaches. Specifically, for 3D models with more than $1e4$ faces, the presented pipeline is twice as fast as methods with equivalent reconstruction error.

Index Terms—3D mesh denoising, data driven normal filtering, variational autoencoders.

I. INTRODUCTION

Industrial sites and manufacturing plants, often require infrastructure upgrades and construction projects. Modifications in the production line lead to downtime, high costs or unwanted delays. Digital twins could allow for improved supervision, inspection and monitoring based on simulation studies. Subsequently, they could enable the identification of errors while handling risk and dealing with liabilities. 3D scanning can facilitate the development of accurate, high-end digital twins of manufacturing processes and factory layouts.

Furthermore, smart manufacturing encompasses “fully-integrated, collaborative manufacturing systems that respond in real-time to meet changing demands and conditions in the factory, in the supply network, and in customer needs” [1], [2]. Accurate digital reconstruction for material inspection, quality control and reverse engineering are challenges in this evolving landscape. Inline 3D scanning could allow the examination of constructed parts, in many stages of the manufacturing process. [3]–[6].

Towards this direction, technologies in metrology have changed, in the past few decades, from stand-alone coordinate measuring machines (CCMs) to portable 3D scanners. The

benefits of 3D scanning become evident in several use cases. Quality assurance protocols in automotive industry [4], [7], maintenance processes in maritime industry [8] and automated reverse engineering prove that error-free representations are a requirement for Industry 4.0 outcomes.

Several use cases appear in the literature, i.e. Artec 3D reports that a foundry [9] uses handheld solutions to scan 3D castings, saving time and increasing productivity. The authors in [10] present a use case of a surface inspection method for wind turbines, employing an autonomous robotic arm equipped with a 3D scanner. Moreover, state-of-the-art 3D industrial scanning outcomes, available online [11] reveal that models of 2 million sampled points require up to 30 minutes for scanning and up to 30 minutes for post-processing. Robust, high accuracy and low-error processing outcomes would reduce the scanning time since the collection of fewer samples would be necessary for equivalent results, while fast processing would reduce the post-processing times.

These challenging issues highlight the need for parallelizable computationally inexpensive, and accurate approaches for mesh denoising. In a classic scenario, scanners yield noisy point clouds that are consequently converted to noisy 3D meshes. Denoising aims to remove the noise while preserving features and multi-scale geometric details. Noise is usually inserted by scanning devices and digitization processes, thus making mesh denoising an important post-processing step. Several methods are available in the literature with significant denoising results [12]–[16]. Yet the need for robust and fast algorithms, able to handle dense models rapidly, becomes essential in industrial applications [17], [18], where they are expected to significantly reduce the operational cost of many manufacturing tasks.

Motivated by the aforementioned challenges, we provide a fast approach for mesh denoising, based on data-driven normal filtering. We employ deep conditional variational autoencoders allowing to handle efficiently dense models. We summarize the contributions of the proposed approach in the following points:

- The network can localize since, training and inference are performed in a sliding patch setup. The filtered face normal vector is generated by providing a patch of neighbouring faces as input, corresponding to a local region around that face.
- It requires a relatively small training set. We propose a preprocessing method that describes each patch with a

This work was supported by European Union Horizon 2020 Research and innovation program “WARMEST -or IoW Altitude Remote sensing for the Monitoring of the state of cultural hEritage Sites: building an inTegrated model for maintenance” under Marie Skłodowska grant agreement No 777981

This work was supported by European Union Horizon 2020 Research and innovation program Ageing@Work under Ageing@Work grant agreement no 826299.

scale, translation and rotation invariant representation.

- Evaluation studies indicate that our approach demonstrates lower complexity and execution times than other non-data-driven state-of-the-art methods. Specifically, for 3D models with more than 10^4 faces, the presented pipeline is twice as fast as methods with an equivalent reconstruction error.
- It is fully parallelizable. We evaluated the execution efficiency with respect to the number of utilized cores and type of processing unit (i.e., CPU or GPU).
- It can be employed for feature-preserving denoising of dense 3D models, with different noise patterns. Such a property would be ideally suited for industrial applications where 3D scanners with different properties are generating dense representations of physical objects.
- It is parameter-free since every used parameter is predefined and the user does not need to search for optimal values per model.

Evaluation studies were carried out using scanned and CAD 3D industrial models. Our results verify the effectiveness of the proposed method, compared to other state-of-the-art approaches, both in terms of denoising quality and computational efficiency.

The rest of this paper is organized as follows: Section II presents state-of-the-art methods and related works. Section III focused on preliminaries. Section IV describes the workflow of the proposed approach in detail. Section V is dedicated to the experimental setup and simulation results, while conclusions are drawn in Section VI.

II. RELATED WORK

Mesh denoising approaches can be organized in the following categories: isotropic and anisotropic mesh filtering based, regularization based and data-driven methods.

Isotropic and anisotropic mesh filtering: Laplacian and Taubin smoothing are well-known approaches that remove noise and artifacts by employing iterative vertex update based on the Laplacian matrix of the geometry. To the same direction, Desbrun et al. used the mean curvature flow [19] to allow treating of irregular surfaces. However, one of their main disadvantages is that they do not preserve geometric features. Another category of methods, namely graph spectral processing [20], employs singular values, eigenvectors and eigenspace projections to separate 3D mesh data from noise. However, their disadvantage is that they are computationally expensive and resource consuming. Other widely accepted feature preserving approaches [21] process vertex positions locally while preserving the geometric features. The main drawback is, in many cases, the deformation of large scale features. Mesh bilateral filtering methods use normal coordinates to estimate the parameters of noise removing filters, with the reconstruction accuracy relying heavily on noise characteristics [22]. Such approaches are based on normal filtering and vertex position update [12]–[14], [23] consisting of two iterative stages. The first stage filters the face normals while the second updates the position of vertices. Although this

category of approaches preserves most of the sharp features, they require heavy parameterization and fail to generalize.

Regularization based: Regularizers are often used for ill-posed problems. Denoising of 3D meshes is in many cases an ill-posed problem due to sensing limitations and non-uniform sampling operations. Zhang et al. [24] minimize the energy of both vertex position and normal error and He et al. [15], propose an L_0 minimization approach. Even though they demonstrate accurate surface reconstruction in Gaussian noise cases, the computational cost is high, and the denoising outcomes deteriorate with other noise types. Furthermore, a cascaded denoising framework is presented by the authors in [25]. Their approach includes multi-scale tensor voting, vertex clustering step for detecting sharp features and a piece-wise fitting step for preserving the identified features.

Learning based: Several studies employ deep networks for mesh denoising [16], [26]–[28]. The authors in [16] suggest a data-driven method for mesh denoising that uses training sets of noisy objects. The objects are scanned by the same devices, thus facilitating the denoising of geometries with similar noise. The geometric features are reconstructed sufficiently. Yet, their main limitation is that the reconstruction accuracy of important details relies heavily on whether they were initially included in the training set. In the same fashion, the authors in [26] present a two-step ELM based denoising approach, where the first step performs coarse denoising and the second step recovers features. Several other research groups use CNNs working on voxelized versions of the geometry [27], or on images derived from local patches [28]. However, our approach aims to be applied directly on the mesh nodes avoiding preprocessing, thus contributing to the field of geometric deep learning where the sampling of the latent space is nonuniform.

III. PRELIMINARIES

A. Preliminaries on deep autoencoders

Deep autoencoders encompass a multi-layer neural network architecture where the hidden layers encode the input to a latent space and decode the latter to a reconstructed output. A deep autoencoder is composed of two symmetrical deep-belief networks [29] that typically have three to five shallow layers for the encoding and the decoding part. The layers are Restricted Boltzmann Machines (RBMs). Variational autoencoders (VAE) [30] assume that the input vectors are generated by some random process of an unobserved continuous random variable \mathbf{z} . The parameters of the VAE are estimated efficiently by the stochastic gradient variational Bayes framework [30]. Furthermore, conditioning input vector \mathbf{x} under label c constitutes the basis of conditional variational autoencoder (CVAE) [31].

B. Preliminaries on 3D meshes

In this work, we focus on triangular meshes \mathcal{M} with n vertices \mathbf{v} and n_f faces f . Each vertex \mathbf{v}_i is denoted by $\mathbf{v}_i = [x_i, y_i, z_i]^T$, $\forall i = 1, \dots, n$. Each f_j face is a triangle that can be described by its centroid:

$$\mathbf{c}_j = (\mathbf{v}_{j1} + \mathbf{v}_{j2} + \mathbf{v}_{j3})/3 \quad (1)$$

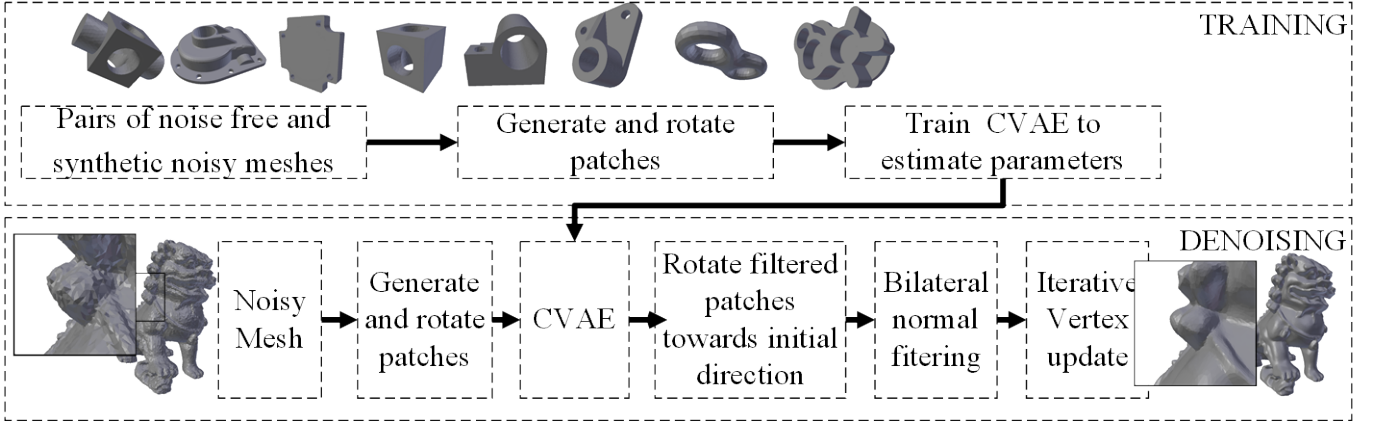


Fig. 1: Pipeline of the proposed approach.

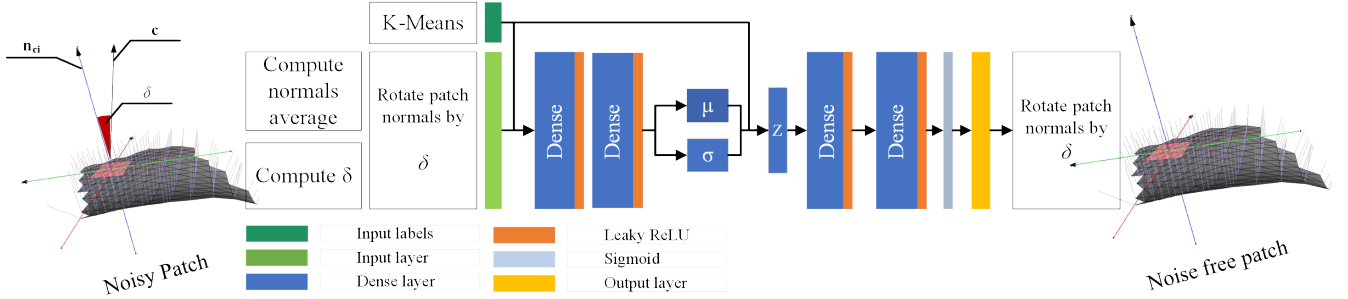


Fig. 2: CVAE architecture and training scheme. Normal coordinates of noisy and noise-free patches are employed. Initially, they are properly rotated and labeled with k-means clustering.

and its outward unit normal:

$$\mathbf{n}_{c_i} = \frac{(\mathbf{v}_{j2} - \mathbf{v}_{j1}) \times (\mathbf{v}_{j3} - \mathbf{v}_{j1})}{\|(\mathbf{v}_{j2} - \mathbf{v}_{j1}) \times (\mathbf{v}_{j3} - \mathbf{v}_{j1})\|} \quad (2)$$

where \mathbf{v}_{j1} , \mathbf{v}_{j2} and \mathbf{v}_{j3} are the position of the vertices that define face $f_j = \{\mathbf{v}_{j1} \ \mathbf{v}_{j2} \ \mathbf{v}_{j3}\}$, $\forall j = 1, \dots, n_f$. The first-ring area of a vertex \mathbf{v}_i is defined as the neighborhood \mathcal{N}_i in which the vertex \mathbf{v}_i is connected to other vertices by only one edge (i.e., with topological degree equal to 1).

IV. AUTOENCODER ARCHITECTURES FOR 3D MESH DENOISING

This section presents the mesh denoising pipeline. Training data were generated from meshes, distorted by noise, using the normal vectors corresponding to the 3D mesh faces. After training the autoencoder, the generated output vector is used for a normal-based vertex update [12] of the mesh vertices.

A. Autoencoder architectures for mesh denoising

This section presents the deep network architecture for mesh denoising. Specifically, a conditional variational autoencoder [31] was employed, as illustrated in Figure 2. A conditional Gaussian encoder with two dense layers is succeeded by a conditional Bernoulli decoder with two dense layers. Each dense layer is succeeded by a layer of leaky rectified linear units (ReLUs) and a dropout layer. We denote \mathbf{X} as the input tensor, \mathbf{Y} the corresponding labels, $\mathbf{X}_{in} = [\mathbf{Y} \parallel \mathbf{X}]$ the concatenation of \mathbf{X} and \mathbf{Y} and \mathbf{X}_{out} the output of the

autoencoder. The conditional Gaussian encoder is described as follows:

$$\mathbf{X}_{in} = [\mathbf{Y} \parallel \mathbf{X}] \quad (3)$$

$$\mathbf{Y}_{E_{H_1}} = \max(0, \mathbf{X}_{in} \mathbf{W}_{E_{H_1}} + \mathbf{B}_{E_{H_1}}) \quad (4)$$

$$\mathbf{Y}_{E_{H_2}} = \max(0, \mathbf{Y}_{E_{H_1}} \mathbf{W}_{E_{H_2}} + \mathbf{B}_{E_{H_2}}) \quad (5)$$

$$\mathbf{Y} = \max(0, \mathbf{Y}_{E_{H_2}} \mathbf{W}_{E_{H_3}} + \mathbf{B}_{E_{H_3}}) \quad (6)$$

$$\mu = [\mathbf{y}_1 \ \mathbf{y}_2], \ \sigma = [\mathbf{y}_3 \ \mathbf{y}_4] \quad (7)$$

$$\mathbf{Z} = \mu + \sigma \cdot \mathcal{N}(0, 1) \quad (8)$$

Subsequently, the conditional Bernoulli decoder is described as follows:

$$\mathbf{Z} = [\mathbf{Y} \parallel \mathbf{Z}] \quad (9)$$

$$\mathbf{Y}_{D_{H_1}} = \max(0, \mathbf{Z} \mathbf{W}_{D_{H_1}} + \mathbf{B}_{D_{H_1}}) \quad (10)$$

$$\mathbf{Y}_{D_{H_2}} = \max(0, \mathbf{Y}_{D_{H_1}} \mathbf{W}_{D_{H_2}} + \mathbf{B}_{D_{H_2}}) \quad (11)$$

$$\mathbf{Y}_{D_{H_2}} = \max(0, \mathbf{Y}_{D_{H_1}} \mathbf{W}_{D_{H_2}} + \mathbf{B}_{D_{H_2}}) \quad (12)$$

$$\mathbf{X}_{out} = \sigma(\mathbf{Y}_{D_{H_1}} \mathbf{W}_{D_{H_2}} + \mathbf{B}_{D_{H_2}}) \quad (13)$$

where $\mathbf{W}_{E_{H_1,2,3}}$, $\mathbf{W}_{D_{H_1,2}}$ and $\mathbf{B}_{E_{H_1,2,3}}$, $\mathbf{B}_{D_{H_1,2}}$ are weighting and bias tensors respectively. For the training of the autoencoder architecture the evidence lower bound (ELBO) error is computed:

$$\mathcal{L} = H(\mathbf{X}_{in}, \mathbf{X}_{out}) + D_{KL}(\mu \parallel \sigma) \quad (14)$$

TABLE I: Summary of Notations

\mathbf{v}_i	Vertex i
\mathbf{c}_i	Centroid of face i
f_i	Face $i, f_i = \{\mathbf{v}_{i1}, \mathbf{v}_{i2}, \mathbf{v}_{i3}\}$
A_i	Face area
n	Number of vertices
n_f	Number of faces
\mathcal{N}_i	Set of neighbouring vertices of vertex i
\mathcal{N}_{f_i}	Set of neighbouring vertices of faces
\mathbf{n}_{c_i}	Normal vector for centroid of face i , $\mathbf{n}_{c_i} \equiv \mathbf{n}_{f_i}$
\mathbf{n}_{f_i}	Normal vector for face i , $\mathbf{n}_{c_i} \equiv \mathbf{n}_{f_i}$
δ_{n_1}	Rotation angle
\mathbf{a}_{n_1}	Rotation axis
\mathbf{a}_c	Arbitrary vector facilitating patch rotation
\mathbf{x}	Hidden layer output
\mathbf{y}	Hidden layer output
\mathbf{b}	Hidden layer bias
\mathbf{z}	Autoencoder output vector
$\mathbf{W}_{EH1,2,3}$	Weighting tensors for the encoder part
$\mathbf{W}_{DH1,2}$	Weighting tensors for the decoder part
$\mathbf{B}_{EH1,2,3}$	Bias tensors for the encoder part
$\mathbf{B}_{DH1,2}$	Bias tensors for the decoder part
$L(\cdot)$	Loss function
$s(\cdot)$	Sigmoid function
$q_D(\mathbf{X}), q_\Phi(\mathbf{X})$	Empirical distributions associated to n training inputs
D_{KL}	Kullback-Leibler Divergence
$\mathbf{H}(\mathbf{a}, \mathbf{b})$	Cross entropy loss
$[\mathbf{A} \mathbf{B}]$	Concatenation of matrices \mathbf{A} and \mathbf{B}
\mathcal{L}	Evidence lower bound (ELBO) error
$\langle \mathbf{a} \mathbf{b} \rangle$	Inner product of \mathbf{a} and \mathbf{b}
\bar{L}_p	Average edge length for 3D mesh
N_B	Number of bilateral filter iterations
N_V	Number of vertex update iterations
$\mathbf{1}$	All ones vector $[1 \ 1 \ 1]$

where $H(\mathbf{X}_{in}, \mathbf{X}_{out})$ is the cross entropy and $D_{KL}(\mu||\sigma)$ the Kullback-Leibler divergence.

B. Training, denoising and post-processing

This section describes the autoencoder training and denoising pipeline, depicted in Figures 1 and 2, and the patch descriptor utilized in the proposed scheme.

1) *Patch descriptor*: The patch descriptor aims to constrain the latent space, to allow for efficient training of the autoencoder. Each patch \mathcal{P}_i is comprised of N topological neighbours of face f_i . Neighbouring faces are sorted using the distance of the face centroid c_i to the face centroid c_j , where f_j is a neighbouring face belonging to patch \mathcal{P}_i , $f_j \in \mathcal{P}_i$. A schematic visualization of the patch descriptor is depicted in Figure 3. Assuming a local coordinate system (Figure 3a), arranging faces constrains the latent space in the z -axis. To further constrain the latent space across x and y axes the patch is rotated by angle δ_{n_i} around rotation axis \mathbf{a}_{n_1} so that:

$$\mathbf{n}_f = \frac{1}{N} \sum_{i \in \mathcal{N}_i} A_i \cdot \mathbf{n}_{c_i} = \mathbf{a}_c \quad (15)$$

where \mathbf{a}_c is a known arbitrarily defined vector. The motivation behind rotating each patch towards the same direction is that it allows efficient training with smaller training sets. Otherwise, we would have to include patches with every possible direction of normals to the training dataset, resulting in large datasets.

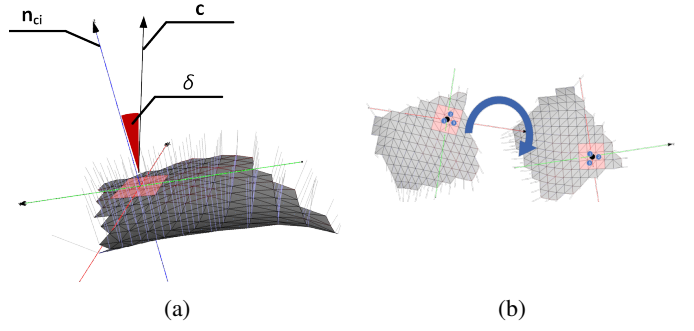


Fig. 3: Example of a patch neighbourhood ($n = 17$), highlighted in orange color, of the face highlighted in cyan color. The normals of the faces are also apparent.




Fig. 4: Visualization of the K-means clustering for a (a) noise-free, and (b) a noisy mesh of the same model.

2) *Training*: The training of the deep network is schematically presented in Figures 1 and 2. The training set contains pairs of noisy and noise-free patches comprised of N neighbouring faces. The corresponding face normals are rotated by δ_{n_i} around rotation axis \mathbf{a}_{n_1} . For the definition of the rotation axis the normals of the noisy patch are used as reference. In order to generate labels for the training set, we perform K-means clustering defining the group centroids for K clusters. The motivation behind applying K-means clustering is that it divides the dataset into groups of patches with high curvature, low curvature, flat areas, and features i.e corners. Thus, different models are trained for each category. Figure 4 presents an example of a 3D mesh and its corresponding noisy version. The K-means clustering of the different surface categories is depicted using different color per different cluster. The coordinates of the normalized normal vectors \mathbf{n}_{c_i} , that comprise patch \mathcal{P}_i , range in $[-1, 1]$. They are transformed to range in $[0, 1]$ by the following equation:

$$\mathbf{n}'_{c_i} = 2 \cdot \mathbf{n}_{c_i} - \mathbf{1} \quad (16)$$

Subsequently, the matrix $\mathbf{N}_i^{3 \times (n+1)}$, consisting of vectors \mathbf{n}'_{c_i} is reshaped to $\mathbf{Z}_i^{3(n+1) \times 1}$. Finally, training is performed with Adam optimizer.

3) *Denoising*: The denoising process is visualized in Figure 1. To use the trained autoencoder for denoising, patches are formed on the noisy mesh. For each patch the average normal \mathbf{n}_f is extracted and the patch is rotated by δ_k so that \mathbf{n}_f is

co-directional to \mathbf{a}_c , to form the input matrix $\mathbf{Z}_i^{3(n+1) \times 1}$. After the autoencoder has generated the filtered output $\mathbf{Z}'_i^{3(n+1) \times 1}$, they are reshaped back to the original form $\mathbf{Z}_i^{3 \times (n+1)}$. The exported filtered normal vector for patch \mathcal{P}_i is the first column of \mathbf{Z}' , and more specifically:

$$\hat{\mathbf{n}}_{c_i}^{3 \times 1} = \mathbf{Z}'_i[:, 0] \quad (17)$$

Finally, each patch is rotated by the opposite angle $-\delta_{n_i}$ and the same axis \mathbf{a}_{n_i} , that they were rotated with in the first place.

4) *Post-processing*: As a final post processing step, we use the bilateral filtering approach according to [12]:

$$\hat{\mathbf{n}}_c = \frac{\sum_{f_j \in \mathcal{P}_i} A_j \mathbf{W}_{1ij} \mathbf{W}_{2ij} \mathbf{n}_{cj}}{\|\mathbf{A}_j \mathbf{W}_{1ij} \mathbf{W}_{2ij} \mathbf{n}_{cj}\|_2} \quad (18)$$

$$\mathbf{W}_{1ij} = \exp\left(-\frac{\|\mathbf{c}_i - \mathbf{c}_j\|^2}{2\sigma_1^2}\right), \quad \mathbf{W}_{2ij} = \exp\left(-\frac{\|\mathbf{n}_{ci} - \mathbf{n}_{cj}\|^2}{2\sigma_2^2}\right) \quad (19)$$

where A_j represents the area of face f_j . Finally, the denoised normals $\hat{\mathbf{n}}_c$ are used to update the vertices according to [14]:

$$\mathbf{v}_i = \mathbf{v}_i + \frac{\sum_{\mathbf{c}_j \in \mathcal{N}_i} \hat{\mathbf{n}}_{cj} (\langle \hat{\mathbf{n}}_{cj} | (\mathbf{c}_j - \mathbf{v}_i) \rangle)}{|\mathcal{N}_i|} \quad (20)$$

where \mathcal{N}_i represents the first-ring area of a vertex \mathbf{v}_i . At this point, it is significant to clarify that we always use the same values for each model, without searching for the ideal parameters per model. More specifically, $\sigma_2 = 0.15$ and for the estimation of σ_1 , we use the following equation:

$$\sigma_{1ij} = \frac{\sum_{\mathbf{c}_j \in \mathcal{N}_i} \|\mathbf{c}_i - \mathbf{c}_j\|_2^2}{|\mathcal{N}_i|} \quad (21)$$

as proposed by [12].

For the rest of the paper, we define N_B the number of bilateral filtering iterations and N_V the number of vertex update iterations that are performed. Furthermore, Gaussian noise will be defined as $N \sim (\mu, \sigma)$, where μ is the mean value, $\sigma = \beta \cdot \bar{L}_p$ is the standard deviation, \bar{L}_p is the average edge length and β a scalar value.

N_V is set to $N_V = 20$, while N_B depends on the noise level. Experimental evaluation showed that for $N \sim (0, 0.1 \cdot \bar{L}_p)$, $N_B = 1$ allows fine-tuning by removing small artifacts. More iterations increase the computational cost, without any additional benefit. Further elaboration, is presented in subsection V-A. Algorithm 1 summarizes the steps of the proposed method. Even though this training process is time-consuming, it takes place only once. Also, in comparison with other data-driven methods [16], the training process is faster, due to the smaller required dataset size.

V. EXPERIMENTAL ANALYSIS AND SIMULATION RESULTS

A. Experimental setup and training

Two different datasets are examined. The first includes meshes originating from the shape repository of the AIM@Shape project [32] with synthetic Gaussian noise. The second utilizes Kinect 2 scans of 3D printed objects provided

Algorithm 1: Data Training and Denoising of 3D Meshes

```
// Training Process
Input : Noisy Dataset of meshes  $\mathcal{M}_p$ , Original Dataset of meshes  $\mathcal{M}'_p \in \mathbb{R}^{n_{fp} \times 3k}$ ,  $\forall p \in \text{dataset}$ ;
Output: Weights  $\mathbf{W}_{E_{H_i}}, \mathbf{W}_{D_{H_i}}, i \in [1, 2, 3]$ ;
1 Estimate centroid normals via Eq. (2);
2 for  $i = 1, \dots, n_{fp}$  do
3   Rotate the normals by angle  $\delta_{n_i}$  around rotation axis  $\mathbf{a}_{n_i}$  via Eq. (15);
4   Normalize the normal vectors in a range of [0,1] via Eq. (16);
5 end
6 Compute labels  $\mathbf{Y}$ 
7 Reshape and create training set  $\mathbf{Z}_i^{3(n+1) \times 1}$  and  $\mathbf{Z}'_i^{3(n+1) \times 1}$ ;
8  $\mathbf{W}_{E_{H_i}}, \mathbf{W}_{D_{H_i}}, i \in [1, 2, 3] = \text{train}(\mathbf{Z}_i, \mathbf{Z}'_i)$ ;
// Denoising Process
Input : Noisy model  $\mathcal{M} \in \mathbb{R}^{n_f \times 3k}$ ;
Output: Denoised model  $\hat{\mathcal{M}} \in \mathbb{R}^{n_f \times 3k}$ ;
9 Estimate centroid normals via Eq. (2);
10 for  $i = 1, \dots, \tilde{n}_f$  do
11   Rotate the normals by angle  $\delta_{n_i}$  around rotation axis  $\mathbf{a}_{n_i}$  via Eq. (15);
12   Normalize the normal vectors in a range of [0,1] via Eq. (16);
13 end
14 Reshape and create the input  $\mathbf{Z}_i^{3(n+1) \times 1}$ ;
15  $\hat{\mathbf{n}}_{ci} = \text{denoising}(\mathbf{Z}_i, \mathbf{W}_{E_{H_i}}, \mathbf{W}_{D_{H_i}}, i \in [1, 2, 3])$ ;
16 Post processing for fine-tuning using bilateral filtering via Eqs. (18)-(19);
17 Reconstruction of the final denoised model using vertices updating via Eq. (20);
```

by Wang et al. [16]. The latter provides noisy scanned outcomes along with ground truth models. To test the denoising capability of our method, we compared our results to guided mesh normal filtering [13], bilateral normal filtering [12], L_0 minimization mesh denoising [15], fast and effective mesh denoising [14], mesh denoising via cascaded normal regression [16] and feature preserving mesh denoising based on graph spectral processing [23].

As an additional comparison, the CVAE part of our pipeline was replaced with traditional autoencoders, referred to as AE. For the latter, a 5-layer deep autoencoder was employed with $N = [256, 128, 64, 128, 256]$ the number of neurons for each layer. An element-wise sigmoid operation succeeds each layer, trained with a mean square error loss function.

1) *Synthetic Gaussian Noise* : Eight meshes were selected for the training of the autoencoder architecture, comprising in total of 1,977,740 patches. Noisy meshes were synthesized by adding Gaussian noise $N \sim (0, 0.1 \cdot \bar{L}_p)$ co-directional to each vertex normal. 1,977,740 training pairs of noisy and noise-free rotated patches were utilized for the training of the autoencoder. Two configurations were tested for patch size, $n = 8$ and $n = 20$ neighbours. $K = 200$ was selected for the K-means clustering of the CVAE. Training was performed with an Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 1e-8$. The training took place for 100 epochs, utilizing an NVIDIA GeForce GTX 1080 graphics card with 8GB VRAM




Fig. 5: Hyper-parameter optimization radar chart

and compute capability 6.1. For the bilateral filtering, we execute only $N_B = 1$ iteration, while for the vertex update operation, we execute $N_B = 20$ iterations. Experimental evaluation showed that for noise level up to $N \sim (0, 0.1 \cdot \bar{L}_p)$ a single bilateral iteration adequately performs fine-tuning, by removing small artifacts, while more iterations increase the computational cost without any additional benefit.

2) *Kinect scans*: Kinect scans were selected from the dataset provided by Wang et al. [16] in order to form 927541 training examples in total. The trained CVAE model was employed to denoise a noisy Kinect scanned model excluded from the training set. The observed noise level of the Kinect scans was computed to $N \sim (-0.18, 1.4 \cdot \bar{L}_p)$, while Figure 7 presents the denoising outcome. Furthermore, different settings were tested to evaluate optimal patch size and number of bilateral iterations. Patch size ranged in $n = 8, 20, 40, 60$, the number of clusters in $K = 6, 10, 50, 100, 200$ and the number of bilateral filtering iterations N_B in $N_B = 0, 1, 4, 8$.

3) *Hyper-parameter optimization*: To define the number of nodes for each layer we performed hyper-parameter optimization. For encoding layers E_1, E_2 and decoding layers D_1, D_2 the number of nodes ranged in $N_{E_1}, N_{E_2}, N_{D_1}, N_{D_2} \in [256, 512, 1024, 2048, 4096]$. The number of clusters was set equal to $K = 200$, the patch size $n = 20$, the learning rate ranged in $lr \in [1e-05, 2e-05, 3e-05]$, and the keep ratio ranged in $kr \in [0.90, 0.95, 0.99]$. We computed the set of parameters that exhibit higher performance, in terms of lower ELBO loss, as shown in Fig. 5: i) Batch size equals 256 ii) keep ratio $kr = 0.99$, iii) learning rate $lr = 3e-05$, iv) decrease ratio $dr = 0.998$, v) $N_{E_1} = 2048$, vi) $N_{E_2} = 2048$, vii) $N_{D_1} = 2048$ and viii) $N_{D_2} = 2048$.

4) *Evaluation models and metrics*: The quality of the reconstructed results is evaluated using a) the Hausdorff distance (HD) which represents the average one-side distance between the reconstructed and the original 3D mesh, b) the metric α which represents the average angle difference between the normals of the ground truth and the reconstructed model c) visualizations which present in different colors the absolute difference between the reconstructed and original meshes.

B. Mesh denoising studies

1) *Evaluation of reconstructed models*: Table II presents the Hausdorff distance and the mean angular difference α of the face normals between original and denoised 3D models correspondingly. In these tables, we use a variety of different initialization approaches and architectures [12]–[16], [23]. More specifically, we deployed two different deep architectures (i.e., AE and CVAE), in two different patch sizes (i.e., with 8 and 20 nearest neighbours (nn)), and with, or without **post-processing step (pp)**. As we can observe, the best performance depends on the model, and none of these approaches is universally the best. Nevertheless, in most of the cases, the CVAE using 8 nearest neighbours seems to have the most stable behaviour. Comparing the reconstructed meshes provided by AE and CVAE, we notice that simple AE gives a smoothed result to the object’s surface, but it negatively affects the preservation of features. On the other hand, CVAE achieves the accurate reconstruction of geometrical features, but the surface of flat areas contains artifacts, as shown in Figure 8. However, this is a problem efficiently tackled by the post-processing step.

Figure 6 presents a visual comparison of the reconstructed models. In this figure, we also provide enlarged details as well as the α metric for easier evaluation. Additionally, Figure 6 illustrates a visualization of the absolute distance and the theta metric between the original and the reconstructed model for each vertex of the meshes. The lowest value (dark blue) denotes that the compared vertices have the same position, in the 3D coordinate system, while a high value (dark red) of the absolute distance denotes that the vertices exhibit high error. Figure 7 presents the denoising result for the Kinect2 scanned models. Our approach accomplishes a lower theta mean value yielding equivalent results with other established data-driven approaches [16].

2) *Impact of patch size, number of clusters and filter parameters*: Figure 9 presents the theta distribution for different settings of selected number of clusters, patch size and bilateral iterations. Purple lines correspond to $N_B = 8$ iterations, blue lines to $N_B = 4$ iterations, green lines to $N_B = 1$ iterations and red lines to $N_B = 0$ iterations (no post processing). As we can observe, 8 iterations significantly improve the result for $N \sim (-0.18, 1.4 \cdot \bar{L}_p)$ noise.

C. Computational complexity evaluation

This subsection presents a comparison of our approach with other methods in terms of computational complexity. To facilitate the performance evaluation we used an open-source implementation in C++ of state-of-the-art methods [12]–[15] available in [33]. To be more specific, the execution was totally performed in C++. For our approach, the autoencoder part of our pipeline is executed in Python TensorFlow, the denoised normal rotation, bilateral normal filtering and vertex update parts in C++. All the evaluation studies took place in a Intel(R) Core(TM) i7-4790 CPU @ 3.60Hz with 32GB of RAM.

As Figure 10 and Table III show, our method is much faster than L_0 minimization [15] and Guided Normal Filtering [13]

TABLE II: Hausdorff distance and metric α (in degree) using disparate approaches of initialization and deep architectures.

	Bilateral [12]	Guided [13]	Fast n Effective [14]	L0 min [15]	Feature aware [23]	Cascaded [16]	CVAE_20_pp	
Carter (100000 F)	7.934 / 0.507	9.16 / 0.602	11.519 / 0.89	7.253 / 0.363	7.209 / 0.56	8.43 / 0.621	5.955 / 0.457	
Pulley (100000 F)	5.476 / 0.365	7.984 / 0.455	8.573 / 0.627	6.341 / 0.255	3.786 / 0.827	6.78 / 0.440	3.591 / 0.305	
Screwdriver (54000 F)	4.191 / 0.003	4.654 / 0.004	6.389 / 0.006	4.829 / 0.003	2.58 / 0.004	3.93 / 0.003	3.161 / 0.003	
	AE				CVAE			
	AE_8_no	AE_8_pp	AE_20_no	AE_20_pp	CVAE_8_no	CVAE_8_pp	CVAE_20_no	CVAE_20_pp
Carter (100000 F)	5.345 / 0.343	5.206 / 0.406	5.159 / 0.471	5.847 / 0.523	5.531 / 0.341	5.364 / 0.39	5.282 / 0.405	5.955 / 0.457
Pulley (100000 F)	4.359 / 0.178	3.371 / 0.219	2.917 / 0.205	2.995 / 0.257	4.562 / 0.248	3.672 / 0.254	3.372 / 0.278	3.591 / 0.305
Screwdriver (54000 F)	4.179 / 0.003	2.923 / 0.003	3.315 / 0.004	3.113 / 0.004	4.112 / 0.002	2.946 / 0.003	3.347 / 0.003	3.161 / 0.003

TABLE III: Execution time, measured in seconds, for presented approaches and noise level $N \sim (0, 0.1 \cdot \bar{L}_p)$

	Bilateral normal [12]	Guided normal [13]	Fast effective [14]	L0 min [15]	Feature aware [23]	CVAE 20 pp	CVAE 8 pp
Sculpt (3669V,7342F)	0.1082	0.6465	0.0591	3.5884	0.26743	0.0772	0.0754
Trimmed star (5192V, 10384F)	0.1529	0.9843	0.0869	4.2748	0.41393	0.0995	0.0995
Rocker Arm (9413V,18826F)	0.3242	2.0561	0.1804	11.1609	1.10021	0.1642	0.1617
Chinese Lion (50000 V, 100000F)	2.0508	21.6360	1.5792	110.6100	16.24114	0.9872	0.9624
Gear (250000V,500000F)	8.5630	221.1910	5.7120	2512.2500	180.77456	3.8858	3.7505

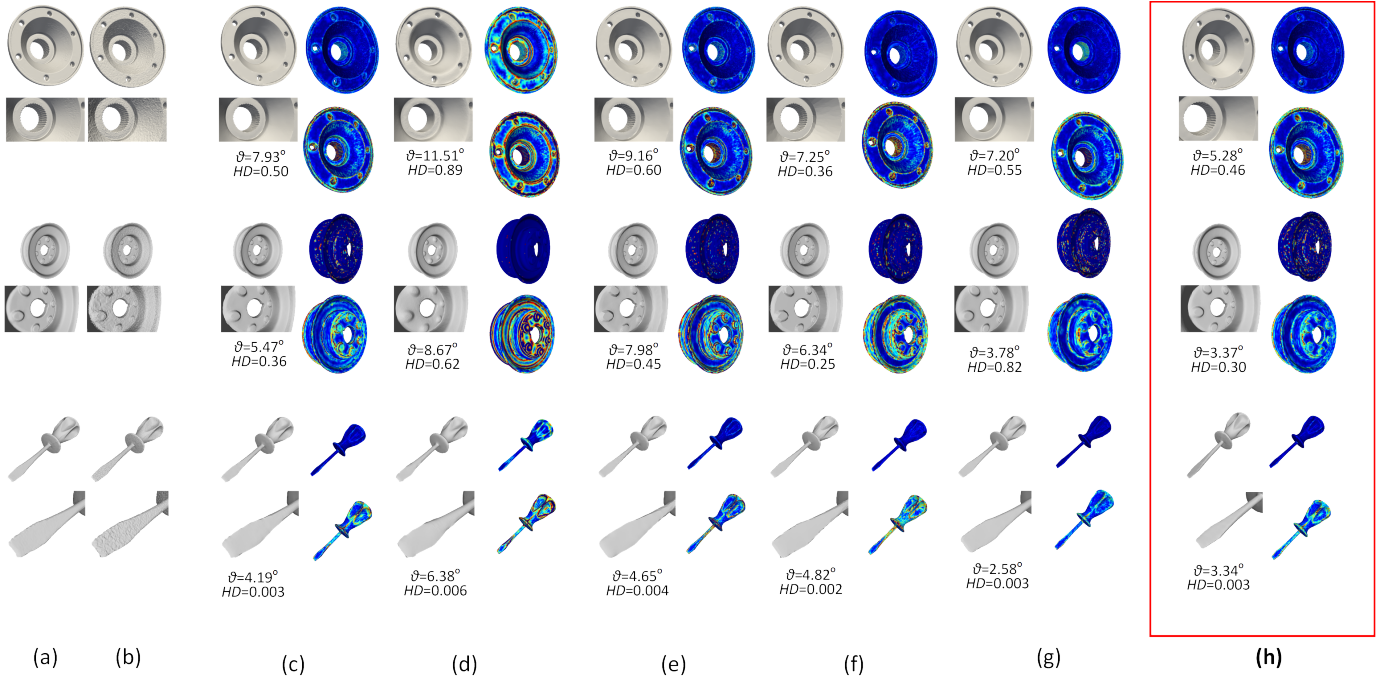


Fig. 6: Denoising results and normal angle difference visualization between reconstructed and original 3D model with Gaussian noise $N \sim (0, 0.1 \cdot \bar{L}_p)$. For each model absolute distance (upper colored mesh) and theta distribution (lower colored mesh) are presented. (a) original mesh, (b) noisy mesh, (c) fast and effective [14], (d) bilateral normal filtering [12], (e) L_0 minimization [15], (f) guided normal filtering [13], (g) Cascaded mesh denoising (h) **our approach**.

and traditional bilateral normal filtering [12]. Compared to fast and effective mesh denoising [14], our method is slower in small models but becomes faster as the number of faces increases. Execution time measurements presented in Table III were computed as the mean value of 10 repetitions. In the case of Rocker Arm counting 18826 faces, our method outperforms all the other approaches. We attribute this observation to the autoencoder complexity. Denoising requires $O(1)$ operations per face removing a large portion of the computational cost.

D. Impact of parallelization

As subsection IV-B3 highlights, the input vector $\mathbf{Z}_i^{3(n+1) \times 1}$ contains the normal vector coordinates of n neighbouring faces for a single patch. The matrix formulated for all patches can be expressed as $\mathbf{Z}^{n_f \times 3(n+1)}$ where n_f is the number of faces for the processed model. Each patch \mathbf{Z}_i is being processed separately through the same processing pipeline. Tensorflow already parallelizes this process and allows to control the number of used CPUs or GPUs. To further elaborate, a noisy model consisting of 100K faces, specifically the "carter" model, was denoised to measure the execution time for the autoencoder part. Six different settings were examined, namely, i) 1 CPU

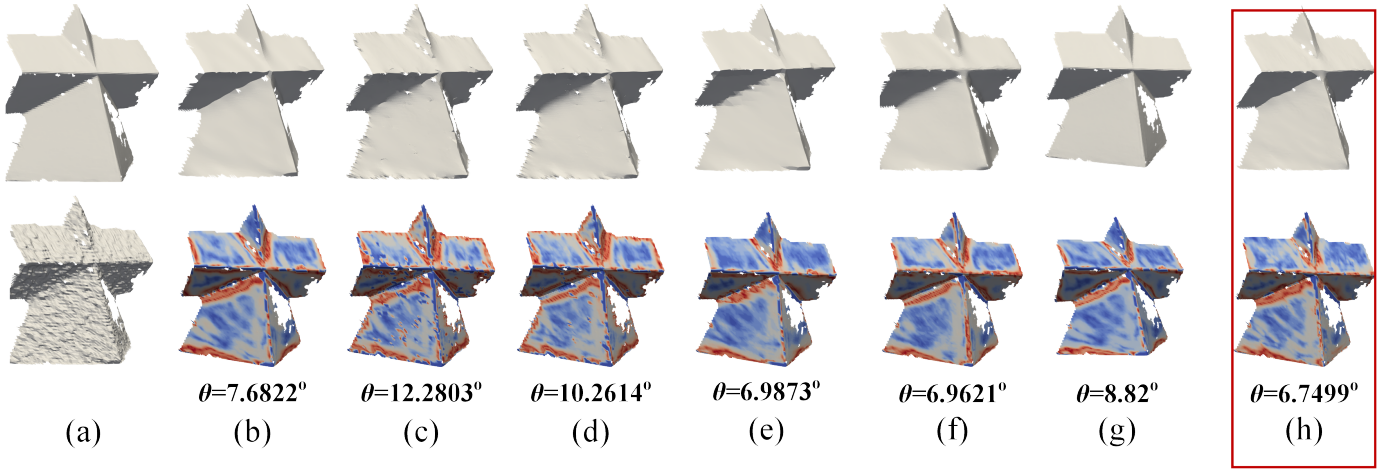


Fig. 7: Denoising results and normal angle difference visualization between reconstructed and original Kinect scan with noise level $N \sim (-0.18, 1.4 \cdot \bar{L}_p)$. (a) original mesh (up) and noisy mesh (down), (b) fast and effective [14], (c) bilateral normal filtering [12], (d) L_0 minimization [15], (e) guided normal filtering [13], (f) Cascaded mesh denoising, (g) Feature aware denoising [23] and (h) **our approach**.




Fig. 8: Denoising results using: (a) AE, (b) CVAE.
Patch size (number of neighbours)




Fig. 9: Evaluation of theta distribution for different settings of the number of clusters, the number of patch neighbours and bilateral filter iterations.

core, ii) 2 CPU cores, iii) 4 CPU cores, iv) 6 CPU cores, v) 8 CPU cores and vi) GPU only. For each different setting, 20 repetitions were performed. Figure 11 presents boxplots summarizing the execution times distribution for each setting.

E. Defect detection in an industrial setting

Saliency maps [34] are essential tools for reliable, accurate and computationally efficient 3D representations, by simplifying the representation of physical. Figure 12 presents visual




Fig. 10: Performance evaluation.




Fig. 11: CVAE execution times for 100K faces model and different settings

confirmation, that proper feature preserving denoising can facilitate defect detection in an industrial setting. The first row presents the 3D mesh and the second row, the result of the defect detection process. The color map is related to the Hausdorff distance of each mesh to the ground truth geometry presented in subfigure 12a. The first column (a,e) presents the original object, the second column (b,f) presents the same object with surface defects that may have originated from the manufacturing process. As Figure 12g reveals, noise prohibits




Fig. 12: Defect detection with saliency maps. a,e) Noise-free mesh b,f) Noise-free mesh with surface defects c,g) Noisy mesh d,h) Denoised mesh

the detection of the defects. The third column (c,g) presents the same 3D mesh with defects and Gaussian noise that may have originated from scanning. Finally, the fourth column presents the denoised object, where the outcome of denoising facilitates accurate detection.

VI. DISCUSSION

In this work, we presented a fast data-driven denoising approach, applying conditional variational autoencoders to filter the normals of noisy 3D mesh patches. These patches are modelled by a scale, translation and rotation invariant patch descriptor exploited during the learning process. A computationally light post-processing step is employed only for fine-tuning purposes. Extensive evaluation studies verify the effectiveness of the proposed method, as compared to other robust and well-known state-of-the-art approaches. We summarize the benefits of the proposed approach in the following points. i) The network can localize since, training and inference are performed in a sliding patch setup. The filtered face normal vector is generated by providing a local neighbouring region as input. ii) It requires a relatively small training set since we employ a preprocessing approach that restricts the input data space. iii) It is fully parallelizable. iv) Evaluation studies indicate that it demonstrates lower complexity and execution times than other non-data-driven state-of-the-art approaches. v) It can be utilized in industrial applications for denoising of dense objects with features such as corners and edges. vi) It is parameter-free since every used parameter is predefined and the user does not need to search for optimal values per model. Furthermore, our study aims to contribute to the field of geometric deep learning where the sampling of the latent space is nonuniform, on contrary to image processing or voxelized versions of 3D volumes. The proposed approach removes the noise from noisy 3D meshes, given that the deep architecture is trained with the same noise distribution. This could prove extremely beneficial for the fast denoising of meshes generated from a particular scanning device operated at a manufacturing production line. However, different levels or types of noise require a different training process.

We should also highlight that the size of the deep neural network has an impact on the denoising performance, in terms of execution times and energy efficiency. As a future step, model compression and acceleration strategies are yet to be investigated. They would allow a smaller network size and lower execution times. Further reduction of execution times could originate from the removal of post-processing steps. Specifically, patches are appropriately rotated and clustered. The bilateral filter copes with problems or irregularities, while extensive evaluation studies reveal that the required number of bilateral iterations is proportional to the noise level. Training the CVAE so that no post-processing is required could improve performance. Furthermore, employing robust clustering, to effectively group patches with similar distributions of normal coordinates, could efficiently reduce the size of training set or boost reconstruction accuracy.

REFERENCES

- [1] Y. Lu, K. C. Morris, and S. Frechette, "Current standards landscape for smart manufacturing systems," *National Institute of Standards and Technology, NISTIR*, vol. 8107, p. 39, 2016.
- [2] NIST. (2014) Smart manufacturing operations planning and control program. [Online]. Available: <https://www.nist.gov/programs-projects/smart-manufacturing-operations-planning-and-control-program>
- [3] D. P. Agrawal, B. B. Gupta, H. Wang, X. Chang, S. Yamaguchi, and G. M. Perez, "Guest editorial deep learning models for industry informatics," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 7, pp. 3166–3169, 2018.
- [4] L. Li, K. Ota, and M. Dong, "Deep learning for smart industry: Efficient manufacture inspection system with fog computing," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 10, pp. 4665–4673, 2018.
- [5] R. Iqbal, T. Maniak, F. Doctor, and C. Karyotis, "Fault detection and isolation in industrial processes using deep learning approaches," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 5, pp. 3077–3084, 2019.
- [6] J. Wang, Y. Ma, L. Zhang, R. X. Gao, and D. Wu, "Deep learning for smart manufacturing: Methods and applications," *Journal of Manufacturing Systems*, vol. 48, pp. 144–156, 2018.
- [7] S. von Enzberg and A. Al-Hamadi, "A multiresolution approach to model-based 3-d surface quality inspection," *IEEE Transactions on Industrial Informatics*, vol. 12, no. 4, pp. 1498–1507, 2016.
- [8] P. Fraga-Lamas, T. M. Fernández-Caramés, Ó. Blanco-Novoa, and M. A. Vilar-Montesinos, "A review on industrial augmented reality systems for the industry 4.0 shipyard," *IEEE Access*, vol. 6, pp. 13 358–13 375, 2018.
- [9] A. 3D. (2019, dec) How one foundry uses "crazy efficient" artec eva to slash up to 75Available: <https://www.artec3d.com/cases/eva-cast-starring-role-leading-foundry>
- [10] H. Aanaes, E. Nielsen, and A. B. Dahl, "Autonomous surface inspection of wind turbine blades for quality assurance in production," in *9th European Workshop on Structural Health Monitoring, EWSHM 2018*, 2018.
- [11] Artec3D. (2019, dec) Industrial 3d models, scanned using artec 3d technology. [Online]. Available: <https://www.artec3d.com/3d-models/industrial-design-and-manufacturing>
- [12] Y. Zheng, H. Fu, O. K.-C. Au, and C.-L. Tai, "Bilateral normal filtering for mesh denoising," *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 10, pp. 1521–1530, 2011.
- [13] W. Zhang, B. Deng, J. Zhang, S. Bouaziz, and L. Liu, "Guided mesh normal filtering," in *Computer Graphics Forum*, vol. 34, no. 7. Wiley Online Library, 2015, pp. 23–34.
- [14] X. Sun, P. Rosin, R. Martin, and F. Langbein, "Fast and effective feature-preserving mesh denoising," *IEEE transactions on visualization and computer graphics*, vol. 13, no. 5, pp. 925–938, 2007.
- [15] L. He and S. Schaefer, "Mesh denoising via l 0 minimization," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 4, p. 64, 2013.
- [16] P.-S. Wang, Y. Liu, and X. Tong, "Mesh denoising via cascaded normal regression," *ACM Trans. Graph.*, vol. 35, no. 6, pp. 232–1, 2016.

- [17] F. Bosché, “Automated recognition of 3d cad model objects in laser scans and calculation of as-built dimensions for dimensional compliance control in construction,” *Advanced engineering informatics*, vol. 24, no. 1, pp. 107–118, 2010.
- [18] H. Srinivasan, O. L. Harrysson, and R. A. Wysk, “Automatic part localization in a cnc machine coordinate system by means of 3d scans,” *The International Journal of Advanced Manufacturing Technology*, vol. 81, no. 5-8, pp. 1127–1138, 2015.
- [19] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr, “Implicit fairing of irregular meshes using diffusion and curvature flow,” in *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. Citeseer, 1999, pp. 317–324.
- [20] H. Zhang, O. Van Kaick, and R. Dyer, “Spectral mesh processing,” in *Computer graphics forum*, vol. 29, no. 6. Wiley Online Library, 2010, pp. 1865–1894.
- [21] P.-S. Wang, X.-M. Fu, Y. Liu, X. Tong, S.-L. Liu, and B. Guo, “Rolling guidance normal filter for geometric processing,” *ACM Transactions on Graphics (TOG)*, vol. 34, no. 6, p. 173, 2015.
- [22] S. Fleishman, I. Drori, and D. Cohen-Or, “Bilateral mesh denoising,” in *ACM transactions on graphics (TOG)*, vol. 22, no. 3. ACM, 2003, pp. 950–953.
- [23] G. Arvanitis, A. S. Lalos, K. Moustakas, and N. Fakotakis, “Feature preserving mesh denoising based on graph spectral processing,” *IEEE transactions on visualization and computer graphics*, vol. 25, no. 3, pp. 1513–1527, 2019.
- [24] H. Zhang, C. Wu, J. Zhang, and J. Deng, “Variational mesh denoising using total variation and piecewise constant function space,” *IEEE transactions on visualization and computer graphics*, vol. 21, no. 7, pp. 873–886, 2015.
- [25] M. Wei, L. Liang, W.-M. Pang, J. Wang, W. Li, and H. Wu, “Tensor voting guided mesh denoising,” *IEEE Transactions on Automation Science and Engineering*, vol. 14, no. 2, pp. 931–945, 2017.
- [26] J. Wang, J. Huang, F. L. Wang, M. Wei, H. Xie, and J. Qin, “Data-driven geometry-recovering mesh denoising,” *Computer-Aided Design*, 2019.
- [27] W. Zhao, X. Liu, Y. Zhao, X. Fan, and D. Zhao, “Normalnet: Learning based guided normal filtering for mesh denoising,” *arXiv preprint arXiv:1903.04015*, 2019.
- [28] K. Sarkar, K. Varanasi, and D. Stricker, “3d shape processing by convolutional denoising autoencoders on local patches,” in *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2018, pp. 1925–1934.
- [29] G. E. Hinton, “Deep belief networks,” *Scholarpedia*, vol. 4, no. 5, p. 5947, 2009.
- [30] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [31] K. Sohn, H. Lee, and X. Yan, “Learning structured output representation using deep conditional generative models,” in *Advances in neural information processing systems*, 2015, pp. 3483–3491.
- [32] B. Falcidieno, “Aim@ shape project presentation,” in *Proceedings Shape Modeling Applications, 2004*. IEEE, 2004, p. 329.
- [33] B. D. Wangyu Zhang. (2015) Mesh denoising ui. [Online]. Available: <https://github.com/bldeng/GuidedDenoising>
- [34] M. C. M. Lau and K. Dev, “Tactile mesh saliency: a brief synopsis,” 2016.