



# Selfredepth

## Self-supervised real-time depth restoration for consumer-grade sensors

Alexandre Duarte<sup>1</sup> · Francisco Fernandes<sup>2</sup> · João M. Pereira<sup>1,2</sup> · Catarina Moreira<sup>2,4</sup> · Jacinto C. Nascimento<sup>1,3</sup> · Joaquim Jorge<sup>1,2</sup>

Received: 14 September 2023 / Accepted: 3 June 2024 / Published online: 4 July 2024  
© The Author(s) 2024, corrected publication 2024

### Abstract

Depth maps produced by consumer-grade sensors suffer from inaccurate measurements and missing data from either system or scene-specific sources. Data-driven denoising algorithms can mitigate such problems; however, they require vast amounts of ground truth depth data. Recent research has tackled this limitation using self-supervised learning techniques, but it requires multiple RGB-D sensors. Moreover, most existing approaches focus on denoising single isolated depth maps or specific subjects of interest highlighting a need for methods that can effectively denoise depth maps in real-time dynamic environments. This paper extends state-of-the-art approaches for depth-denoising commodity depth devices, proposing SelfReDepth, a self-supervised deep learning technique for depth restoration, via denoising and hole-filling by inpainting of full-depth maps captured with RGB-D sensors. The algorithm targets depth data in video streams, utilizing multiple sequential depth frames coupled with color data to achieve high-quality depth videos with temporal coherence. Finally, SelfReDepth is designed to be compatible with various RGB-D sensors and usable in real-time scenarios as a pre-processing step before applying other depth-dependent algorithms. Our results demonstrate our approach's real-time performance on real-world datasets shows that it outperforms state-of-the-art methods in denoising and restoration performance at over 30 fps on Commercial Depth Cameras, with potential benefits for augmented and mixed-reality applications.

**Keywords** Deep learning · Self-supervised learning · Image denoising · Image reconstruction · RGB-D sensors

**Mathematics Subject Classification** 68T07 · 94A08

---

✉ Joaquim Jorge  
jaj@inesc-id.pt

Alexandre Duarte  
alexandre.a.duarte@tecnico.ulisboa.pt

Francisco Fernandes  
francisco.fernandes@tecnico.ulisboa.pt

João M. Pereira  
jap@inesc-id.pt

Catarina Moreira  
catarina.pintomoreira@uts.edu.au

Jacinto C. Nascimento  
jacinto.nascimento@tecnico.ulisboa.pt

<sup>1</sup> Instituto Superior Técnico, Universidade de Lisboa (IST-UL), 1000-029 Lisbon, Portugal

<sup>2</sup> Instituto de Engenharia de Sistemas e Computadores, Investigação e Desenvolvimento (INESC-ID), 1000-029 Lisbon, Portugal

<sup>3</sup> Institute for System and Robotics (ISR), Instituto Superior Técnico, Universidade de Lisboa (IST-UL), 1049-001 Lisbon, Portugal

<sup>4</sup> Human Technology Institute, University of Technology Sydney, Sydney, Australia

## 1 Introduction

Depth information is pivotal in many applications, from digital entertainment to virtual and augmented reality [21]. It is the backbone for digital object and environment modeling [8, 42] and cost-effective motion capture solutions [18].

Pose estimation derived from depth data finds utility in diverse fields such as physiotherapy [5, 17], video surveillance [34, 63], and human–computer interaction [46]. Depth data also aids autonomous navigation [15] and enhances security measures through facial recognition [43].

Consumer depth devices, often employing low-cost LiDAR, structured light, or time-of-flight technologies, are instrumental in these applications. Among these, the Microsoft Kinect v2 stands out for its balance of quality, availability, and affordability. However, consumer-grade sensors like Kinect v2 still grapple with noisy and incomplete data issues.

Efforts to address these quality issues span traditional smoothing techniques to data-driven machine learning algorithms. Many adopt supervised learning with neural networks, training models on noisy-clean data pairs  $(\hat{x}, y)$  to minimize empirical risk.

However, acquiring clean training data is non-trivial. Recent attention has thus shifted towards self-supervised techniques, such as Noise2Noise [27], which leverages noisy-noisy data pairs  $(\hat{x}, \hat{y})$  for training, and minimizing the cost function  $g(\theta) = \operatorname{argmin}_{\theta} \sum_i \mathcal{L}(f_{\theta}(\hat{x}_i), \hat{y}_i)$ , where the network  $f_{\theta}(\hat{x}_i)$  is parameterized by  $\theta$ .

Despite their efficacy in various domains, self-supervised methods for depth data restoration remain underexplored, largely due to the intricate noise patterns in consumer-grade sensors.

Our paper introduces **SelfReDepth (SRd)**, a novel self-supervised, real-time depth data restoration technique optimized for the Kinect v2. SelfReDepth introduces a convolutional autoencoder architecture inspired by U-Net, specifically designed to process sequential depth frames efficiently. This design choice directly responds to the need for maintaining temporal coherence in dynamic scenes, a gap often left unaddressed by traditional single-frame denoising approaches. Furthermore, SelfReDepth incorporates RGB data into the depth restoration process as an innovative way to enhance the accuracy of inpainting missing pixels by providing contextual color information. This method significantly improves the restoration quality by providing additional context that depth data alone lacks. Our contributions are fourfold: (1) We employ a convolutional autoencoder with an architecture akin to U-Net [47] to process sequential frames. (2) Our method achieves real-time performance and temporal coherence

by adopting a video-centric approach. (3) We incorporate RGB data to guide an inpainting algorithm during training, enhancing the model’s ability to complete missing depth pixels. (4) Our approach maintains a 30 fps real-time rate while outperforming state-of-the-art techniques.

## 2 Background and related work

In recent years, depth-sensing technology has emerged as a pivotal tool in various applications, from gaming to augmented reality and robotics. The promise of capturing the third dimension, depth, has opened up new horizons in computer vision, augmented reality, and human–computer interaction. Next, we introduce some concepts and methodologies related to the present work.

**Denoising vs. inpainting:** The distinction between *denoising* and *inpainting* is important to be stressed, as these terms will be used throughout this work constituting important stages of the proposed methodology. Denoising and inpainting are two core image processing problems. As the name suggests, denoising removes noise from an observed noisy image, while inpainting aims to estimate missing image pixels. Both denoising and inpainting are inverse problems: the common goal is to infer an underlying image from incomplete/imperfect observations. Formally, in both problems the observed image  $\mathbf{Y} \in \mathbb{R}^{M' \times N'}$  is modeled as  $\mathbf{Y} = \mathcal{F}(\mathbf{X}) + \eta$  where  $\mathbf{X} \in \mathbb{R}^{M \times N}$  is the unknown (original) image and  $\eta$  is the observed noise. The difference between the denoising and the inpainting emerges from the mapping  $\mathcal{F} : \mathbb{R}^{M \times N} \mapsto \mathbb{R}^{M' \times N'}$  that expresses a linear degradation operator that could represent a convolution or a masking process. In concrete, the denoising process means that  $\mathcal{F}$  is an *identity projector*, having  $\mathcal{F} = \mathbf{H}$ , such that  $\mathbf{H} = \mathbf{I}$ , (with  $\mathbf{I}$  the identity matrix). In the inpainting process,  $\mathcal{F}$  is a *selection operator*. In practice, this corresponds to having the same  $\mathbf{H}$  as before. However, this matrix only contains a subset of the rows of  $\mathbf{I}$ , accounting for the loss of pixels.

Having formally defined the two concepts, we stress that both are used in restoration problems, as we propose in this work. It follows that, in this paper, two major contributions are offered for restoration, concretely: (i) a new denoising method. Contrasting with Noise2Noise [27] that applies denoising in traditional images, we extend the framework to depth images that requires a new learning strategy to handle depth information (see top branch in Fig. 2 and Sect. 3.2), and (ii) inpainting approach where we integrate a new two-stage pipeline comprising a RGB-Depth registration and a fast marching method stages (see bottom branch in Fig. 2).

**Problems with low-quality depth:** Despite all the progress made in in-depth sensing hardware for consumer devices,

depth cameras (such as the Kinect v2) still suffer from many of the same problems that previous iterations also did, namely, noisy measurements and depth holes [38].

These depth holes typical of time-of-flight (ToF) devices have multiple causes [20, 50] including: (1) measuring regions that are outside the distance range of the sensor, (2) highly reflective objects in the scene, (3) measurements near the edges of the camera's field-of-view (FoV). Smaller holes, on the other hand, can appear in one of two types: (1) isolated points caused by physical and lighting interferences on the sensor, (2) thin outlines around objects due to the scattering of infrared rays at shallow angles and sharp edges.

Besides missing depth values, the measurement inaccuracy in the successfully captured points is a concerning issue, leading to noisy depth maps. The noise produced by Kinect v2 is significantly less severe than in the previous generation. Nevertheless, it is still very much present. Experimental analysis [26, 54, 61] has shown that there is a direct correlation between the noise observed in Kinect v2's depth maps and various physical factors, including distance, angle, material, color, warm-up time, to quote a few. Furthermore, there is a general consensus [1, 12, 25] that this noise can be described as the sum of two different sources of noise: (1) random noise, associated with pixel-based local distortions caused by physical factors like color and others mentioned above, (2) systematic bias, associated with the wiggling error that radially increases as the measurements get closer to the edges of the sensor's FoV.

**Deterministic denoisers**, or manual denoising algorithms that do not rely on machine learning were the first noise reduction techniques to be developed targeting depth data [14]. These can generally be divided into three main categories: (i) *filter-based* denoisers, (ii) *outlier removal* techniques, and (iii) *calibration* methods.

Filter-based denoisers work by applying smoothing and sharpening filters, such as bilateral filters [37, 62], joint bilateral filters [9, 10], anisotropic filters [35] and zero block filters [35], to leverage spatial pixel neighborhoods through sliding pixel windows (i.e. kernels). The preservation of edge sharpness is particularly difficult to achieve using filters; thus, some denoisers introduce specialized techniques, such as RGB-D alignment [10] and contextual image partitioning [9].

Other works, instead, focus exclusively on removing incorrect or low-quality depth points rather than correcting them. This can be seen, for instance, in Ref. [12] for cleaning hand depth scans and in Ref. [57] for cleaning body scans, later merged to form a complete body point cloud. Finally, some works tackle the denoising problem from a calibration standpoint, focusing on alleviating systematic errors affecting consumer-grade sensors by fitting planes or splines to the

raw measurements [25, 26] or generating specialized noise correction maps [1].

**Self-supervised denoisers:** Noise2Noise [27] pioneered self-supervised image denoising, showing that a denoising model trained with only noisy data can achieve quality results on par with supervised learning strategies. The shift in the learning method from supervision to self-supervision resides primarily in the training data. Specifically, Noise2Noise [27] uses input-target pairs of the form  $(\hat{x}, \hat{y}) = (x + n_1, x + n_2)$ , where  $x$  is the base signal (the undamaged data that we want to uncover) and  $n_1$  and  $n_2$  are two independent noise instances following the same statistical distribution. From the above, the Noise2Noise strategy differs from the noise-clean data pairs  $(\hat{x}, y)$  used in supervised learning and has the advantage of not requiring clean target images.

Nonetheless, Noise2Noise [27] has some data limitations, encouraging subsequent works to propose further improvements. Towards this challenge, [4] proposes two data permutation techniques to increase the number of noisy training pairs. On the other hand, Noise2Void [23] eliminates the need for quasi-similar input-target pairs (not always easy to obtain) by training the model to predict a central pixel using noisy-void training pairs  $(\hat{x}, -)$  and a blind-spot mask to avoid learning the identity. Noise2Self [2] later expanded on this by proposing a more generalized model.

Going further, some works, namely Probabilistic Noise2Void [24], SURE [39], Noisier2Noise [41] and NoiseBreaker [28], managed to improve denoising for specific distributions. On the other hand, GAN2GAN [6] combines a generative model, Self2Self [45] introduces training with a single data sample, and GainTuning [40] proposes an ever-adapting model.

**Spatio-temporal denoisers** are an extension of image denoising where the coherence of temporal locality is also considered to provide visual continuity in the final denoised videos. Likewise, clean data may also not be easy to obtain for this task, and thus, blind training comes with great interest. A simple multi-frame self-supervised strategy for denoising can be extrapolated from Noise2Stack [44], in which a self-supervised approach is proposed to denoise MRI data using adjacent sets from a stack of layered MRI brain scans.

Self-supervised denoising techniques targeting color videos have also been developed, using the multi-frame input concept described in Noise2Stack [44] combined with additional self-supervision techniques. Multi-Frame-to-Frame (MF2F) [13] (see Fig. 1) takes the FastDVDNet [52] supervised video denoising network, composed by cascaded U-Net [47] autoencoders, and applies its own self-supervised loss.



Fig. 1 Multi-frame-to-frame (MF2F) [13] architecture with distinct strategies from the training and inference steps

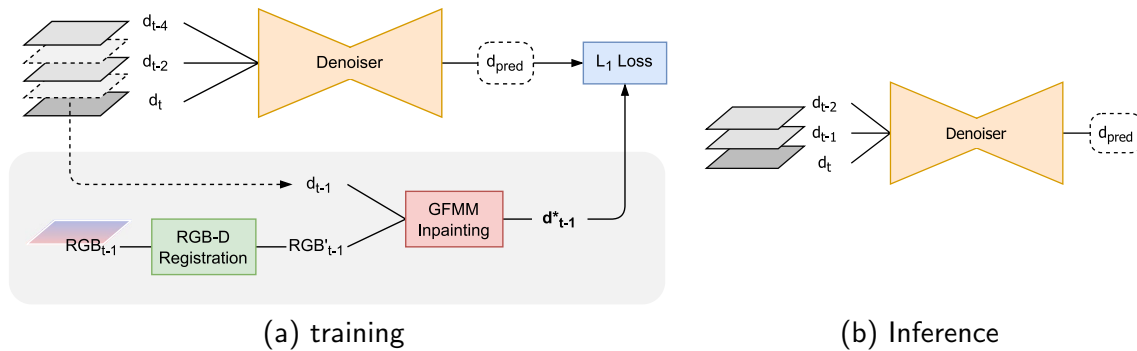


Fig. 2 Full overview of SelfReDepth's architecture

Similarly, UDVD [48] uses a cascaded structure akin to FastDVDNet [52] but performs the network pass 4-fold, each with the input frames at a different rotation (0°, 90°, 180° and 270°). The four outputs generated, one for each of the four rotations, are then rotated back to 0° and combined to form the final production.

**Depth completion:** Alongside inaccurate data points, low-quality depth maps also suffer from missing or invalid data. Depth completion, also known as hole-filling, is a well-known and vastly researched area that falls under the umbrella of image inpainting [32, 59, 60].

The effect of using Noise2Noise and similar algorithms over images with missing data without any prior inpainting is that the majority of depth holes remain untreated in the final images, and even in methods that deal with multiple consecutive frames, there is insufficient data to fill these gaps in most cases.

As in-depth denoising, depth completion has been approached using traditional and deep neural methodologies. Traditional techniques typically rely on either filtering algorithms, which classify the holes and apply dedicated filters, such as PDJB, DJBF, and FCRN, or boundary-extending algorithms, based on FMM or the Navier–Stokes equation [3].

The fast marching method (FMM) inpainting, in particular, was originally proposed for color image inpainting [53] and works by progressively shrinking the boundaries of hole

regions inwards, until all pixels have been filled, using the equation

$$I(p) = \frac{\sum_{q \in N(p)} w(p, q) \cdot [I(q) + \nabla I(q) \cdot (p - q)]}{\sum_{q \in N(p)} w(p, q)} \quad (1)$$

where  $p$  is the pixel being inpainted,  $N(p)$  is a neighborhood pixels of  $p$ ,  $w(p, q)$  is a function that determines how much pixel  $q$  contributes to the inpainting of  $p$ ,  $I$ , and  $\nabla I$  represent the image and discrete gradient of the image, respectively. This algorithm was extended to depth completion, introducing improvements like the use of aligned color as a guiding factor for the weight function and to define the order of computations [33], and the use of a pixel-wise confidence factor [30].

Sparse depth maps, generally captured with LiDAR sensors, suffer especially from large patches of missing depth data and have particular time limitations, as they are commonly linked with autonomous driving. Thus, more advanced techniques have been developed, relying on both supervised [29] and self-supervised [11, 16, 36] deep convolutional neural networks assisted by color information to fill large depth gaps.

### 3 Our approach

Building upon recent advancements in self-supervised data denoising research, the proposed SReD offers a novel approach for denoising and inpainting low-quality depth maps. It leverages the flexibility and adaptability of deep learning models while eliminating the need for reference data - a highly desirable feature also found in deterministic denoisers. SReD was designed with a specific practical use case in mind, incorporating several additional requirements during its design and development. Specifically, our technique aims to: (i) denoise and restore as much of the initial depth maps as possible, (ii) operate with a single RGB-D device, (iii) facilitate direct sensor data streaming, and (iv) strive for temporal coherence and real-time performance.

Naturally, these requirements posed challenges that influenced the architectural decisions. For example, to achieve depth video denoising with temporal coherence, it is logical to design a method that utilizes multiple sequential frames, similar to MF2F [13]. However, given real-time constraints, only frames up to the most recent one are considered. This contrasts with MF2F, which incorporates two subsequent frames at the cost of adding considerable lag. Additionally, an architecture with faster inference is preferable over a more complex one to meet real-time performance criteria.

#### 3.1 SelfReDepth's architecture

SReD's architecture, particularly its neural network and learning method, takes inspiration from previous self-supervised denoisers, mainly Noise2Noise [27], MF2F [13] and Noise2Stack [44], adapting their proposed denoising models to the specific case of online denoising in depth map sequences. As depicted in Fig. 2, the architecture differs between its training and inference stages. The model uses a dilated input during the training stage, as proposed in MF2F [13]. The autoencoder is trained with noisy pairs  $(\hat{x}, \hat{y}) = ([d_{t-4}, d_{t-2}, d_t], d_{t-1})$ , where  $d_{t-k}$  is the depth frame at time instant  $t - k$ , with  $k \in \{0, 1, 2, 4\}$ . This technique improves the denoising results during inference and prevents the network from learning the identity by hiding frame  $d_{t-1}$  from the input. Since  $d_t$  and  $d_{t-1}$  are consecutive time frames, it is also plausible to assume they are similar in content while having different instances of noise, making them a suitable image pair for the Noise2Noise-style training.

Moreover, noisy depth frames frequently have regions persistently composed of depth holes in both the input and target frames, making these regions impossible to "denoise" using a standalone denoising network. As such, during training, the target frame  $d_{t-1}$  is inpainted with an FMM

inpainting algorithm guided by the registered color frame  $RGB_{t-1}$ , providing a way for the denoiser to learn how to fill depth-holes.

In summary, SReD's training architecture has two distinct main blocks, depicted in Fig. 2a: (i) a denoising convolutional autoencoder with dilated input; and (ii) a target generation pipeline responsible for creating inpainted targets. During inference (Fig. 2b), the target generation is removed, contributing to faster performance and the denoiser shifts to non-dilated input (i.e., taking the frames  $[d_{t-2}, d_{t-1}, d_t]$ ), estimating a denoised/inpainted instance of the frame  $d_t$ .

#### 3.2 Target generation

The denoiser requires a learning strategy to handle depth holes. To achieve this, SReD generates target frames through deterministic inpainting. This deterministic approach consists of two stages: (i) computing a registered RGB image and (ii) using the previous result to apply guided inpainting to the damaged depth frame. The selection of this strategy rests on three primary reasons: (1) The prevalence of depth holes in general consumer depth data is sufficiently low for a deterministic approach to yield acceptable inpainting results. (2) It avoids the need for reference data. (3) From a temporal performance perspective, it only introduces computational time during training.

**RGB-D registration:** RGB-D devices collect color and depth with physically separate sensors/cameras, and often also different resolutions and FoV. Therefore, aligning the RGB and depth frames simultaneously captured must be done with a registration algorithm and requires acquiring the extrinsic and intrinsic parameters of the device, namely:

- Focal length  $f_d = [f_{d,x} \ f_{d,y}]^T$  and principal point  $c_d = [c_{d,x} \ c_{d,y}]^T$  of the depth/IR sensor,
- Focal length  $f_{rgb} = [f_{rgb,x} \ f_{rgb,y}]^T$  and principal point  $c_{rgb} = [c_{rgb,x} \ c_{rgb,y}]^T$  of the RGB sensor,
- Rotation matrix  $R$ , which encodes the rotation from the RGB sensor view to the depth/IR sensor,
- Translation vector  $T$  translates from the RGB sensor's position to the depth/IR sensor's position.

Following [64], RGB-D registration is performed through a series of coordinate transformations that map depth values captured from the depth sensor's point-of-view to color values in the RGB camera's point-of-view. To achieve this, the depth data, given as a depth map, is first converted to a point format where for each pixel coordinate  $[x_d \ y_d]^T$  exists a 3-d point  $X_d = [x_d \ y_d \ z_d]^T$  with  $z_d = \text{depth}(x_d, y_d)$ , and then

transformed from *Depth Image Coordinate Space* to *RGB Image Coordinate Space*,  $X_{rgb}$ , using the following equalities:

$$X'_d = \begin{bmatrix} x'_d \\ y'_d \\ z'_d \end{bmatrix} = \begin{bmatrix} \frac{(x_d - c_{d,x}) \cdot z_d}{f_{d,x}} \\ \frac{(y_d - c_{d,y}) \cdot z_d}{f_{d,y}} \\ z_d \end{bmatrix} \tag{2}$$

$$X'_{rgb} = R^{-1} \cdot (X'_d - T) \tag{3}$$

$$X_{rgb} = \begin{bmatrix} x_{rgb} \\ y_{rgb} \\ z_{rgb} \end{bmatrix} = \begin{bmatrix} \frac{x'_{rgb} \cdot f_{rgb,x}}{z'_{rgb}} + c_{rgb,x} \\ \frac{y'_{rgb} \cdot f_{rgb,y}}{z'_{rgb}} + c_{rgb,y} \\ z'_{rgb} \end{bmatrix} \tag{4}$$

Points  $X_{rgb}$  in (4) can then be mapped to a 2D  $W_{rgb} \times H_{rgb}$  size image, forming a registered depth image. This process summarizes the standard registration algorithm. However, the computation of registered RGB images is needed for target generation. So the  $RGB \mapsto D$  mappings produced by Eqs. (2)–(4) are reversed to build a 2D  $W_d \times H_d$  image of color values instead. Of course, doing this still leaves depth holes with no RGB value attributed, weakening the whole purpose of performing RGB-D registration. To overcome this, depth holes are filled with pixel interpolation and blurring to create smooth transitions between edges of known depth regions.

**Inpainting:** After completing the RGB-D registration, we employ a color-guided fast marching method (FMM) inpainting algorithm to generate the target frames for training the denoising network. Following the original FMM inpainting technique [53], the algorithm starts by delineating the boundaries of all hole regions within the image. Subsequently, it performs inpainting from the outer pixels of these boundaries inwards, ensuring that all hole regions are filled.

Our FMM inpainting technique combines ideas presented in Refs. [30, 33, 53] and introduces novel elements that enable better results in consumer depth maps. Specifically, the pixel weighting function (see Eq. (5)) is different from the original FMM inpainting [53]. Concretely, we prioritize the distance factor  $w_{dst}$  while dropping the factors  $w_{lev}$  and  $w_{dir}$ . Additionally, we include two novel weights:  $w_g$ , relating to color guidance [33]; and  $conf$ , a confidence factor as in Ref. [30]. All these new insights contribute to the following novel functions:

$$w(p, q) = w_{dst}^2(p, q) \cdot w_g(p, q) \cdot conf(q) \tag{5}$$

$$w_{dst}(p, q) = \frac{d_0^2}{\|p - q\|^2} \tag{6}$$

$$w_g(p, q) = \exp\left(-\frac{\|G(p) - G(q)\|^2}{2 \cdot \sigma_g^2}\right) \tag{7}$$

$$conf(q) = \frac{1}{1 + 2 \cdot T_{out}(q)} \tag{8}$$

where  $d_0$  is the minimum inter-pixel distance, usually 1,  $G$  denotes the guiding image, and  $\sigma_g^2$  is its standard deviation. Additionally,  $T$  is a distance map that stores the distance of each pixel to the closest initial hole patch boundary, and  $T_{out}$  is a function that zeroes pixels in the set of initial holes  $\Omega$  and assigns  $T$  to the remaining.

Furthermore, as in GFMM [33], the pixel inpainting priority takes into account both the distance to the initial hole boundary, given by  $T(p)$ , and the guidance value of neighboring pixels, so that homogeneous areas are inpainted before other regions more likely to be transitive or edge areas. However, the priority function used in SReD’s inpainting for target generation,  $Pr(p)$ , introduces a new normalization variable  $T_{max}$ , leading to the final equation:

$$Pr(p) = (1 - \lambda) \cdot \frac{T(p)}{T_{max}} + \lambda \cdot (1 - S_g(p)) \tag{9}$$

$$S_g(p) = \frac{1}{|N(p)|} \cdot \sum_{q \in N(p)} w_g(p, q), \tag{10}$$

where  $S_g(p)$  (Eq. (10)) gives the local guide similarity at pixel  $p$ ,  $|N(p)|$  denotes the number of known pixels in the neighborhood of  $p$ ,  $T_{max}$  is the greatest value in distance map  $T$ , and  $\lambda$  is a mixing parameter. (Note: lower  $Pr$  values denote greater priority.)

### 3.3 Denoising network

The denoising neural network implemented in SReD adopts a convolutional autoencoder architecture based on the U-Net design [47]. This architecture is primarily influenced by features from MF2F [13], FastDVDnet [52], and Noise-2Noise [27]. During inference, the network takes as input three sequential depth frames, specifically  $d_{t-2}, d_{t-1}, d_t$ . Conversely, during training, the input frames are dilated, namely  $d_{t-4}, d_{t-2}, d_t$ . This setup enables using an inpainted version of frame  $d_{t-1}$  as the target, thereby preventing the network from learning the identity function. The network employs the mean absolute error (MAE or  $L_1$ ) loss function to measure the discrepancy between the inpainted target  $d_{t-1}^*$  and the

input depth frame  $d_t$ . In noisy regions, this approach replicates the effects of Noise2Noise [27], and for depth holes, the network learns inpainting techniques.

Using a U-Net [47] helps with image denoising. This is because the skip-connections enable passing higher frequency details from the encoding stage to the decoding stages via layer concatenation. This propagation allows the network to “flatten” noise areas while still preserving sharp image features, such as object contours. Regarding the layer structure, the general layout loosely follows the model presented in Noise2Noise [27], differing mainly in the number of channels at each network block and the downsampling/upsampling layers. Instead of max pooling and 2D upsampling layers, SReD uses 2D convolutions with stride two and transposed 2D convolutions, giving the model more learning flexibility.

Additionally, like in FastDVDnet [52], the network applies a final residual operation between the input frame  $d_t$  and the frame resulting from the last convolutional layer in the model  $d_{\text{last}}$ , yielding the depth frame prediction  $d_{\text{pred}} = d_t - d_{\text{last}}$ .

## 4 Evaluation

We conducted a comprehensive evaluation of our method to assess the algorithm’s performance. A quantitative evaluation was performed using a reference-independent noise metric to measure SReD’s depth-restoration capabilities objectively. We also performed several tests using a synthetic dataset that provides usable artificial ground-truth data. The algorithm’s time performance was also assessed to determine its suitability for real-time applications. Furthermore, we evaluated the method’s temporal coherence using a specialized metric. Finally, we compared SReD to other relevant reference-independent restoration algorithms to situate its performance within the broader landscape of available techniques.

### 4.1 Data and metrics

Identifying an appropriate combination of data and metrics for evaluating SReD proved to be a non-trivial task. Ideally, we would have access to a consumer-grade depth video dataset featuring raw frame sequences and reference depth data, perhaps captured using a high-precision laser sensor. However, such a dataset is not readily available. This very challenge underscores the importance of developing self-supervised depth denoisers like SReD.

We evaluated SReD on a depth video dataset devoid of reference depth. The evaluation also used appropriate reference-independent metrics. We conducted comprehensive tests on the CoRBS dataset [58], explicitly focusing

on the Kinect v2 subset. These data include five distinct RGB-D frame sequences capturing a stationary scene with a mobile camera, resulting in an aggregate of approximately 14,000 depth frames. In terms of evaluation metrics, the denoised depth frames underwent quantitative assessment concerning noise through a “non-reference metric for image denoising” [22] (NMID), a robust measure based on structure similarity maps from both homogeneous and highly-structured regions, in the absence of the original clean data.

Additional tests relied on synthetic depth data from the InteriorNet dataset published in Ref. [31], which provides computer-rendered RGB and depth images for various indoor scenes. For evaluating this synthetic ground-truth data, we employed proper reference metrics for the comparisons: MSE, PSNR, and SSI. Since this dataset does not provide noisy data, we introduced synthetic noise using the developed Kinect v1 noise model from Handa et al. [19], which combines Gaussian noise, bilinear interpolation, and quantization to produce noisier pixels at higher distances and missing depth values at pixels whose corresponding normals are close to perpendicular to the camera’s viewing direction.

To further assess our approach’s feasibility on these data, we evaluated it against the total variation (TV) method [7]. This denoising technique reduces the total magnitude of the image’s color intensity gradient while simultaneously trying to keep object boundaries. The regularization parameter weight used in the algorithmic implementation from [56], which controls the denoising strength at the expense of fidelity to the original image, was set to 0.4 as this value maximized the mean scores among both datasets in our experiments.

Furthermore, we assess temporal coherence using straightforward image differences,  $M_{\text{temp}} = \text{mean}(I_{t+1} - I_t)$ . We address depth value oscillations over time by analyzing granular noise values on a frame-by-frame basis over contiguous video sequences from the dataset.

### 4.2 Experimental setup

We ran all experiments on a Windows 10 desktop machine with an NVIDIA GeForce RTX 3080 GPU, a Ryzen 7 3700x 8-core CPU, 16 GB of RAM, and an SSD disk drive. We developed and tested SReD using Python 3.10.8 and tensorflow-gpu 2.10, along with CUDA 11.2 and cuDNN 8.1.

### 4.3 Results

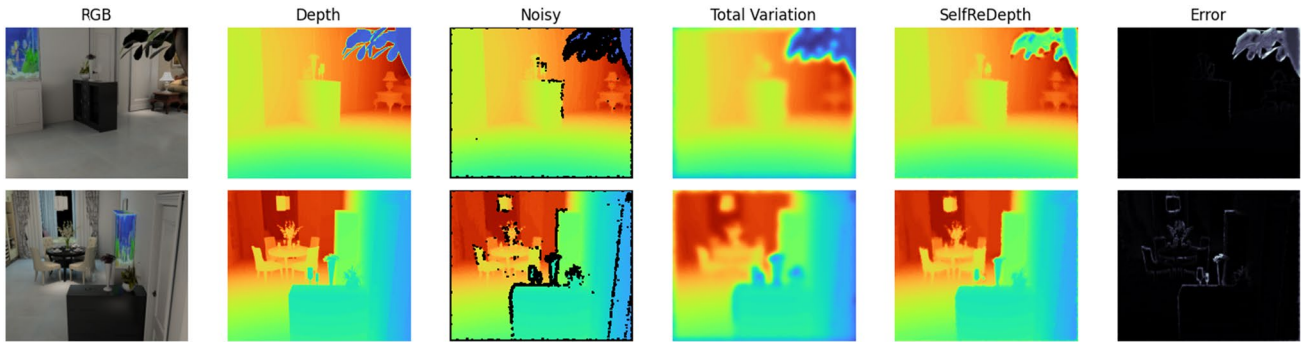
We trained SReD with batch sizes of 16 and 200 epochs on the CoRBS [58] dataset, which we also thoroughly shuffled and set with validation and test splits of 0.1 and 0.04, respectively.

In the produced denoised depth maps, in Fig. 5, it can be seen that the model learned how to attenuate the noise in the original depth map and fill depth holes. Moreover, on the hardware used for evaluation, the model takes, on average, 9 ms to denoise each depth map. Given that a regular RGB-D sensor, such as the Kinect v2, records data at a frequency of 30 frames-per-second (33 ms per frame), this evaluation confirms that the model can achieve the desired real-time performance during inference.


We also compared SReD against other approaches, including two deep-learning methods, Noise2Stack [44] and Noise2Noise [27] and two deterministic approaches, the Total Variation method [7] and a combination of a pair of methods that applies FMM inpainting [53] followed by bilateral filtering [55] denoising. We chose Noise2Noise to evaluate how the implemented technique differs from the original self-supervised U-Net [47] denoiser and what benefits were secured by targeting specifically the denoising of depth. Similarly, we chose Noise2Stack [44] to compare SReD against another spatio-temporal depth denoiser. Last, we used the deterministic FMM + BF combination to evaluate how SReD fares against more traditional approaches that perform both denoising and inpainting.

As already mentioned, we used a non-reference noise metric, NMID [22], to quantify the denoising quality, and applied a direct image difference metric to evaluate temporal coherence on contiguous depth videos. The results can be seen in Table 1. From the measured values, we note that SReD attained promising results, rivaling the significantly more computationally expensive deterministic algorithms with the NMID metric and achieving the best results with the temporal coherence in both datasets as expected, since it relies on multiple consecutive frames for its inference process. In addition, results in Fig. 4 show that original per-frame noise discrepancies are mostly fixed, and yield temporarily consistent values after denoising.

As shown in Fig. 5, Noise2Noise and Noise2Stack can only perform pixel denoising and not depth completion. As for the deterministic algorithm combination, while it was capable of denoising and inpainting the depth maps, it can also be visually seen that both the edge preservation and depth completion results are inferior to those of SReD.




**Fig. 3** Example scenes from the ground-truth dataset demonstrating (from left to right) RGB color image, real depth map, depth map with synthetic noise added (missing values in black color), results from the total variation method, SelfReDepth, and error map from our approach



**Fig. 4** Temporal analysis of the mean depth value differences across frames. The noisy data were generated from a sample video sequence extracted from the InteriorNet synthetic dataset and restored using the SelfReDepth and total variation methods





**Fig. 5** Visual comparison between four image restoration algorithms applied to an example image taken from real data from the CoRBS dataset. From left to right: SelfReDepth, Noise2Noise, Noise2Stack and FMM + BF

**Table 1** Benchmark of several metrics for all the tested methods: NMID (higher is better), temporal difference (lower is better), PSNR (higher is better), MSE and SSI (lower is better)

	CoRBS		InteriorNet				
	NMID	Temporal	NMID	Temporal	PSNR	MSE	SSI
SelfReDepth	<b>0.735</b>	<b>0.858</b>	0.1611	<b>0.0024</b>	38.663	0.00036	0.830
Noise2Noise	-0.154	1.085	0.1009	0.0039	39.020	0.00013	0.937
Noise2Stack	-0.166	1.012	-0.0504	0.0038	37.582	0.00020	0.918
FMM + BF	<b>0.735</b>	0.926	<b>0.2153</b>	0.0032	<b>43.753</b>	0.00008	0.971
Total variation	0.165	1.057	0.0129	0.0026	42.979	<b>0.00006</b>	<b>0.980</b>

Results highlighted in bold are the best in each category

The last three metrics are reported for both real (CoRBS) and synthetic (InteriorNet) datasets while the first two are only available for the synthetic dataset since this is the only one which provides ground-truth depth data

#### 4.4 Discussion

Based on the experiments and metrics, SReD effectively learned to reduce noise in depth maps. However, some image details still need to be recovered, evident in the blurred features of the doll in the CoRBS [58] dataset. In these cases where large black depth holes are present, inpainting struggles to effectively reconstruct this missing data due to the absence of depth information, resulting in over-smoothed restorations. While the results are promising, they highlight the need for further work on detail preservation.

Additionally, our method seems to struggle with accurately restoring the depth of objects very close to the camera, as seen in the top row of Fig. 3 and its rightmost error profile. This scenario where the objects are almost touching the camera was not seen in the training data but is very frequent in the synthetic dataset, signaling the need to extend the training set to a wider range of scenarios.

The metrics pitched SReD promisingly against the other four algorithms evaluated. The visual analysis of the denoised data aligns well with the NMID metric values, reinforcing its reliability. Although not optimal, our method also performed very favorably in the synthetic ground-truth dataset regarding MSE, PSNR, and SSI scores. The subpar results on this synthetic dataset could be related to the a

posteriori added synthetic noise based on a Kinect v1 noise model, while SReD was trained on Kinect v2 noise. This discrepancy, limited by the nonexistence of a usable v2 noise model implementation, might explain the better results achieved in this dataset by more general image restoration approaches not based on deep learning.

In qualitative terms, upon visual inspection, SReD achieved both consistent inpainting and denoising behavior and outperforms both deterministic approaches, namely the FMM + BF algorithm, particularly when filling missing areas and sharpening object boundaries, and TV, as this last method over-smooths the overall depth image, failing to preserve object details.

#### 4.5 Real-time performance

Finally, the implemented algorithm can produce denoised frames at frequencies higher than 30 frames per second, thus making SReD suitable for real-time use. Indeed, on the computer hardware used for evaluation, the model requires, on average, 9 ms to denoise each depth map. Given that commercial off-the-shelf RGB-D sensors, such as the Kinect v2, generate data at 30 frames-per-second (33 ms per frame), our technique can achieve real-time performance at even higher frame rates. The modular design of the U-Net architecture allows for straightforward scalability to accommodate larger

image sizes without a significant impact on computational time, thereby maintaining real-time performance, as will be detailed in the following section.

**SelfReDepth complexity:** Our approach is made up of three main blocks as follows: inpainting (Eqs. (1) and (5)), registration (Eqs. (2)–(4)) and denoising (U-Net network) procedures. As already stated in Sect. 3.1, our proposal is designed to satisfy real-time requirements, as we can modify the SReD architecture during the inference stage, which is the one that has a direct impact on the complexity budget. Specifically, (i) remove the target generation, and (ii) the denoise shift to non-dilated input. Thus, only the U-Net (denoiser) needs to be carefully addressed since it is the unique block that affects a constrained time budget requirement.

**Time complexity:** We detail the architecture adaptation under a given complexity budget. The choice of the U-Net provides flexibility because it is possible to adapt its architecture under a predefined budget. The designs of network architectures should exhibit tradeoffs among several of its components, i.e. depth, numbers of filters, and filter sizes, from which the scalability is accomplished. From the above, depth is the most influential concerning the accuracy. Although it is not a straightforward observation, previous work [49, 51] has demonstrated its impact. The total time complexity of all the convolutional layers is given as  $O\left(\sum_{l=1}^d c_{l-1} \cdot s_l^2 \cdot f_l \cdot m_l^2\right)$  where  $d$  is the depth of the network, (i.e. the number of convolutional layers),  $l$  indexes the convolutional layer,  $c_l$  is the number of input channels in the  $l$ -th layer,  $f_l$  is the number of filters in the  $l$ -th layer, (i.e., the width),  $s_l$  and  $m_l$  are the spatial size of the filter and the size of the output of the feature map, respectively. The time cost of fully connected and pooling layers is not considered since these layers take about 5-10% computational time. The time complexity above is the basis of the network designs, from which we consider the tradeoffs between the depth  $d$  and filter sizes  $f_l$ , inferring how the network scales in time.

Concretely, we design a model by replacing the layers in our experimental evaluation. This means that when we replace a few layers with some other layers, we must guarantee that the complexity is preserved without changing the remaining layers in the architecture. To design such a replacement, we progressively modify the model and observe the changes in accuracy. Our method addresses the following tradeoffs:

1. depth  $d$  and filter sizes  $s_l$ ,
2. depth  $d$  and width  $f_l$ , and
3. width  $f_l$  and filter sizes  $s_l$ .

We illustrate one of the steps above, the remaining with an analogous procedure, only changing the corresponding parameters accordingly. For instance, as an illustrative of step 1 (tradeoff between depth  $d$  and filter size  $s$ ), we replace a larger filter, say  $s_1$ , with a cascade of smaller filters, say  $s_2$ . Denoting the layer configuration as above,  $L_{\text{conf}} = c_{l-1} \cdot s_l^2 \cdot f_l$  and considering two instances of filter sizes, e.g.  $s_1 = 3s_2 = 2$ , and  $c_{l-1} = N$ ,  $f_l = N$ , we have the following complexities:

$$\begin{aligned} \mathcal{O}_1 &= N^2 \cdot s_1^2 \\ \mathcal{O}_2 &= 2 \cdot (N^2 \cdot s_2^2) \end{aligned}$$

This replacement a  $s_1 \times s_1$  layer with  $N$  input/output channels is replaced by two  $s_2 \times s_2$  layers with  $N$  input/output channels. After the above replacement, the complexity involved in these layers is nearly unchanged, with the reduction fraction of  $2s_1^2/s_2^2 \approx 1$ .

With the strategy above, we can “deepen” the network under the same complexity time budget. This allows us to obtain several architectures and pick the best accuracy. In concrete, from our experimental evaluation yielding the times mentioned in Sect. 4.5, we found the best accuracy using the configuration in our final network, which has 31 layers, 1729 convolutional size three filters, yielding 1,260,865 trainable parameters.

**Scalability:** Now, we delve into how the architecture scales with the size of input images. First, let us introduce some basic notation:

- $\text{conv2D}_{F,\text{st}}^{\text{st}}$ : 2D (contraction) convolution with  $F$  number of filters and with stride  $st$
- $\text{conv2D}_{F,\text{st}}^{\text{T}}$ : 2D (expansion) transpose convolution with  $F$  number of filters and with stride  $st$

Our U-Net network includes the following main blocks:

- First block: 2 x  $\text{conv2D}_{32,1}$
- $i$ th down block: 1 x  $\text{conv2D}_{F_i,2}$  + 1 x  $\text{conv2D}_{F_i,1}$
- $i$ th up block: 2 x  $\text{conv2D}_{F_i,1}$  + 1 x  $\text{conv2D}_{F_i,2}^{\text{T}}$
- Last block: 2 x  $\text{conv2D}_{32,1}$  + 1 x  $\text{conv2D}_{1,1}$

We use five blocks for each down and up stage, thus having  $i \in \{1, \dots, 5\}$ . The number of filters for each block is  $F = [F_0 \dots F_5] = [32 \ 32 \ 48 \ 48 \ 64 \ 128]$ , where  $F_0$  accounts for the filter in the first and last blocks.

Now, it is straightforward to determine convolutions. Assuming an image size of  $W \times H$ , we have:

- First Block:  $2 \cdot F_0 \cdot W \cdot H$
- Down Block  $i$ :  $2 \cdot F_i \cdot W \cdot H \cdot 2^{-2i}$
- Up Block  $i$ :  $3 \cdot F_i \cdot W \cdot H \cdot 2^{-2i}$

- Last Block:  $(2 \cdot F_0 + 1) \cdot W \cdot H$
- $\mathbf{T} \quad \mathbf{o} \quad \mathbf{t} \quad \mathbf{a} \quad \mathbf{l} \quad :$   
 $\left(1 + 4F_0 + 5 \sum_{i=1}^5 F_i \cdot 2^{-2i}\right) W \cdot H = 189.625 W \cdot H$

This means that, for a  $\Delta$ -increment in the image resolution  $(W + \Delta)(H + \Delta)$ , we have a complexity of  $\mathcal{O}(\Delta^2)$ . So, roughly speaking, a twofold increase in image resolution would entail a fourfold increase in image processing time using the same architecture and memory footprint. Assuming that in the worst case, 80% of the CPU time is spent on running the Neural Network, the processing time per frame would be around 30 ms for an effective frame rate of 30 Hz, which is still reasonable.

## 5 Limitations

While our technique has proven to be very effective at restoring depth values from noisy RGB-D images, it can be improved in several ways. A notable limitation involves adequately addressing high-frequency temporal noise. While effective for general noise reduction, averaging pixel values across frames falls short in capturing and mitigating these rapid fluctuations. This suggests potential for future refinement. More sophisticated techniques should be capable of discerning and smoothing out high-frequency temporal noise without compromising the dynamic content of the scenes.

## 6 Conclusions and future work

We introduced SelfReDepth, a self-supervised approach for denoising and completing low-quality depth maps generated from consumer-grade sensors. Our technique advances self-supervised learning in-depth data denoising, offering a precise, data-driven architecture without reference data. This flexibility makes SelfReDepth easily adaptable across various environments and applications.

SelfReDepth's architecture features two main elements: a denoising network and a target generation component. The denoising network is inspired by the original Noise2Noise [27] and MF2F [13] video denoisers and is responsible for learning how to denoise depth data without the need for reference data. Meanwhile, the target generation component fills in the gaps in target depth frames using color-guided FMM inpainting. The technique can denoise inaccurate depth values and paint out missing ones with this structure.

We also implemented and assessed SelfReDepth for both denoising efficacy and time performance. Results indicate real-time noise elimination and successful inpainting of depth gaps. Future work will focus on preserving image details compromised by denoising. Training with synthetic

data might also improve depth inpainting performance and dampen oscillations.

In future work, we aim to explore controllable image denoising to generate clean sample frames with human perceptual priors and balance sharpness and smoothness. In most common filter-based denoising approaches, this can be straightforwardly achieved by regulating the filtering strength. However, for deep neural networks (DNN), regulating the final denoising strength requires performing network inference each time. This of course, hampers the real-time user interaction. Further work will address real-time controllable denoising, to be integrated into a video denoising pipeline that provides a fully controllable user interface to edit arbitrary denoising levels in real-time with only one-time DNN inference.

SelfReDepth represents a significant advancement in data denoising, tackling noise and depth hole challenges with notable efficiency. The outcomes of our research are encouraging, illustrating the algorithm's capacity to mitigate these problems. However, the concomitant loss of certain image details in the process highlights areas for potential improvement. This observation underscores the need for additional investigation while pointing to clear pathways for refining future algorithm iterations. Such enhancements aim to improve the balance between our denoising algorithm's robustness and critical image detail preservation, enhancing its already remarkable efficiency and making it more applicable to very complex scenarios.

**Acknowledgements** The work reported in this article was partially supported under the auspices of the UNESCO Chair on AI and VR by national funds through Fundação para a Ciência e a Tecnologia with references DOI:10.54499/UIDB/50021/2020, DOI:10.54499/DL57/2016/CP1368/CT0002 and 2022.09212.PTD (XAVIER project). The SelfReDepth source code is publicly available at: <https://github.com/alexduarte23/sred>

**Author contributions** A.D. Did most of the Software Development and Tests and contributed to the initial write-up F.F. Contributed to supervision, technical discussions and the initial write-up J.M.P. Contributed supervision, technical discussions and write-up C.M. Contributed to discussions, Machine Learning Architecture and Proofreading several versions of the manuscript J.C.N. Contributed to discussions, technical arguments, wrote the section and analysis of complexity and reviewed the manuscript J.J. Contributed to supervision, the original idea, discussions, revising the manuscript, and wrote parts of the section on Discussion and Conclusions.

**Funding** Open access funding provided by FCTIFCCN (b-on).

**Data availability** All source code and data developed for the project are available via Github at: <https://github.com/alexduarte23/sred>

## Declarations

**Conflict of interest** The authors declare that they have no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Basso, F., Menegatti, E., Pretto, A.: Robust intrinsic and extrinsic calibration of rgb-d cameras. *IEEE Trans. Rob.* **34**(5), 1315–1332 (2018)
- Batson, J., Royer, L.: Noise2self: blind denoising by self-supervision. In: *Proceedings of the 36th International Conference on Machine Learning*, pp. 524–533 (2019)
- Bertalmio, M., Bertozzi, A.L., Sapiro, G.: Navier–stokes, fluid dynamics, and image and video inpainting. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, vol. 1, p. I (2001)
- Calvarons, A.F.: Improved noise2noise denoising with limited data. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 796–805 (2021)
- Capecchi, M., Ceravolo, M.G., Ferracuti, F., Iarlori, S., Kyrki, V., Longhi, S., Romeo, L., Verdini, F.: Physical rehabilitation exercises assessment based on hidden semi-Markov model by kinect v2. In: *IEEE-EMBS International Conference on Biomedical and Health Informatics*, pp. 256–259 (2016)
- Cha, S., Park, T., Kim, B., Baek, J., Moon, T.: Gan2gan: generative noise learning for blind denoising with single noisy images. *arXiv preprint arXiv:1905.10488* (2019)
- Chambolle, A.: An algorithm for total variation minimization and applications. *J. Math. Imaging Vis.* **20**, 89–97 (2004)
- Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., Zhang, Y.: Matterport3d: learning from rgb-d data in indoor environments. In: *2017 International Conference on 3D Vision* (2017)
- Chaudhary, R., Dasgupta, H.: An approach for noise removal on depth images. *arXiv preprint arXiv:1602.05168* (2016)
- Chen, L., Lin, H., Li, S.: Depth image enhancement for kinect using region growing and bilateral filter. In: *ICPR2012*, pp. 3070–3073 (2012)
- Choi, J., Jung, D., Lee, Y., Kim, D., Manocha, D., Lee, D.: Self-deco: self-supervised monocular depth completion in challenging indoor environments. In: *IEEE International Conference on Robotics and Automation*, pp. 467–474 (2021)
- Dai, Y., Fu, Y., Li, B., Zhang, X., Yu, T., Wang, W.: A new filtering system for using a consumer depth camera at close range. *Sensors* **19**(16), 3460 (2019)
- Dewil, V., Anger, J., Davy, A., Ehret, T., Facciolo, G., Arias, P.: Self-supervised training for blind multi-frame video denoising. In: *IEEE Winter Conference on Applications of Computer Vision*, pp. 2724–2734 (2021)
- Essmaeel, K., Gallo, L., Damiani, E., De Pietro, G., Dipanda, A.: Temporal denoising of kinect depth data. In: *Eighth International Conference on Signal Image Technology and Internet Based Systems*, pp. 47–52. *IEEE* (2012)
- Feng, D., Rosenbaum, L., Dietmayer, K.: Towards safe autonomous driving: capture uncertainty in the deep neural network for lidar 3d vehicle detection. In: *2018 21st International Conference on Intelligent Transportation Systems*, pp. 3266–3273 (2018)
- Feng, Z., Jing, L., Yin, P., Tian, Y., Li, B.: Advancing self-supervised monocular depth learning with sparse lidar. *arXiv preprint arXiv:2109.09628* (2021)
- Gabel, M., Gilad-Bachrach, R., Renshaw, E., Schuster, A.: Full body gait analysis with kinect. In: *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 1964–1967 (2012)
- Gao, Z., Yu, Y., Zhou, Y., Du, S.: Leveraging two kinect sensors for accurate full-body motion capture. *Sensors* **15**(9), 24297–24317 (2015)
- Handa, A., Whelan, T., McDonald, J., Davison, A.J.: A benchmark for rgb-d visual odometry, 3d reconstruction and slam. In: *IEEE international conference on Robotics and automation*, pp. 1524–1531. *IEEE* (2014)
- Jiang, L., Xiao, S., He, C.: Kinect depth map inpainting using a multi-scale deep convolutional neural network. In: *Proceedings of the 2018 International Conference on Image and Graphics Processing*, pp. 91–95 (2018)
- Jorge, J., Anjos, R.K.D., Silva, R.: Dynamic occlusion handling for real-time ar applications. In: *Proceedings of the 17th International Conference on Virtual-Reality Continuum and Its Applications in Industry* (2019)
- Kong, X., Li, K., Yang, Q., Wenyin, L., Yang, M.H.: A new image quality metric for image auto-denoising. In: *IEEE International Conference on Computer Vision*, pp. 2888–2895 (2013)
- Krull, A., Buchholz, T.O., Jug, F.: Noise2void—learning denoising from single noisy images. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2124–2132 (2019)
- Krull, A., Vičar, T., Prakash, M., Lalit, M., Jug, F.: Probabilistic noise2void: unsupervised content-aware denoising. *Front. Comput. Sci.* **2**, 5 (2020)
- Kweon, I.S., Jung, J., Lee, J.Y.: Noise aware depth denoising for a time-of-flight camera. In: *20th Korea–Japan Joint Workshop on Frontiers of Computer Vision* (2014)
- Lachat, E., Macher, H., Landes, T., Grussenmeyer, P.: Assessment and calibration of a rgb-d camera (kinect v2 sensor) towards a potential use for close-range 3d modeling. *Remote Sens.* **7**(10), 13070–13097 (2015)
- Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T.: Noise2noise: learning image restoration without clean data. In: *International Conference on Machine Learning*, pp. 2965–2974. *PMLR* (2018)
- Lemarchand, F., Findeli, T., Nogues, E., Pelcat, M.: Noise-breaker: gradual image denoising guided by noise analysis. In: *IEEE 22nd International Workshop on Multimedia Signal Processing*, pp. 1–6 (2020)
- Li, A., Yuan, Z., Ling, Y., Chit, W., Zhang, S., Zhang, C.: Fast-completion: a cascade network with multiscale group-fused inputs for real-time depth completion. In: *25th International Conference on Pattern Recognition*, pp. 866–872 (2021)
- Li, L., Wu, H., Chen, Z.: Depth image restoration method based on improved fmm algorithm. In: *2021 13th International Conference on Machine Learning and Computing, ICMLC 2021*, pp. 349–355 (2021)
- Li, W., Saeedi, S., McCormac, J., Clark, R., Tzoumanikas, D., Ye, Q., Huang, Y., Tang, R., Leutenegger, S.: Interiornet: mega-scale multi-sensor photo-realistic indoor scenes dataset. *arXiv preprint arXiv:1809.00716* (2018)
- Liu, G., Reda, F.A., Shih, K.J., Wang, T.C., Tao, A., Catanzaro, B.: Image inpainting for irregular holes using partial convolutions. In: *Computer Vision—ECCV 2018*, pp. 89–105 (2018)

33. Liu, J., Gong, X., Liu, J.: Guided inpainting and filtering for kinect depth maps. In: Proceedings of the 21st International Conference on Pattern Recognition, pp. 2055–2058 (2012)
34. Liu, J., Liu, Y., Zhang, G., Zhu, P., Chen, Y.Q.: Detecting and tracking people in real time with rgb-d camera. *Pattern Recogn. Lett.* **53**, 16–23 (2015)
35. Liu, S., Chen, C., Kehtarnavaz, N.: A computationally efficient denoising and hole-filling method for depth image enhancement. In: Real-time image and video processing 2016, vol. 9897, pp. 235–243. SPIE (2016)
36. Ma, F., Cavalheiro, G.V., Karaman, S.: Self-supervised sparse-to-dense: self-supervised depth completion from lidar and monocular camera. In: International Conference on Robotics and Automation, pp. 3288–3295 (2019)
37. Maimone, A., Bidwell, J., Peng, K., Fuchs, H.: Enhanced personal autostereoscopic telepresence system using commodity depth cameras. *Comput. Graph.* **36**(7), 791–807 (2012)
38. Mallick, T., Das, P.P., Majumdar, A.K.: Characterizations of noise in kinect depth images: a review. *IEEE Sens. J.* **14**(6), 1731–1740 (2014)
39. Metzler, C.A., Mousavi, A., Heckel, R., Baraniuk, R.G.: Unsupervised learning with stein’s unbiased risk estimator. arXiv preprint [arXiv:1805.10531](https://arxiv.org/abs/1805.10531) (2018)
40. Mohan, S., Vincent, J.L., Manzorro, R., Crozier, P., Fernandez-Granda, C., Simoncelli, E.P.: Adaptive denoising via gaintuning. In: Thirty-Fifth Conference on Neural Information Processing Systems (2021)
41. Moran, N., Schmidt, D., Zhong, Y., Coady, P.: Noisier2noise: learning to denoise from unpaired noisy data. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 12061–12069 (2020)
42. Newcombe, R.A., Izadi, S., Hilliges, O., Molyneaux, D., Kim, D., Davison, A.J., Kohi, P., Shotton, J., Hodges, S., Fitzgibbon, A.: Kinectfusion: real-time dense surface mapping and tracking. In: 10th international symposium on mixed and augmented reality, pp. 127–136 (2011)
43. Oyedotun, O.K., Demisse, G., El Rahman Shabayek, A., Aouada, D., Ottersten, B.: Facial expression recognition via joint deep learning of rgb-depth map latent representations. In: IEEE International Conference on Computer Vision Workshops, pp. 3161–3168 (2017)
44. Papkov, M., Roberts, K., Madissoon, L.A., Shilts, J., Bayraktar, O., Fishman, D., Palo, K., Parts, L.: Noise2stack: improving image restoration by learning from volumetric data. In: International Workshop Machine Learning for Medical Image Reconstruction, pp. 99–108 (2021)
45. Quan, Y., Chen, M., Pang, T., Ji, H.: Self2self with dropout: learning self-supervised denoising from single image. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1887–1895 (2020)
46. Ren, Z., Yuan, J., Zhang, Z.: Robust hand gesture recognition based on finger-earth mover’s distance with a commodity depth camera. In: Proceedings of the 19th international conference on Multimedia, pp. 1093–1096 (2011)
47. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention, pp. 234–241 (2015)
48. Sheth, D.Y., Mohan, S., Vincent, J.L., Manzorro, R., Crozier, P.A., Khapra, M.M., Simoncelli, E.P., Fernandez-Granda, C.: Unsupervised deep video denoising. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 1759–1768 (2021)
49. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014)
50. Song, W., Le, A.V., Yun, S., Jung, S.W., Won, C.S.: Depth completion for kinect v2 sensor. *Multimed. Tools Appl.* **76**(3), 4357–4380 (2017)
51. Szegegy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–9 (2015)
52. Tassano, M., Delon, J., Veit, T.: Fastdvdnet: Towards real-time deep video denoising without flow estimation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1354–1363 (2020)
53. Telea, A.: An image inpainting technique based on the fast marching method. *J. Graph. Tools* **9**(1), 23–34 (2004)
54. Tölgyessy, M., Dekan, M., Chovanec, L., Hubinský, P.: Evaluation of the azure kinect and its comparison to kinect v1 and kinect v2. *Sensors* **21**(2), 413 (2021)
55. Tomasi, C., Manduchi, R.: Bilateral filtering for gray and color images. In: 6th International Conference on Computer Vision (IEEE Cat. No. 98CH36271), pp. 839–846 (1998)
56. Van der Walt, S., Schönberger, J.L., Nunez-Iglesias, J., Boulogne, F., Warner, J.D., Yager, N., Gouillart, E., Yu, T.: scikit-image: image processing in python. *PeerJ* **2**, e453 (2014)
57. Wan, Y., Li, Y., Jiang, J., Xu, B.: Edge voxel erosion for noise removal in 3d point clouds collected by kinect. In: Proceedings of the 2020 2nd International Conference on Image, Video and Signal Processing, pp. 59–63 (2020)
58. Wasenmüller, O., Meyer, M., Stricker, D.: Corbs: comprehensive rgb-d benchmark for slam using kinect v2. In: IEEE Winter Conference on Applications of Computer Vision, pp. 1–7 (2016)
59. Xiong, W., Yu, J., Lin, Z., Yang, J., Lu, X., Barnes, C., Luo, J.: Foreground-aware image inpainting. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 5833–5841 (2019)
60. Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Generative image inpainting with contextual attention. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 5505–5514 (2018)
61. Zennaro, S., Munaro, M., Milani, S., Zanuttigh, P., Bernardi, A., Ghidoni, S., Menegatti, E.: Performance evaluation of the 1st and 2nd generation kinect for multimedia applications. In: IEEE International Conference on Multimedia and Expo, pp. 1–6 (2015)
62. Zhang, B., Allebach, J.P.: Adaptive bilateral filter for sharpness enhancement and noise removal. *IEEE Int. Conf. Image Process.* **4**, 417–420 (2007)
63. Zhang, X., Yan, J., Feng, S., Lei, Z., Yi, D., Li, S.Z.: Water filling: Unsupervised people counting via vertical kinect sensor. In: IEEE 9th International Conference on Advanced Video and Signal-based Surveillance, pp. 215–220 (2012)
64. Zhou, X.: A study of microsoft kinect calibration. Department of Comp. Science, George Mason University, Fairfax (2012)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Alexandre Duarte** received the MSc degree in Computer Science and Engineering from Instituto Superior Técnico, Universidade de Lisboa, specialising in Artificial Intelligence and Data Processing and Analysis. His research interests include Computer Graphics, Artificial Intelligence and Computer Vision.



**Francisco Fernandes** received the MSc degree in Applied Mathematics and Computation and the PhD degree in Information Systems and Computer Engineering at Instituto Superior Técnico, Universidade de Lisboa. He is currently a Senior Researcher at INESC-ID at the Graphics and Interaction group. His research interests include advanced algorithms and data structures, artificial intelligence, bioinformatics and computational biology.



**João M. Pereira** is Associate Professor with Habilitation at Computer Science Department of Instituto Superior Técnico from University of Lisbon where he teaches Computer Graphics. He holds a PhD in Electrical and Computers Engineering from IST/UTL, and also received a MSc and a BsEE degrees in Electrical and Computers Engineering from IST/UTL. He is the Scientific Coordinator of the Graphics and Interaction scientific area at INESC-ID and is Associated Editor of Journal of

Real-Time Processing. His main research fields are Real-Time Rendering, Photorealistic Rendering, Parallel Computer Graphics, Augmented Reality, Networked Virtual Environments and Serious Games. He has been involved with several European and National projects, is author or co-author of more than 200 peer-reviewed scientific papers presented at national and international events and journals, and is member of the Eurographics Association, IEEE (Senior Member), ACM (Senior Member) and GPCG.



**Catarina Moreira** is a senior lecturer at the School of Information Systems, Queensland University of Technology, Australia, where she also holds the role of Deputy Higher Degree Research Academic Lead. She is a computer scientist and passionate about investigating data science, machine learning/deep learning technologies and their impacts on society. She is the co-founder and co-leader of the XAMI research initiative (eXplainable Analytics for Machine Intelligence), where she focuses on

innovative human interactive probabilistic models to promote more fair, transparent, and trustworthy AI models. She is also a pioneer in quantum probabilistic technologies for cognition and decision-making, where she has developed models to accommodate paradoxical human decisions in scenarios with high levels of uncertainty.



**Jacinto C. Nascimento** has been with the Faculty of the Department of Informatics and Computer Engineering, IST, where he is currently a Associate Professor. He is the recipient of several research grants and fellowships from Science and Technology Foundation (FCT-Portugal). His research interests include statistical image processing, pattern recognition, machine learning, medical imaging analysis, video surveillance, general visual object classification. He has coauthored more than 60 publications in international journals (many of which are sponsored by the IEEE), has served on program committees of many international conferences, and has been a reviewer for several international journals.

international journals (many of which are sponsored by the IEEE), has served on program committees of many international conferences, and has been a reviewer for several international journals.



**Joaquim Jorge** founded the Interaction and Graphics research group at INESC-ID. A Full Professor of Computer Graphics and Interactive Techniques at Instituto Superior Técnico, he holds the UNESCO Chair on AI\ and XR at the University of Lisbon. He is Editor-in-Chief of the Computers and Graphics Journal (Elsevier), a Fellow of the Eurographics Association and a Distinguished Speaker and Member of ACM, and an IEEE Senior Member, Distinguished Visitor and Contributor. An organizer of

50+ scientific events, including IEEE VR 2020 (papers co-chair), 2021 and 2022 (conference co-chair), Eurographics 2016 as IPC CO-Chair, INTERACT 2011 as co-chair, and ACM IUI 2012 as IPC co-chair, he has served on 200+ program committees and (co)authored over 300 publications in international refereed journals, conferences, and books. His research interests include virtual reality, graphical modelling, multimodal user interfaces, and advanced learning techniques.