Reconstructing occluded Elevation Information in Terrain Maps with Self-supervised Learning

Stolzle, Maximilian; Miki, Takahiro; Gerdes, Levin; Azkarate, Martin; Hutter, Marco

DOI

[10.1109/LRA.2022.3141662](https://doi.org/10.1109/LRA.2022.3141662)

Publication date

2022

Document Version

Final published version

Published in

IEEE Robotics and Automation Letters

Citation (APA)

Stolzle, M., Miki, T., Gerdes, L., Azkarate, M., & Hutter, M. (2022). Reconstructing occluded Elevation Information in Terrain Maps with Self-supervised Learning. *IEEE Robotics and Automation Letters*, 7(2), 1697-1704. <https://doi.org/10.1109/LRA.2022.3141662>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Reconstructing Occluded Elevation Information in Terrain Maps With Self-Supervised Learning

Maximilian Stölzle[✉], Takahiro Miki[✉], Levin Gerdes[✉], Martin Azkarate[✉], *Member, IEEE*,
and Marco Hutter[✉], *Member, IEEE*

Abstract—Accurate and complete terrain maps enhance the awareness of autonomous robots and enable safe and optimal path planning. Rocks and topography often create occlusions and lead to missing elevation information in the Digital Elevation Map (DEM). Currently, these occluded areas are either fully avoided during motion planning or the missing values in the elevation map are filled-in using traditional interpolation, diffusion or patch-matching techniques. These methods cannot leverage the high-level terrain characteristics and the geometric constraints of line of sight we humans use intuitively to predict occluded areas. We introduce a self-supervised learning approach capable of training on real-world data without a need for ground-truth information to reconstruct the occluded areas in the DEMs. We accomplish this by adding artificial occlusion to the incomplete elevation maps constructed on a real robot by performing ray casting. We first evaluate a supervised learning approach on synthetic data for which we have the full ground-truth available and subsequently move to several real-world datasets. These real-world datasets were recorded during exploration of both structured and unstructured terrain with a legged robot, and additionally in a planetary scenario on Lunar analogue terrain. We state a significant improvement compared to the baseline methods both on synthetic terrain and for the real-world datasets. Our neural network is able to run in real-time on both CPU and GPU with suitable sampling rates for autonomous ground robots. We motivate the applicability of reconstructing occlusion in elevation maps with preliminary motion planning experiments.

Index Terms—AI-enabled robotics mapping.

I. INTRODUCTION

AS WE empower mobile robots to autonomously navigate to their goal, they rely on maps of the surrounding environment for traversability analysis and motion planning. 2.5D

Manuscript received September 9, 2021; accepted January 4, 2022. Date of publication January 10, 2022; date of current version January 19, 2022. This letter was recommended for publication by Associate Editor L. Paull and Editor J. Civera upon evaluation of the reviewers' comments. This work was supported by the PERASPERA activity of the European Commission, Horizon 2020 projects No. 852044 and No. 101016970, and the Swiss-European Mobility Programme. (Corresponding author: Maximilian Stölzle.)

Maximilian Stölzle is with the Robotic Systems Lab, ETH Zürich, 8092 Zürich, Switzerland, with the Planetary Robotics Lab, European Space Agency, 2201 AZ Noordwijk, Netherlands, and also with Cognitive Robotics, Delft University of Technology, 2628 CD Delft, Netherlands (e-mail: M.W.Stolzle@tudelft.nl).

Takahiro Miki and Marco Hutter are with the Robotic Systems Lab, ETH Zürich, 8092 Zürich, Switzerland (e-mail: tamiki@ethz.ch; mahutter@ethz.ch).

Levin Gerdes and Martin Azkarate are with the Planetary Robotics Lab, European Space Agency, 2201 AZ Noordwijk, Netherlands (e-mail: gerdeslevin@gmail.com; martin.azkarate@esa.int).

Digital Object Identifier 10.1109/LRA.2022.3141662

Digital Elevation Maps (DEMs) represent a memory-efficient and accurate approximation for most terrains as they project the 3D structure into two-dimensional grid cells. Usually, robots use depth sensors such as LiDAR or stereo cameras to construct DEMs [1]. To enable safe and optimal path planning, we strive for complete and accurate elevation maps. In practice, our robots often encounter occlusions caused by terrain discontinuities such as rocks, obstacles or convex terrain characteristics which hide an area patch from the sensor's viewpoint. Further, depth measurements can be degraded due to reflections, stereo matching failures, dust, or textureless surfaces which frequently lead to additional missing elevation information in the DEM. This motivates the need for solutions to reconstruct the missing elevation information in the terrain map. We are inspired by the application of data-driven methods for inpainting of images [2]–[4] even for irregular holes [5] and aim to tailor these methods for the application of filling occlusion in terrain maps. Our method allows us to exploit prior information about the deployment area such as known terrain characteristics as well as the geometric constraints of line of sight for the reconstruction of missing elevation information using a neural network. The method is based on a U-Net [6] similarly to the state-of-the-art in neural network-based inpainting for vision problems [2]–[5]. We propose a new strategy for self-supervised learning applied to terrain maps based on adding artificial occlusion with ray casting. As non-occluded ground-truth data is very difficult to acquire, this enables us to use incomplete real-world elevation maps for training our neural network. We would like to point out that this self-supervised learning approach breaks ground towards a future where robots can actively improve their performance while they explore their environment in the line of the connected research on lifelong learning [7].

Existing implementations to complete sparse terrain maps for mobile robots either rely on searching terrain patches with close resemblance in an offline library [8] or very recently on neural networks trained with ground-truth data estimated offline [9]. The later approach [9] is based on the assumption that by Simultaneous Localization and Mapping (SLAM) also having access to future measurements, the map will be more complete and accurate than the DEM at the current time-step which is generated solely from current and past depth measurements. It should be noted that learned features will most likely be coupled with pose estimation errors and future flawed depth measurements through the *ground-truth* map.

We conduct extensive bench-marking of both supervised and self-supervised learning approaches against various baseline methods. We evaluated on multiple synthetic terrain types as well as real-world datasets collected with the legged robot ANYmal [10] and the planetary rover test-bed HDPR [11]. We state a decrease between 52% and 82% in Mean Squared Error (MSE) error with our self-supervised method on the real-world datasets compared to the respective best performing baseline approach. Finally, we observe in preliminary motion planning experiments that using reconstructed DEMs paths can be planned further ahead compared to planning with incomplete elevation maps. Consequently, the robot needs to stop less and avoids having to wait for re-planning of the path to finish.

We would like to emphasize the following contributions of this paper: We propose a self-supervised learning method for filling-in sparse DEMs by adding artificial occlusion with ray casting enabling training on partially occluded real-world datasets. We introduce an iterative algorithm for realistic occluded map data generation and to guide the sampling of vantage points for ray casting towards occlusion masks with suitable amounts of occlusion.

II. RELATED WORK

Image inpainting: Inpainting techniques utilise the structure and texture of an image to fill-in missing parts and can be used to for example remove unwanted objects from the image [12]. Traditional image inpainting methods can be separated into two categories: While diffusion-based inpainting methods leverage isophote lines to continuously propagate high-order derivatives of local pixel density from the exterior into the hole of the image [13], patch-based methods such as PatchMatch [14] search the known regions of the image at the patch-level and copy the best-matching patches into the missing region. Nowadays, data-driven learning-based approaches are the state-of-the-art for image inpainting [2]–[5]. Most of them rely on an adoption of the U-Net [6] model architecture. Many recent papers [2], [5], [15] include perceptual and style losses [16]. They rely on high-level features extracted using a pretrained VGG-16 network. These perceptual and style losses work similarly to adversarial losses by not just enforcing convergence of the reconstruction on a pixel-by-pixel basis, but also enforcing encoded high-level characteristics to be similar. Partial convolutions can be appropriate for irregular holes as convolutional strategies might fail to adjust to the changing shapes of the hole boundaries [5]. A Generative Adversarial Network (GAN)’s ability to generate high-fidelity images motivates the use of adversarial methods to produce sharper inpaintings [2], [4], [17].

Self-supervised learning: Self-supervised learning is essential to alleviate the gap between synthetic data and real-world data (Syn2Real) and enable training on real-world datasets for which complete ground-truth information is challenging to acquire. Zhan *et al.* [18] propose a self-supervised approach for de-occlusion in the framework of scene understanding by overlaying semantically-extracted objects onto images. Dai *et al.* [19] uses a self-supervised learning approach for scene completion of sparse real-world RGB-D scans using a GAN by randomly removing some of the scans.

Elevation mapping: Robot-centric elevation mapping uses pose estimates from Inertial Measurements Unit (IMU) or odometry and local distance measurements from a laser range sensor, structured light, or stereo camera to derive a local 2.5D elevation map of the environment with the robot centered in the grid [1]. The uncertainty of elevation values is predicted through sensor noise models and localization drift estimates and can be very valuable for motion planning [1].

Completing sparse elevation maps: Relatively few works consider the challenge of completing sparse robotic elevation maps. One early work by Kolter *et al.* [8] proposes a non-parametric algorithm based on texture synthesis methods to fill-in the missing terrain portions. This method is comparable to patch-based methods in image inpainting as it searches for patches in a library that closely resemble the partially occluded terrain patch with the additional enforcement of geometric constraints with respect to line of sight for the posterior distribution over the missing region. In recent years, data-driven methods were explored for completing sparse DEMs. Qiu *et al.* [17] leverage a deep convolutional GAN to fill the voids of large-scale geospatial data DEM with the loss function being composed of pixel-wise, contextual and perceptual constraints. For mobile robots however, there is even more information available than just map reconstruction losses as the robot pose is estimated and recorded throughout the trajectory [9]. The robot pose hints at a likely local terrain shape which can be leveraged as part of a weakly-supervised KKT-loss term. This is in particular relevant for partially flexible terrain for which lidar measurements over-estimate the support height. The supervision through recorded poses is extended by training a pose regressor leading to a pose prediction-loss for areas of the terrain map not traversed by the robot. The reconstruction loss itself is trained by estimating a ground-truth map offline which also includes future measurements and thus is more complete and likely accurate than the current DEM [9]. The confidence in occupancy of parts of the map can furthermore be used to actively guide the depth measurement sensor to explore sparse areas of the map [20].

III. METHODOLOGY

First, we give a brief overview of our method as visualized in Fig. 1: we consider sparse 2.5D elevation maps and strive to fill-in the missing elevation information to match the ground-truth DEM as accurately as possible. We propose to use an U-Net [6]-like neural network to inpaint the occluded 2.5D elevation map with the inputs consisting of the DEM and a binary occlusion mask. Supervised learning requires us to know ground-truth data to compute a training loss and after back-propagation optimize the neural network weights. As complete and accurate ground-truth information is rarely available for real-world datasets, we introduce a self-supervised method which leverages ray casting to further occlude the DEM given by the robot sensors. This allows us to compute a MSE and Total Variation (TV) loss between the occluded DEM and the less occluded input DEM.

A. Problem Statement

We consider a 2.5D DEM $\mathbf{m}_{\text{occ}} \in \mathcal{R}^{n \times m}$ of a terrain patch. The terrain patch of length l and width w is discretized at a

Fig. 1. Self-supervised learning using artificial occlusion generated via ray casting for inpainting occlusion in terrain maps: partially occluded DEM are sampled via robot-centric elevation mapping. Artificial occlusion is added by performing ray casting from a random vantage point which allows a neural network to be trained to reconstruct the DEM without occlusion. The composed DEM consists of patching the original occluded DEM with the reconstruction in the occluded area. The ray casting graphic is adapted from Kolter *et al.* [8].

resolution r to compose a regular grid of dimension $n \times m$, which encodes the elevation at every grid cell. However, some elevation information is missing because of geometrical occlusion or other reasons such as degraded depth measurements. We create a binary occlusion mask $\mathbf{M}_{\text{occ}} \in [0; 1]^{n \times m}$ which assigns a value of 1 to grid cells with missing elevation information and 0 to grid cells with known elevations. We define our problem as such, that we want to estimate a reconstructed DEM $\mathbf{m}_{\text{rec}}$ which matches the ground-truth DEM $\mathbf{m}_{\text{gt}}$ as closely as possible. During inference, we construct a composition $\mathbf{m}_{\text{comp}}$ between the original input $\mathbf{m}_{\text{occ}}$ for the non-occluded grid cells and the reconstruction $\mathbf{m}_{\text{rec}}$ for the occluded grid cells.

B. Ray Casting

We use ray casting to compute occluded DEMs based on the synthetic ground-truth DEMs in addition to generating artificial occlusion in the framework of self-supervised learning. We developed a lightweight C++ component to perform fast ray casting of an entire grid map from a given vantage point. The ray casting algorithm takes a DEM and a vantage point $\mathbf{x}_v$ as inputs and then iterates through every cell in the grid and checks for each cell whether the cell is visible from the vantage point.

C. Self-Supervised Learning

We propose a self-supervised learning methodology to train on real-world datasets without ground-truth to learn to fill-in missing elevation information of DEMs. We reformulate the training setup by using our partly occluded real-world elevation map as a target and creating additional artificial occlusion which is used as an input into the neural network. We considered a dilation of the already occluded area in addition to randomly occluding pixels. However, this approach does not render realistic and diverse occlusion masks. We decided to add artificial occlusion by ray casting from a randomly sampled vantage point.

We employ an iterative algorithm to generate useful artificial occlusion masks with occlusion ratios between 0.1% and 50% and choose the elevation offset o of the vantage point accordingly. First, we sample a random vantage point from the grid with a uniform distribution. Then after a random elevation offset for the vantage point $o \sim \mathcal{U}(o_{\min}, o_{\max})$ is sampled, we perform ray casting for the chosen vantage point and evaluate the resulting occlusion ratio (e.g., number of occluded grid cells over number of total grid cells).

We require that the occlusion ratio r_{occ} lies within the interval $[r_{\text{occ,min}}, r_{\text{occ,max}}]$. This threshold interval was tuned in a selection study and set to [0.1%, 50%]. If the occlusion ratio does not lie within the interval, we sample a new elevation offset. This time, we adjust the range of the uniform distribution: if $r_{\text{occ}} > r_{\text{occ,max}}$, we increase $o_{\min}$ to the previously sampled elevation offset o , and if $r_{\text{occ}} < r_{\text{occ,min}}$, we decrease $o_{\max}$ to the previously sampled elevation offset o . If $\|o_{\max} - o_{\min}\| < 0.05$ m, we enforce a minimal sampling range by respectively decreasing $o_{\min}$ by 0.05 m or increasing $o_{\max}$ by 0.05 m. We repeat this algorithm until either we have satisfied our occlusion ratio constraint $[r_{\text{occ,min}}, r_{\text{occ,max}}]$ or we have reached the maximum number of iterations (15). We reformulate our loss function to only consider the areas of artificial occlusion and ignore the already occluded areas in the target DEM.

D. Model

We adopt an U-Net [6] as the architecture of our neural network which is often used in the literature for image and video inpainting [2], [5]. The input is composed of two channels: the occluded elevation map $\mathbf{m}_{\text{occ}}$ and the binary occlusion mask $\mathbf{M}_{\text{occ}}$. As the occluded DEM contains missing elevation values represented as NaNs computationally, we replace them with a floating point number by linearly interpolating for synthetic datasets and with the constant 0.0 for real-world datasets after a selection study. A reason could be that linear interpolation is performing worse for the higher occlusion ratios of the real-world

datasets. We also implement input and output normalization: we compute the mean of the non-occluded elevation values of a DEM and subtract this mean from each elevation value in the grid. After receiving the output of the model, we add this mean back to each elevation value before computing the loss. We adapt our model from the vanilla U-Net architecture and make slight adjustments to the number of max-pooling steps because we are working with smaller input images (64x64px instead of 572x572px) than in the original paper [6] and we want to keep the network as lightweight as possible. Thus, we limit the number of max-pooling operations to 3 (instead of 5 in the original paper) and treat the number of channels in each hidden dimension as a hyperparameter for which we select 64, 128, and 256 channels for our hidden dimensions.

E. Loss Function

We enforce a MSE loss between the ground-truth DEM $\mathbf{m}_{\text{gt}}$ and the reconstructed DEM $\mathbf{m}_{\text{rec}}$ as our pixel-by-pixel reconstruction loss. We separate the MSE loss for the occluded region of the DEM MSE_{occ} and the non-occluded region of the DEM MSE_{noocc} . A TV loss $\mathcal{L}_{\text{tv}}$ [21] has shown to be valuable as a smoothing penalty on the reconstruction of the occluded area. We compute the final loss as a weighted sum of all loss components with weights of 1 and 10 for the MSE loss of the non-occluded and occluded area respectively, and scale of 0.1 for the TV loss.

F. Training

We base our software to train the models on PyTorch [22]. After every training epoch on the training set, a pixel-by-pixel MSE loss is evaluated for the occluded area $\mathcal{L}_{\text{MSE,occ}}$ as our validation loss. The training is stopped if either a maximum number of epochs is reached (100) or the validation loss did not improve during a specified number of epochs (50). We use the Adam [23] optimizer with a learning rate of 0.0001, a weight decay of 0.001, and the beta coefficients (0.9, 0.999).

IV. EXPERIMENTS AND RESULTS

We evaluate our method both quantitatively and qualitatively on multiple synthetic and real-world datasets. More specifically, we compare the performance of the self-supervised learning approach to several trivial baselines such as linear [24] and cubic interpolation and traditional inpainting methods with hand-tuned heuristics such as Navier-Stokes [13] and Telea [25].

A. Datasets

Synthetic dataset: We evaluate our methods on synthetic datasets using three different terrain types: hills, standard stairs, and random boxes. This allows us to train with supervised learning but also to benchmark our methods against an accurate ground-truth, as it is usually very hard to acquire a matching ground-truth for a large dataset of real-world occluded terrain maps. We generate terrain maps of size 64x64px and with a resolution of 0.04 m/px. We attribute 25,000 samples to the training, and 2,500 samples each to the validation and test set. The hills

terrain is generated using a distribution of Perlin noise [26]. We also implement terrains with highly regular patterns in standard stairs and random boxes. While the standard stairs terrain is generated by stacking boxes in a uniform direction with a regular step size using fixed box dimension, boxes are stacked at various angles with randomly sampled dimensions for the random boxes terrain. After spawning the terrain map for the chosen terrain type, we perform ray casting to gather the occluded elevation map $\mathbf{m}_{\text{occ}}$ and the binary occlusion mask $\mathbf{M}_{\text{occ}}$. We sample a random robot position from a uniform distribution $x \sim \mathcal{U}(-1.25 \text{ m}, 1.25 \text{ m})$ to perform ray casting. We subsequently infer the elevation of the robot position from the DEM at the corresponding pixel and sample a random elevation offset of the vantage point for ray casting $o_{\text{vp}} \sim \mathcal{U}(0.2 \text{ m}, 0.5 \text{ m})$ for the hills terrain and $o_{\text{vp}} \sim \mathcal{U}(0.2 \text{ m}, 0.3 \text{ m})$ for the standard stairs and random boxes terrains respectively.

ANYmal datasets: We consider datasets recorded using the ANYmal [27] legged robot in three different terrains. The ANYmal robot is equipped with a dome LiDAR sensor and an IMU to perceive 3D point clouds of its environment and odometry information respectively. 2.5D DEMs of size 300x300px and resolution 4 cm are derived using robot-centric elevation mapping [1], [28]. The terrains include several datasets recorded while the robot is traversing on a staircase and on an obstacle course at ETH Zürich, and one dataset capturing the subterranean exploration of the Gonzen mine in Switzerland [29]. For the ANYmal datasets, we use approximately 80% of the subgrid DEMs for the training set and 10% each for the validation and test sets respectively. The ETH stairs, ETH obstacles course, and Gonzen mine datasets contain in total 26,233, 37,274 and 16,459 samples. We divide the 300x300px DEMs into 16 subgrids à 75x75px each, to increase generalization capabilities, reduce the GPU memory requirements and exclude empty parts. We uploaded a video of the inference of our method on the Gonzen mine dataset [29] to YouTube.¹

Tenerife Lunar Analogue dataset: We evaluate our methods on a lunar analogue dataset which was collected during a field test campaign in June 2017 at Minas de San José on the Teide Volcano on the island of Tenerife using the Heavy Duty Planetary Rover (HDPR) [11], a lab rover testbed with resemblance to the Rosalind Franklin rover [30] flying to Mars in 2022 during the ExoMars mission [31]. The dataset contains Global Navigation Satellite System (GNSS) recordings, the rovers raw inertial data, and images from three stereo cameras captured during several traverses over the duration of multiple days including a variety of lighting conditions and different pre-planned paths. We choose elevation maps generated during a traverse on an isolated side track as our test set because they are completely unseen during all other traverses. We apply the GA SLAM [32] technique to extract DEMs from the raw dataset. As we want to minimize all errors in the local maps introduced by drifts in the localization, we use the ground-truth robot pose provided by the GNSS as opposed to the visual odometry from the LocCam to translate the local map to the next time step and rotate the incoming point

¹[Online]. Available: <https://youtu.be/2Khxeto62LQ>

TABLE I
RESULTS FOR SYNTHETIC DATASETS REPORTING THE RECONSTRUCTION PERFORMANCE FOR THE OCCLUDED AREA AVERAGED OVER THREE RANDOM SEEDS

Method	Terrain	$\ell_{1, \text{rec}, \text{occ}}$	$MSE_{\text{rec}, \text{occ}}$	$PSNR_{\text{rec}, \text{occ}}$	$SSIM_{\text{rec}}$	$SSIM_{\text{comp}}$
Linear interpolation [24]	Hills	0.0123	0.00052	37.95	0.99261	0.99261
Cubic interpolation		0.0103	0.00046	38.54	0.99326	0.99326
Telea [25]		0.0231	0.00223	31.65	0.98686	0.98686
Navier-Stokes [13]		0.0226	0.00302	30.33	0.98882	0.98882
Supervised		0.0121 $\pm$ 0.0005	0.00033 $\pm$ 0.00001	40.0 $\pm$ 0.2	0.9888 $\pm$ 0.0004	0.9950 $\pm$ 0.0002
Self-supervised		0.0116 $\pm$ 0.0003	0.00035 $\pm$ 0.00001	39.7 $\pm$ 0.2	0.987 $\pm$ 0.001	0.9949 $\pm$ 0.0003
Linear interpolation [24]	Stairs	0.0534	0.00472	32.43	0.97273	0.97273
Cubic interpolation		0.0633	0.00937	29.44	0.95952	0.95952
Telea [25]		0.0608	0.00752	30.40	0.95543	0.95538
Navier-Stokes [13]		0.0601	0.00877	29.73	0.95399	0.95399
Supervised		0.012 $\pm$ 0.003	0.0003 $\pm$ 0.0001	44 $\pm$ 1	0.992 $\pm$ 0.002	0.9982 $\pm$ 0.0005
Self-supervised		0.013 $\pm$ 0.002	0.0005 $\pm$ 0.0001	42.2 $\pm$ 0.8	0.992 $\pm$ 0.002	0.9980 $\pm$ 0.0003
Linear interpolation [24]	Random Boxes	0.0660	0.00721	09.35	0.78321	0.78321
Cubic interpolation		0.0674	0.00784	08.98	0.77970	0.77970
Telea [25]		0.0678	0.00794	08.93	0.7784	0.7784
Navier-Stokes [13]		0.0655	0.00761	09.12	0.7806	0.7806
Supervised		0.0118 $\pm$ 0.0009	0.00078 $\pm$ 0.00003	19.0 $\pm$ 0.1	0.88 $\pm$ 0.03	0.959 $\pm$ 0.004
Self-supervised		0.0151 $\pm$ 0.0007	0.0011 $\pm$ 0.0001	17.7 $\pm$ 0.5	0.72 $\pm$ 0.09	0.949 $\pm$ 0.006

The best performing method for each dataset / terrain and error metric.

cloud before stereo processing. We only select keyframe elevation maps with a Peak signal-to-noise ratio (PSNR) smaller than 50 dB between the current and the last accepted sample which drastically decreases the computational load during training and also prevents overfitting issues which would otherwise arise from having too many similar samples. We divide the 600x600px DEMs into 64 subgrids à 75x75px. This policy results in 42,600 samples for the training set, 1,000 samples for the validation set and 7,950 samples for the test set.

B. Evaluation Metrics

Similar to other publications on the topic of image inpainting such as [5], [33], we state the ℓ_1 loss, the PSNR and the Structural Similarity Index Measure (SSIM) [34]. We use the separated test set for all evaluation results.

The PSNR is a function of the total MSE loss of the occluded area and the maximum dynamic range of the grid-values L determined by computing the delta between the highest and lowest elevation of the ground-truth DEMs of the entire test dataset. We only report the evaluation metrics for the occluded area as we argue that a composed DEM incorporating the non-occluded parts of the input DEM $\mathbf{m}_{\text{comp}}$ can be easily created and thus the reconstruction capability of the non-occluded area by the model is not essential.

We compute the SSIM between the ground-truth DEM $\mathbf{m}_{\text{gt}}$ and the reconstructed DEM $\mathbf{m}_{\text{rec}}$ as $SSIM_{\text{rec}}$ and similarly between the ground-truth and the composed DEM $\mathbf{m}_{\text{comp}}$ as $SSIM_{\text{comp}}$. The metric is quantified over the entire DEM and not just the occluded region as it is not defined for irregular patches. We use the constants $C_1 = (k_1 L)^2$, $C_2 = (k_2 L)^2$ and $C_3 = C_2/2$, where L signifies the dynamic range as previously used for the PSNR, and $k_1 = 0.01$ and $k_2 = 0.03$. We calculate the average μ and the variance σ^2 with a Gaussian filter with a 1D kernel of size 11 and sigma 1.5 for every DEM separately. We set the weight exponents α , β , and γ all to one. The SSIM metric can only be evaluated on complete DEMs and thus requires knowledge of the entire ground-truth.

C. Experiments on Synthetic Datasets

We evaluate both baseline methods, supervised and self-supervised approaches on the three synthetic datasets and report the results in Table I including the the SSIM for both the reconstructed and the composed DEM. We train and evaluate experiments on three different random seeds and report a confidence interval with mean and standard deviation for those experiments. For supervised learning, we directly use the non-occluded ground-truth as our training targets, while for self-supervised learning we adopt the occluded DEM from the dataset as our training target and add artificial occlusion using ray casting.

We show 24%, 90%, and 85% reductions in MSE error for the occluded area compared best performing baseline approach on each dataset using self-supervised learning based on generating artificial occlusion with ray casting for the synthetic hills, standard stairs, and random boxes terrains respectively. Substantial improvements can also be noted for the ℓ_1 and SSIM error metrics. We would like to point out larger improvements for highly-structured terrain such as standard stairs and random boxes compared to the baseline methods. Surprisingly, the performance of self-supervised learning is only slightly worse (approx. 0.3 dB-2.3 dB in PSNR) compared to fully supervised learning. While (cubic) interpolation baseline methods perform very well on the smooth hills terrain, their performance is much worse on more structured terrain and similar to Telea [25] and Navier-Stokes [13].

D. Experiments on Real-World Datasets

We evaluate our proposed self-supervised learning approach on several real-world datasets. As complete ground-truth DEMs are hard to acquire in the real world, we evaluate on the test set of the real-world data with artificial occlusion created using ray casting. We run experiments using five different random seeds and report a confidence interval with mean and standard deviation.

Evaluating quantitatively on real-world datasets containing artificial occlusion generated with ray casting, we state a decrease of 64%, 52%, 82%, and 57% in MSE error compared

Fig. 2. Snapshots of inference using the trained neural network on the Gonzen mine dataset recorded with the ANYmal legged robot. The occluded DEM is marked in a color scale from white to blue (lowest to highest elevation), while the reconstruction is visualized from white to red with a slight opacity. To improve the clarity of presentation, we only inpaint subgrids which are less than 85% occluded as the neural network requires sufficient input elevation information.

TABLE II

RESULTS FOR REAL-WORLD DATASETS EVALUATED USING ARTIFICIAL OCCLUSION GENERATED WITH RAY CASTING ON THE TEST SET AVERAGED OVER FIVE RANDOM SEEDS FOR THE BASELINE METHODS LINEAR- / CUBIC INTERPOLATION, TELEA [25] AND NAVIER-STOKES [13], AND SELF-SUPERVISED LEARNING. THE CHOSEN UNIT OF ELEVATION IS METERS.

Method	Terrain	ℓ_1 _{occ}	MSE _{occ}	$PSNR$ _{R_{occ}}
Linear	ETH St.	0.1014	0.05544	25.98
Cubic		0.1193	0.08715	24.02
Telea		0.0958	0.04980	26.45
Navier-St.		0.0956	0.05339	26.14
Self-super.		0.071 ± 0.004	0.018 ± 0.001	30.9 ± 0.3
Linear	Obstacles	0.2981	0.34646	23.77
Cubic		0.3453	0.54131	21.83
Telea		0.2873	0.35048	23.72
Navier-St.		0.2801	0.34751	23.76
Self-super.		0.173 ± 0.002	0.166 ± 0.003	26.97 ± 0.07
Linear	Gonzen	0.1977	0.14473	14.87
Cubic		0.2242	0.20689	13.32
Telea		0.2035	0.16535	14.29
Navier-St.		0.2020	0.16787	14.23
Self-super.		0.081 ± 0.002	0.026 ± 0.001	22.3 ± 0.2
Linear	Tenerife	0.0881	0.02577	40.81
Cubic		0.0929	0.03106	40.00
Telea		0.0898	0.02714	40.58
Navier-St.		0.0881	0.02705	40.60
Self-super.		0.0502 ± 0.0009	0.0110 ± 0.0002	44.51 ± 0.09

The best performing method for each dataset / terrain and error metric.

to the best performing baseline approach on each dataset using self-supervised learning for the ANYmal ETH stairs, obstacles course, Gonzen mine [29], and Tenerife lunar analogue datasets as listed in Table II. The ℓ_1 and PSNR error metrics are substantially improved likewise. When we analyse error distribution of our self-supervised approach, we note that 75% of samples have an ℓ_1 error smaller than 5 cm for the ETH stairs, 4 cm for the ETH obstacle course, 4 cm for the Gonzen mine, and 3 cm for the Tenerife dataset. We visualize samples of an inference with full-size DEMs on the Gonzen mine dataset in Fig. 2.

E. Qualitative Comparison

We visualize qualitative subgrid samples comparing our self-supervised learning approach with the ground-truth and the Telea [25] baseline method in Fig. 3. Please note that the ground-truth is already partially occluded and we applied additional artificial occlusion using ray casting to generate our input DEM. In general, we can say that our approach is able to reconstruct the DEM with more details and at a lower error compared

to Telea [25]. In particular, our method is showing promising indications in leveraging not only the information contained in the input DEM, but also of the geometric constraints concerning line of sight by estimating that if a patch is fully occluded, it probably has a lower elevation than the surrounding area as seen in the reconstruction in the upper right corner of the third row. Here, we would like to mention one potential drawback in using a self-supervised approach when training on real-world data: as the neural network is trained on data which contains noise, it will also learn to reconstruct similar noise patterns for the occluded areas.

F. Run-Time

We evaluate the run-time of some baselines methods on CPU and the proposed U-Net neural network on CPU and GPU. The methods running on CPU had access to one core of a 2.3 GHz Intel Core i9 processor in a Mid-2019 MacBook Pro. We used an NVIDIA TITAN Xp GPU to evaluate the run-time of the neural network. The scenario we use to evaluate the run-time is inspired by the ANYmal datasets: We assume occluded DEMs of size 300x300px as our input. We divide the DEM into 16 subgrids each of size 75x75px and subsequently downsample the subgrids to 64x64px. We then pass the DEM in two batches each of batch size 8 to the neural network and report the run-time for the processing of the entire full-size DEM. The baseline methods Telea [25] and Navier-Stokes [13] can be run at around 50 Hz and our proposed U-Net with a sampling rate of 3.5 Hz on CPU and 30 Hz on GPU for a full-size DEM enabling real-time inference on autonomous mobile robots.

G. Preliminary Study on Motion Planning

As part of a preliminary investigation into using our approach for navigation, we utilize the completed DEMs as an input for motion planning for the ANYmal legged robot in a simulated environment. The elevation maps are inpainted with a neural network trained on the ETH Stairs dataset. We point out the impressive generalization from the real-world stairs dataset to the simulation environment depicted in Fig. 4. The locomotion [35] and navigation [36] components used in simulation were used by team CERBERUS during the DARPA Subterranean Challenge [29]. First, we execute motion planning using

Fig. 3. Qualitative samples comparing the Telea [25] baseline method and self-supervised learning on various real-world test sets. The ground-truth DEM describes the local, partially occluded terrain map subgrid sampled by the robot. We add artificial occlusion with ray casting to derive the occluded DEM which serves as the input to both inpainting methods. We visualize the error compared to the partially occluded ground-truth map for Telea and our method. All units are in meter. **First row - ETH stairs:** the Telea reconstruction is blurring the stairs and additionally inpainting two holes with lower elevation values on the left side. **Second row - Tenerife lunar analogue:** our method reconstructs the obstacle in the center with higher accuracy leveraging knowledge about the terrain characteristics. **Third row - Gonzen mine:** self-supervised learning respects the geometric constraints of line of sight and reconstructs the occluded area in the top-right corner to have a lower elevation.

Fig. 4. Samples visualizing motion planning for an ANYmal legged robot in a simulated environment. The blue line represents the vector from the start pose (purple arrow) on the right side to the goal on the left side. **First row:** incomplete DEMs are used for motion planning. **Second row:** DEMs reconstructed by a neural network trained with self-supervision are employed for motion planning.

sensor-generated, incomplete DEMs while not planning any path through areas without elevation information. Currently, this baseline scenario is very common in literature [36], [37]. Our experiments show that this approach cannot plan far-ahead and that the paths are not cost-optimized over the entire map. Accordingly, the robot needs to stop and wait frequently until the path re-planning is completed. We show examples for this behaviour in the first row of Fig. 4. On the contrary, paths planned based on completed DEMs reach further ahead and only need to be slightly adjusted when the belief of the occluded area

is updated. We present this behaviour in the second row of Fig. 4 and in a video comparison on YouTube.²

V. CONCLUSION

This work proposes a self-supervised learning approach for completing sparse 2.5D terrain maps. The method leverages artificial occlusion generated with an iterative ray casting algorithm to train a neural network on real-world data with an incomplete

²[Online]. Available: https://youtu.be/_SAZm7tMMUI

ground-truth. We have evaluated our self-supervised learning approach on a variety of synthetic and real-world datasets. Evaluating quantitatively on real-world datasets containing artificial occlusion generated with ray casting, we state a decrease of between 52% and 82% in MSE error compared to the respective best performing baseline approach using self-supervised learning based on ray casting. Qualitative samples show that our proposed method is able to both leverage terrain characteristics but also information inherently encoded in the occlusion masks such as full-filling the geometric constraints of line of sight indicating the advantage compared to traditional baseline methods. The neural network is able to run with sampling rates of 3.5 Hz and 30 Hz respectively for CPU and GPU inference. We utilize the reconstructed DEMs for motion planning in a proof-of-concept. We observe that the path can be planned further ahead using reconstructed elevation maps compared to planning with incomplete elevation maps. Consequently, the robots needs to stop less frequently and wait less for re-planning to finish.

For future work, we could envision several interesting research opportunities: Instead of using ray casting, it could be worthwhile to train a separate GAN network to learn realistic occlusion patterns to generate additional artificial occlusion as part of the self-supervised learning. We are not able to use perceptual and styles losses [16] for our self-supervised learning approach as the adopted pretrained VGG-16 [38] requires a complete ground-truth. As a replacement, we would like to encourage the addition of a discriminator which could be trained to distinguish between reconstructed and ground-truth DEM. Finally, it is crucial to explore model and data uncertainty estimation methods to predict the variance in the neural network predictions to enable safe and optimal motion planning. The cost function optimized during path planning should be influenced by the uncertainty estimation of each grid cell.

REFERENCES

- [1] P. Fankhauser, M. Bloesch, and M. Hutter, "Probabilistic terrain mapping for mobile robots with uncertain localization," *IEEE Robot. Automat. Lett.*, vol. 3, no. 4, pp. 3019–3026, Oct. 2018.
- [2] N. Wang *et al.*, "Multistage attention network for image inpainting," *Pattern Recognit.*, vol. 106, 2020, Art. no. 107448.
- [3] O. Elharrouss *et al.*, "Image inpainting: A review," *Neural Process. Lett.*, vol. 51, no. 2, pp. 2007–2028, 2020.
- [4] U. Demir and G. Unal, "Patch-based image inpainting with generative adversarial networks," 2018, *arXiv:1803.07422*.
- [5] G. Liu *et al.*, "Image inpainting for irregular holes using partial convolutions," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 85–100.
- [6] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervention*, 2015, pp. 234–241.
- [7] S. Thrun and T. M. Mitchell, "Lifelong robot learning," *Robot. Auton. Syst.*, vol. 15, no. 1/2, pp. 25–46, 1995.
- [8] J. Z. Kolter, Y. Kim, and A. Y. Ng, "Stereo vision and terrain modeling for quadruped robots," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2009, pp. 1557–1564.
- [9] V. Šalanský, K. Zimmermann, T. Petříček, and T. Svoboda, "Pose consistency KKT-loss for weakly supervised learning of robot-terrain interaction model," *IEEE Robot. Automat. Lett.*, vol. 6, no. 3, pp. 5477–5484, Jul. 2021.
- [10] M. Hutter *et al.*, "ANYmal—A highly mobile and dynamic quadrupedal robot," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 38–44.
- [11] E. Boukas *et al.*, "HDPR: A mobile testbed testbed for current and future rover technologies," in *Proc. 13th Int. Symp. Artif. Intell. Robot. Automat. Space*, 2016, pp. 1–6.
- [12] S. Mehra *et al.*, "From textural inpainting to deep generative models: An extensive survey of image inpainting techniques," *J. Comput. Sci.*, vol. 16, pp. 35–49, 2020.
- [13] M. Ebrahimi and E. Lunasin, "The navier-stokes-voight model for image inpainting," *IMA J. Appl. Math.*, vol. 78, no. 5, pp. 869–894, 2013.
- [14] C. Barnes *et al.*, "PatchMatch: A randomized correspondence algorithm for structural image editing," *ACM Trans. Graph.*, vol. 28, no. 3, 2009, Art. no. 24.
- [15] Y. Song *et al.*, "Contextual-based image inpainting: Infer, match, and translate," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 3–19.
- [16] L. A. Gatys, A. S. Ecker, and M. Bethge, "A neural algorithm of artistic style," Poster presented at 16th Annu. Meeting Vision Sci. Soc. (VSS 2016), St. Pete Beach, FL, USA, 2016.
- [17] Z. Qiu, L. Yue, and X. Liu, "Void filling of digital elevation models with a terrain texture learning model based on generative adversarial networks," *Remote Sens.*, vol. 11, no. 23, 2019, Art. no. 2829.
- [18] X. Zhan, X. Pan, B. Dai, Z. Liu, D. Lin, and C. C. Loy, "Self-supervised scene de-occlusion," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 3783–3791.
- [19] A. Dai, C. Diller, and M. Nießner, "SG-NN: Sparse generative neural networks for self-supervised scene completion of RGB-D scans," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 849–858.
- [20] K. Zimmermann, T. Petricek, V. Šalanský, and T. Svoboda, "Learning for active 3D mapping," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 1539–1547.
- [21] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 694–711.
- [22] A. Paszke *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 8024–8035.
- [23] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [24] T. Nouza, "Safe adaptive traversability learning for mobile robots," Ph.D. dissertation, Master's thesis, Dept. Cybern., Fac. Elect. Eng., Czech Technical University in Prague, Prague, Czechia, 2014. [Online]. Available: <https://core.ac.uk/download/pdf/47179169.pdf>
- [25] A. Telea, "An image inpainting technique based on the fast marching method," *J. Graph. Tools*, vol. 9, no. 1, pp. 23–34, 2004.
- [26] A. Lagae *et al.*, "A survey of procedural noise functions," *Comput. Graph. Forum*, vol. 29, no. 8, pp. 2579–2600, 2010.
- [27] M. Hutter *et al.*, "Anymal-toward legged robots for harsh environments," *Adv. Robot.*, vol. 31, no. 17, pp. 918–931, 2017.
- [28] P. Fankhauser *et al.*, "Robot-centric elevation mapping with uncertainty estimates," in *Proc. Mobile Service Robot.*, 2014, pp. 433–440.
- [29] M. Tranzatto *et al.*, "CERBERUS: Autonomous legged and aerial robotic exploration in the tunnel and urban circuits of the DARPA subterranean challenge," *J. Field Robot.*, 2021. [Online]. Available: <https://www.research-collection.ethz.ch/handle/20.500.11850/489726>
- [30] J. L. Vago *et al.*, "Habitability on early mars and the search for biosignatures with the exomars rover," *Astrobiology*, vol. 17, no. 6/7, pp. 471–510, 2017.
- [31] M. Azkarate, L. Gerdes, L. Joudrier, and C. J. Pérez-del-Pulgar, "A GNC architecture for planetary rovers with autonomous navigation," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2020, pp. 3003–3009.
- [32] D. Geromichalos *et al.*, "Slam for autonomous planetary rovers with global localization," *J. Field Robot.*, vol. 37, no. 5, pp. 830–847, 2020.
- [33] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, and T. S. Huang, "Generative image inpainting with contextual attention," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 5505–5514.
- [34] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [35] T. Miki *et al.*, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Sci. Robot.*, 2022.
- [36] L. Wellhausen and M. Hutter, "Rough terrain navigation for legged robots using reachability planning and template learning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2021, pp. 6914–6921.
- [37] L. Gerdes *et al.*, "Efficient autonomous navigation for planetary rovers with limited resources," *J. Field Robot.*, vol. 37, no. 7, pp. 1153–1170, 2020.
- [38] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014, *arXiv:1409.1556*.