

Conditional GANs for Sonar Image Filtering with Applications to Underwater Occupancy Mapping

Tianxiang Lin, Akshay Hinduja, Mohamad Qadri, and Michael Kaess

Abstract—Underwater robots typically rely on acoustic sensors like sonar to perceive their surroundings. However, these sensors are often inundated with multiple sources and types of noise, which makes using raw data for any meaningful inference with features, objects, or boundary returns very difficult. While several conventional methods of dealing with noise exist, their success rates are unsatisfactory. This paper presents a novel application of conditional Generative Adversarial Networks (cGANs) to train a model to produce noise-free sonar images, outperforming several conventional filtering methods. Estimating free space is crucial for autonomous robots performing active exploration and mapping. Thus, we apply our approach to the task of underwater occupancy mapping and show superior free and occupied space inference when compared to conventional methods.

I. INTRODUCTION

Autonomous Underwater Vehicles (AUVs) are useful in a broad range of applications that are otherwise tedious or potentially dangerous for humans to perform. A major aspect of AUVs is to perform underwater mapping which assists in tasks like subsea infrastructure inspection, ship hull inspection [1], seafloor surveying, and bathymetry [2].

In most conditions, AUVs cannot rely on optical sensors like cameras and laser range scanners due to limitations arising from turbidity and light absorption. Visibility in deep water is often constrained to a few meters (1-2 meters) at best. Thus, acoustic sensors like sonar are better suited for tasks involving mapping and free space estimation, often providing range information from several meters (10 meters and up) depending on the type and frequency of the sonar.

Different tasks need different sonar types. For seafloor mapping, side scan sonar (SSS) is commonly used as in [2], [3], whereas in more structured and complex 3D environments, imaging sonars are found to be more prevalent as shown in [4]–[7]. Teixeira et al. [8] also used an imaging sonar with a concentrator lens in profiling mode, utilizing submaps to obtain occupancy and subsequently perform simultaneous localization and mapping (SLAM).

While sonar is the preferred sensing modality for underwater environments, it is far from ideal due to noise from random and systematic errors. The sources for these errors can be due to the environment, temperature, an imperfection in calibration, cross-talk between transducers and receivers,

The authors are with The Robotics Institute, Carnegie Mellon University, USA. {tianxian, ahinduja, mqadri, kaess}@andrew.cmu.edu

This work was partially supported by the Office of Naval Research grant N00014-21-1-2482.

The authors would like to thank Easton Potokar for sharing real-world sonar data used for training our network.




Fig. 1: Real-world experiments performed in a test tank.




Fig. 2: Real-world sonar image filtering by our method.

multipath reflections, as well as delayed signal returns during vehicle motion [9]. The most common noise pattern observed is speckle noise. It mainly originates from coherent and random patterns of constructive and destructive interference of backscattered signals creating pixels with low and high intensity in the image. This is mainly due to wave propagation characteristics and the material properties of the objects in the frustum of the sonar. It is also a primary source of visual noise for many wave propagation-based sensors [10]–[12].

Reduced noise in sonar images is important for AUVs. It enables the robot to infer its environment better, through observing features or objects and performing SLAM. It is also important in the estimation of free space, as noise in the image can be incorrectly defined as an occupied space. Conventional speckle noise filters operate on the premise that the actual signal and noise are statistically independent and

can be differentiated on this basis. While most conventional filters are capable of performing some degree of noise reduction, often it is at the cost of degradation of important image features such as corners and high contrast boundaries.

The remainder of the paper is organized as follows. After discussing related work in the next section, Section III introduces imaging sonar geometry in both Cartesian and polar space. Section IV subsequently describes our filtering method, followed by an evaluation of the filter performance in Section V. In Section VI we apply our method toward underwater occupancy mapping and show that our method outperforms conventional methods used for sonar denoising for both simulated and real-world sonar image data. Finally, Section VII closes with our concluding remarks and future directions for this work.

II. RELATED WORK

There have been several attempts at filtering sonar images through a multitude of methods. Some of the earliest attempts have been the Lee, Kuan, and Frost filters for synthetic aperture sonar images as described by [13]. The filters' measure of image homogeneity are based on the multiplicative speckle model, which is not always accurate causing failures for small details in images like corners. Lopes et al. [14] improved upon these methods by adding two thresholds on the coefficient of variation, allowing better homogeneous averaging and heterogeneous feature preservation. Variations of wavelet transforms have also performed well in denoising sonar images as seen in [15], [16]. The transforms work by concentrating the original signal and image features in a few large-magnitude wavelet coefficients, where the smaller magnitude coefficients represent the noise. Negahdripour et al. [17] proposed using multiple images taken from a single viewpoint and averaging to obtain refined intensity returns. Anisotropic diffusion has shown success for feature-based methods [18], [19] as well as for dense 3D reconstruction using sonar [20]. Teixeira et al. [8] utilize a two-dimensional Wiener filter [21] to deconvolve a custom point spread function for their sonar. Despite the variety of filtering algorithms, there is no consensus on what method is best since performance varies with the application of interest.

More recently, machine learning-based methods have gained popularity for general image denoising. Lu et al. [22] use a deep convolutional network to reduce speckle noise. Their results show a great improvement for SSS images. Imaging sonars, especially forward-looking sonars would typically only have partial occupation from the viewed object with the rest of the image being empty with noise. Whether this network could handle this phenomenon is unknown.

The work of Isola et al. [23], pix2pix, showed how a conditional adversarial network can be used for image-to-image translation, performing tasks like generating photos from edge maps, style transfer, and background removal to name a few examples. For applications to underwater imaging sonar, Lee et al. [24] have recently used pix2pix to produce imitation sonar images to augment their training data for a fully convolutional object segmentation network.

Our proposed method utilizes pix2pix in a different way. We train the conditional Generative Adversarial Network (cGAN) to generate binary masks of the actual returns. These masks when used to segment the original raw image help preserve the high-intensity objects detected in our sonar image, while eliminating most of the noise in the process as seen in Fig. 2. The effectiveness of our method is showcased by using the filtered images to construct an occupancy map.

III. BACKGROUND: IMAGING SONAR GEOMETRY

Consider a point $P(\theta, r, \phi)$ in the field of view of the imaging sonar, parameterized in the local spherical sonar coordinate system as seen in Fig. 3. Here, θ , r , and ϕ represent the bearing, range, and elevation angle of the point respectively. The conversion of P to the Cartesian frame to point $C(x, y, z)$ and vice versa is

$$C = \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} r \cos \theta \cos \phi \\ r \cos \theta \sin \phi \\ r \sin \theta \end{bmatrix} \quad (1)$$

$$P = \begin{bmatrix} \theta \\ r \\ \phi \end{bmatrix} = \begin{bmatrix} \arctan 2(y, x) \\ \sqrt{x^2 + y^2 + z^2} \\ \arctan 2(z, x^2 + y^2) \end{bmatrix} \quad (2)$$

Imaging sonars generate partial spherical measurements by sending out acoustic signals into a frustum. Time of flight measured from the reflected signals observed by the transceivers provide the range r and bearing θ , of the reflecting surface. However, these measurements are unable to disambiguate the elevation ϕ of the reflected signal. Due to this, all detected returns from a single elevation arc project onto the same pixel of a range image $I(\theta, r)$. For a pixel corresponding to a certain bearing and range in I , the pixel's intensity corresponds to the intensity of the reflected signal from all returns along the elevation.




Fig. 3: Geometry of a single sonar image. Point P is represented by range r , elevation ϕ , and bearing θ . r_{\max} , r_{\min} , ϕ_{\max} , ϕ_{\min} , θ_{\max} , and θ_{\min} are respectively the maximum and minimum ranges, elevation angles, and azimuth angles of the imaging sonar.

Along with elevation ambiguity, shadow zones also appear when objects closer to the sonar obstruct the view of obstacles behind them. Hence, low pixel intensity values in the sonar images do not necessarily mean the absence of obstacles. A pictorial representation of shadow zones and elevation ambiguity can be seen in Fig. 4




Fig. 4: Elevation ambiguity and shadow zone of imaging sonar. Green points, B and C, intercepted by sonar lie in the shadow zone and are thus not visible in the generated range image. Blue points (D and E) and red points (F and G), lying on the same elevation arc, cannot be distinguished in the image.

IV. SONAR IMAGE FILTERING WITH CGANS

We present an approach which uses cGANs to filter imaging sonar images. GANs consist of two networks, a generator G and a discriminator D . The generator's task is to produce data that closely matches the training distribution and the discriminator has the task of distinguishing real images from those produced by the generator. The generator in GANs tries to learn the underlying training image distribution with no additional constraints. Hence, the generators are free to learn any mapping from the noise distribution $p(z)$ to p_{data} , which are the sonar images and corresponding mask pairs.

Pix2pix, like other cGANs, conditions the output of the generator model on an input image through the discriminator model. In cGANs, the training data is in the form of a pair (x, y) , where x and y represent the original image and labeled image that x is conditioned on respectively. The training objective can be represented as

$$Loss = \min_G \max_D \mathcal{L}_c(G, D) + \lambda \mathcal{L}_{L1}(G) \quad (3)$$

where λ is the weight applied to the L1 loss term, and the objective of cGANs \mathcal{L}_c and L1 loss term \mathcal{L}_{L1} are

$$\mathcal{L}_c(G, D) = E_{x, y \sim p_{data}} [\log D(x, y)] + E_{x \sim p_{data}, z \sim p(z)} [\log (1 - D(x, G(x, z)))] \quad (4)$$

$$\mathcal{L}_{L1} = E_{x, y \sim p_{data}, z \sim p(z)} \|y - G(x, z)\|_{L1} \quad (5)$$

The discriminator's task in cGANs remains the same, which is to differentiate real image pairs (x, y) from fake pairs $(x, G(x, z))$. The added task for the generator is to make sure the produced images are nearer to the label y , encouraged by the loss term \mathcal{L}_{L1} .

Our proposed work provides noisy sonar images (as x) and binary masks (as y) to pix2pix to generate close-to-ground truth binary masks as an output. This approach is in inverse to Lee et al. [24], who give binary masks (as x) and noisy sonar images (as y) to the generator to get realistic imitations of sonar images.

Fig. 5 describes our training procedure of how the cGAN is trained to obtain the inferred mask to filter the raw image for real-world data. We provide a series of noisy raw data as the input image, x , and a binary mask of the actual surfaces for the same scenes as our target image, y . When G and

D are trained together with a sufficient number of image pairs, we get a network that produces a well-defined binary image mask as our output for raw input data. This mask is then used to segment out the actual return in the sonar image while eliminating a majority of the image noise. As the generated mask typically has a buffer zone enveloping the object in view, an unsharp filter is then applied to the trained model's outputs to further contrast the reflected signal from the surrounding residual noise.

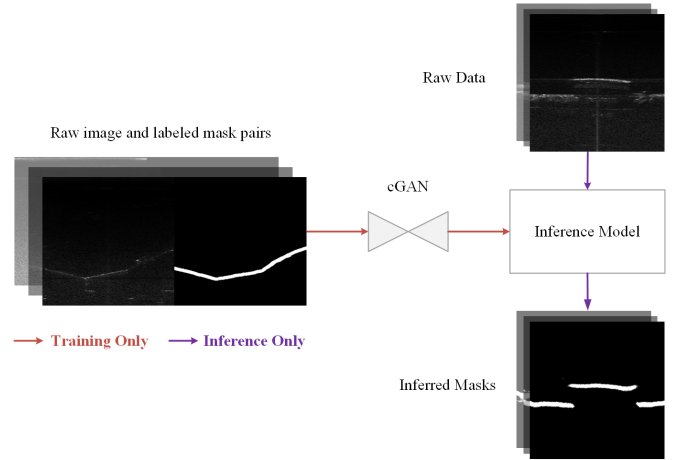


Fig. 5: Conditional GAN for filtering sonar images. Raw and labeled mask pairs are used to train the network. Passing raw data through the network gives a binary mask of the segmented returns.

V. FILTERING RESULTS

In this section, we describe our prepared dataset and training parameters before presenting results for both simulation and real-world sonar images. For our simulated data, we use the HoloOcean simulator [25], [26] and for our real-world data, we use a Bluefin Hovering Autonomous Underwater Vehicle (HAUV) [27] equipped with a 1.2MHz Teledyne/RDI Workhorse Navigator Doppler velocity log (DVL), a Honeywell HG1700 inertial measurement unit (IMU), and a BluePrint Subsea M1200d imaging sonar [28] as seen in Fig. 1. The DVL/IMU navigation solution has a time-based drift, estimated to be 1.1 meters in 20 minutes. We train two models, one for simulated data, and the other for real-world data separately.

To train our simulated data model, we generate 3000 frame pairs of noisy raw data, and noise-free images converted to binary masks. For the real-world data, we randomly select 100 frames from a data log, and manually prepare binary mask images to serve as the target frames. While the HoloOcean simulator is modeled on the M1200d, we observed different noise patterns in simulation and real data, in part due to the small, enclosed nature of the test tank which amplified multipath reflections. Thus attempts to fine-tune a network trained on simulated data with real-world data were unsuccessful. However, a separately trained model worked exceedingly well, even with a modest amount of hand-labeled image pairs. Images used for training were from

different environments than those used for testing, for both simulation and real-world data.

Both models were trained using pix2pix [23] with an Intel(R) Core(TM) i7-7820X CPU running at 3.60GHz and a NVIDIA GeForce RTX 3080 Ti GPU for training. The training time observed for 3000 simulated image frames was roughly 4.5 hours.




Fig. 6: Filtering results on simulated sonar data with added speckle noise. The enhanced Lee and wavelet filter perform best for conventional filters. Our method is able to remove background noise efficiently and give well-defined borders for the returned signal.

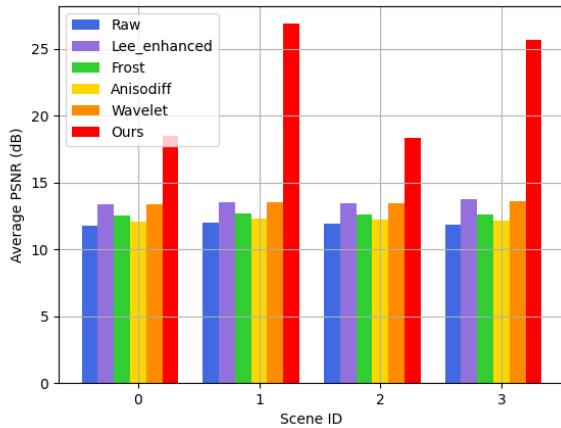


Fig. 7: Comparison of PSNR of our and selected conventional methods on four different scenes from the HoloOcean simulator. Our method consistently shows higher PSNR values for all scenes.

A. Simulation Results

We compare our method against four filters: 1. the original Frost filter, 2. the enhanced Lee filter, 3. Anisotropic diffusion, and 4. wavelet transforms using the VisuShrink threshold [29]. Parameters were adjusted to attain the best performance of each filter.

We choose the peak signal-noise ratio (PSNR) as the error metric to compare the filtration results. As speckle noise is generally formed due to the constructive and destructive interference of reflected signals from largely uneven surfaces, it has been observed that HoloOcean’s simulated images were unable to replicate the effect due to the smooth simulated surfaces on the environment models. To make our comparison more realistic we add normally distributed speckle noise to all images across four scenes from HoloOcean’s list of available environments. Fig. 6 shows the qualitative comparative results of filtration of a sample frame. We note

that our method provides a much sharper distinction of object boundaries, while conventional filters cannot achieve so to a satisfactory degree. In Fig. 7 we compare the baseline filters and our method through the average PSNR per scene. Across all the scenes, our method produces a higher PSNR value, indicating better performance and robustness to noise.

TABLE I: Average run time per frame for filtering methods.

Filters	Lee Enhanced	Frost	Anisodiff	Wavelet	Ours
Time per frame (s)	8.75	22.54	0.03	0.04	0.56

We also compare the time taken for obtaining a filtered image from raw data for different methods as seen in Table. I. All methods were run on the CPU. While our method takes significantly more time for inference compared to methods such as the wavelet transform and anisotropic diffusion, the higher filtering performance is more desirable. When it comes to slow-moving Autonomous Underwater Vehicles (0.5-1.5 knots [27]), the time taken per frame is not as detrimental.

B. Real World Results

For our real-world data, we use the Blueprint Oculus M1200d in its high frequency (2.1MHz) mode which gives a minimum range of 0.1 meters and is configured to a max range of 5 meters, with a 60° horizontal field of view. Fig. 8 compares the filtration performance of our model for a single frame versus the same aforementioned conventional filters. We see that our method is better able to filter surrounding noise and give a sharp contrast to the returned signal.

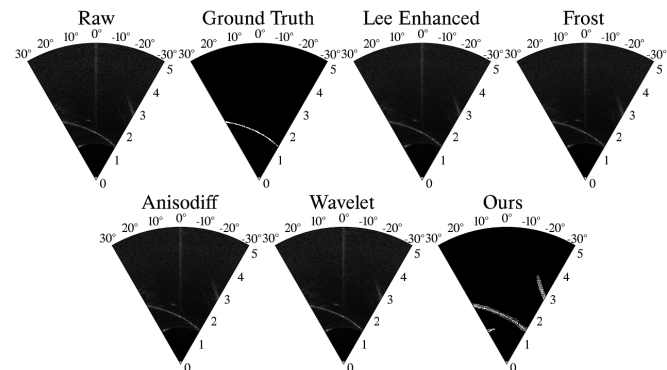


Fig. 8: Filtering results on real-world sonar images from test tank experiments. Ground truth image is prepared by manually removing noise and sharpening the actual return from the tank wall. While contours resulting from multipath propagation exist, our result filters out most noise compared to other conventional methods.

VI. APPLICATION TO OCCUPANCY MAPPING AND EVALUATION

Before an AUV can infer its surroundings for mapping purposes, the estimation of free space is important. For active exploration and mapping, a robot must know the obstacle-free space surrounding itself to plan safe and efficient paths to unexplored regions. Thus, in this section, we apply our method to underwater occupancy mapping. The complete framework we use for underwater occupancy mapping with

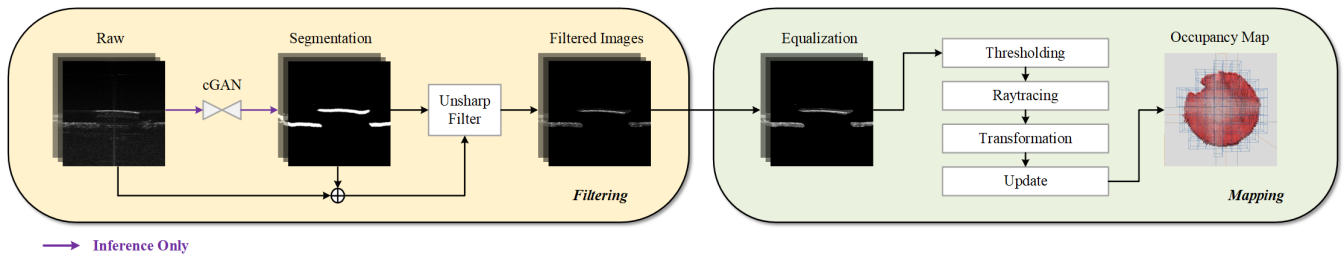


Fig. 9: Underwater occupancy mapping framework. Raw images passed through the network produce a binary mask which is used to segment the original raw image. After passing through an unsharp filter, the image goes through a histogram equalization to generate a higher contrast image for thresholding. Using vehicle poses, we obtain our final occupancy map.

imaging sonar is shown in Fig. 9. Raw sonar image frames are passed through the filtering module to provide denoised images which are then passed to the mapping module. With the filtered range images and known robot poses from the DVL/IMU, the Cartesian coordinates of obstacles can be computed using Equation 1 and utilized in the 3D Bresenham line drawing algorithm [30] to update the values within the corresponding occupancy grid cells.

We begin with the inverse sensor model for our occupancy mapping framework, and then present and evaluate the occupancy maps generated with both simulation and real-world data.

A. Inverse Sensor Model

Once we have the denoised sonar image, histogram equalization is performed [31] which increases the contrast, giving the regions of interest a higher intensity. The intensity of corresponding pixels within sonar range images will determine the occupancy of the grid at the pixel location. We define the occupancy of a grid cell in an occupancy grid map \mathbf{m} in the log-odd form as:


$$l_{x,y} = \log \frac{p(\mathbf{m}_{x,y}|\mathbf{z})}{1 - p(\mathbf{m}_{x,y}|\mathbf{z})} \quad (6)$$

where $\mathbf{m}_{x,y}$ denotes a binary status of a grid cell at position x, y , and \mathbf{z} are the measurements from the sensor from time 1 through time T [32].


For the given sonar range image data, we place higher confidence on the estimated free grid cells rather than occupied cells. This is due to the elevation ambiguity observed in imaging sonars, and the possibility of cells estimated as occupied might not have a significant obstacle along the entire elevation arc. We model our occupancy using the following inverse sensor model:

$$\text{inverse_sensor_model}(\mathbf{m}, \mathbf{x}, \mathbf{z}) = \begin{cases} l_{\text{free}} & \text{if } \mathbf{z} < t, \\ l_{\text{occ}} & \text{otherwise,} \end{cases} \quad (7)$$

where \mathbf{m} represents the status of an occupancy grid cell, \mathbf{x} denotes the state of the sensor itself, \mathbf{z} is the measurement of intensity from sonar images, and t is the threshold used to detect the obstacles within the range image. We use the following probability ratios for free and occupied cells, which performed the best for simulations: $p_{\text{free}} = 0.55$, and $p_{\text{occ}} = 0.05$ [32].



(a) Average False Positive Rate.



(b) Average False Negative Rate.

Fig. 10: Average False Positive (a) and False Negative (b) rates for generated occupancy map using different filtering algorithms and pixel intensity thresholds for simulated data.

B. Occupancy Map Evaluation for Simulation Data

For measuring the accuracy of our occupancy map, we use the assumption that a positive cell value from our inverse model indicates that the grid cell is free. Hence, a false positive occurs when a grid cell is indicated as free but is actually occupied. Similarly, a false negative means that a cell is occupied while it is literally free. We do not make assumptions about unknown grid cells. Using this model we plot the average False Positive Rate (FPR) and False Negative Rate (FNR) for all scenes across different pixel intensity thresholds, t from 0 to 60 for `uint8` images. These results are as seen in Fig. 10. As anticipated, our method performs the best across all threshold values. As the

threshold increases, the enhanced Lee filter gets closer to the performance of our method. But given its filtration run-time, it is not good for online estimation of free space.

The qualitative results of the generated occupancy map for the wavelet transform and our method compared to the ground truth and raw map are shown in Fig. 11, where we see that our method can provide better free and occupied space information compared to the wavelet transform. We demonstrate the result from the wavelet transform since the chosen method performs the best among conventional filtration approaches for simulations.

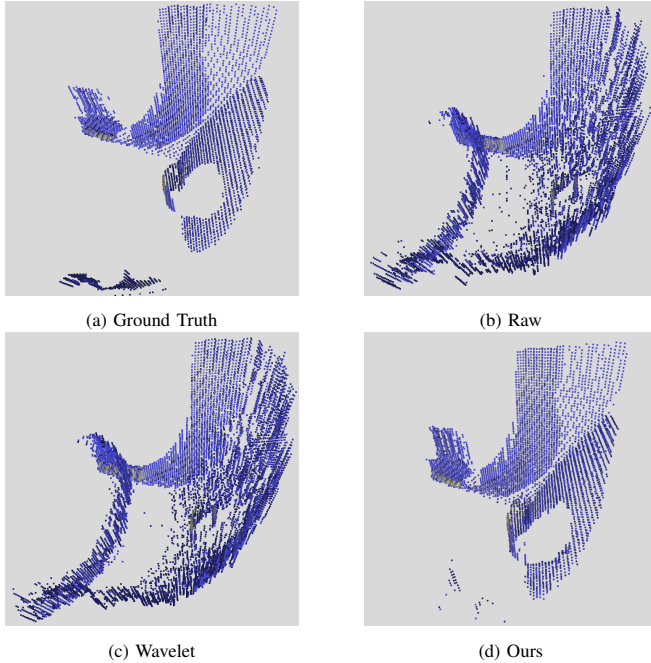


Fig. 11: Final occupancy maps obtained from different filtering strategies using simulated data

C. Occupancy Map Evaluation for Real-world Data

Real-world data collected from our test tank, where a static obstacle board was placed adjacent to the tank wall, as seen in Fig. 1 was used to qualitatively compare the free and occupied space maps generated by our method and the wavelet transform baseline as seen in Fig. 12. We find that our method is more accurate for the experiment environment in determining both, free and occupied cells.

We also experiment to analyze the number of manually labeled image pairs needed for a robust estimation of free space. In Fig. 13 we compare the free space map generated by models using 25, 50, and 100 image pairs for training. We see that the segmentation quality of the trained model improves quickly with the addition of more training pairs, giving more accurate free space estimates. We note that our algorithm is sample efficient as even a small number of hand-labeled data is enough to generate accurate maps of real-world environments like a test tank.




Fig. 12: Occupancy maps generated through wavelet filtering and our method from real data in the test tank. Our method is able to infer available free space more accurately

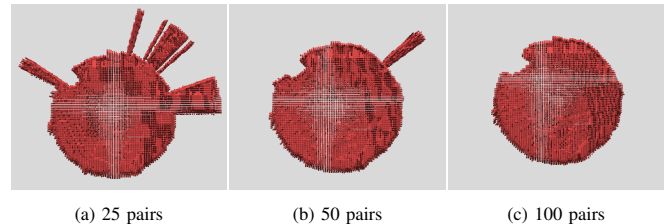


Fig. 13: Occupancy maps by our filter using Conditional GAN. Each model is trained by different pairs of raw data and hand-labeled masks.

VII. CONCLUSION AND FUTURE WORK

In this paper, we presented a novel application of cGANs to filter noisy sonar images. Compared to conventional filters, our approach can recognize and filter noise patterns better by distinguishing between obstacles and image artifacts. These results are attainable even with a small training set of raw data and hand-labeled mask pairs.

Using both simulated and real-world data, we showcase the applicability of our method to the downstream task of occupancy mapping, highlighting how our denoising method has a significantly better inference of free and occupied space compared to conventional methods.

For future directions of this work, we aim to find a solution to account for multipath reflections, something our method cannot disambiguate presently. Apart from autonomous planning and exploration, we aim to study the suitability of our method for feature-based SLAM and 3D object reconstruction using imaging sonar.

REFERENCES

- [1] M. Kaess, H. Johannsson, B. Englot, F. Hover, and J. J. Leonard, "Towards autonomous ship hull inspection using the Bluefin HAUV," in *Proc. 9th Intl. Sym. on Technology and the Mine Problem*, Monterey, CA, USA, May 2010.
- [2] M. F. Fallon, M. Kaess, H. Johannsson, and J. J. Leonard, "Efficient AUV navigation fusing acoustic ranging and side-scan sonar," in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Shanghai, China, May 2011, pp. 2398–2405.
- [3] E. Chen, S.-W. Huang, W.-H. Wang, and J.-H. Guo, "Side scan sonar grid map for unmanned underwater vehicle navigation," in *Proc. IEEE/MTS OCEANS Conf. and Exposition*, Waikoloa Village, HI, USA, Sep. 2011.
- [4] E. Hernandez, P. Ridao, A. Mallios, and M. Carreras, "Occupancy grid mapping in an underwater structured environment," in *Proc. 8th IFAC Conf. on Manoeuvring and Control of Marine Craft*, Guarujá, Sao Paulo, Brazil, Sep. 2009, pp. 286–291.
- [5] H. Cho, B. Kim, and S.-C. Yu, "AUV-based underwater 3-D point cloud generation using acoustic lens-based multibeam sonar," *IEEE J. of Oceanic Engineering*, vol. 43, no. 4, pp. 856–872, Oct. 2018.
- [6] G. Marani, S. K. Choi, and J. Yuh, "Underwater autonomous manipulation for intervention missions AUVs," *Ocean Engineering*, vol. 36, no. 1, pp. 15–23, Jan. 2009.
- [7] F. S. Hover, R. M. Eustice, A. Kim, B. J. Englot, H. Johannsson, M. Kaess, and J. J. Leonard, "Advanced perception, navigation and planning for autonomous in-water ship hull inspection," *Intl. J. of Robotics Research (IJRR)*, vol. 31, no. 12, pp. 1445–1464, Oct. 2012.
- [8] P. V. Teixeira, M. Kaess, F. S. Hover, and J. J. Leonard, "Underwater inspection using sonar-based volumetric submaps," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Daejeon, Korea, Oct. 2016, pp. 4288–4295.
- [9] M. VanMiddlesworth, "Toward autonomous underwater mapping in partially structured 3D environments," Master's thesis, Massachusetts Institute of Technology, Feb. 2014.
- [10] J. Jaybhay and R. Shastri, "A study of speckle noise reduction filters," *Signal & Image Processing: An Intl. J. (SIPIJ)*, vol. 6, no. 3, pp. 71–80, Jun. 2015.
- [11] S. Karabchevsky, D. Kahana, O. Ben-Harush, and H. Guterman, "FPGA-based adaptive speckle suppression filter for underwater imaging sonar," *IEEE J. of Oceanic Engineering*, vol. 36, no. 4, pp. 646–657, Aug. 2011.
- [12] J. Grabek and B. Cyganek, "Speckle noise filtering in side-scan sonar images based on the Tucker tensor decomposition," *Sensors*, vol. 19, no. 13, Jun. 2019.
- [13] N. D. A. Mascarenhas, "An Overview of Speckle Noise Filtering in SAR Images," in *Proc. Image Processing Techniques, First Latino-American Seminar on Radar Remote Sensing*, Buenos Aires, Argentina, Dec. 1996, pp. 71–79.
- [14] A. Lopes, R. Touzi, and E. Nezry, "Adaptive speckle filters and scene heterogeneity," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 28, no. 6, pp. 992–1000, Nov. 1990.
- [15] A. Isar, M. Sorin, and D. Isar, "A new denoising system for SONAR images," *EURASIP J. on Image and Video Processing*, vol. 2009, no. 1, Oct. 2009.
- [16] L. Sendur and I. W. Selesnick, "Bivariate shrinkage functions for wavelet-based denoising exploiting interscale dependency," *IEEE Trans. on Signal Processing*, vol. 50, no. 11, pp. 2744–2756, Nov. 2002.
- [17] S. Negahdaripour, P. Firoozfam, and P. Sabzmejdani, "On processing and registration of forward-scan acoustic video imagery," in *Proc. 2nd Canadian Conf. on Computer and Robot Vision*, Victoria, BC, Canada, May 2005, pp. 452–459.
- [18] Y.-S. Shin, Y. Lee, H.-T. Choi, and A. Kim, "Bundle adjustment from sonar images and SLAM application for seafloor mapping," in *Proc. IEEE/MTS OCEANS Conf. and Exposition*, Washington, D.C., USA, Oct. 2015.
- [19] E. Westman, A. Hinduja, and M. Kaess, "Feature-based SLAM for imaging sonar with under-constrained landmarks," in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Brisbane, Australia, May 2018, pp. 3629–3636.
- [20] E. Westman and M. Kaess, "Wide aperture imaging sonar reconstruction using generative models," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Macau, China, Nov. 2019, pp. 8067–8074.
- [21] E. E. Hundt and E. A. Trautenberg, "Digital processing of ultrasonic data by deconvolution," *IEEE Trans. on Sonics and Ultrasonics*, vol. 27, no. 5, pp. 249–252, Sep. 1980.
- [22] Y. Lu, R. W. Liu, F. Chen, and L. Xie, "Learning a deep convolutional network for speckle noise reduction in underwater sonar images," in *Proc. 11th Intl. Conf. on Machine Learning and Computing (ICMLC)*, Zhuhai, China, Feb. 2019, p. 445–450.
- [23] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, USA, Jul. 2017, pp. 5967–5976.
- [24] E.-h. Lee, B. Park, M.-H. Jeon, H. Jang, A. Kim, and S. Lee, "Data augmentation using image translation for underwater sonar image segmentation," *PLOS ONE*, vol. 17, no. 8, Aug. 2022.
- [25] E. Potokar, S. Ashford, M. Kaess, and J. G. Mangelson, "HoloOcean: An underwater robotics simulator," in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Philadelphia, PA, USA, May 2022, pp. 3040–3046.
- [26] E. Potokar, K. Lay, K. Norman, D. Benham, T. B. Neilsen, M. Kaess, and J. G. Mangelson, "HoloOcean: Realistic sonar simulation," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Kyoto, Japan, Oct. 2022, pp. 8450–8456.
- [27] General Dynamics Mission Systems, "Bluefin HAUV," <https://gdmissionsystems.com/products/underwater-vehicles/bluefin-hauv>.
- [28] Blueprint Subsea, "Blueprint subsea oculus M1200d," <https://blueprintsubsea.com/oculus/oculus-m-series>.
- [29] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 81, no. 3, pp. 425–455, Sep. 1994.
- [30] X.-W. Liu and K. Cheng, "Three-dimensional extension of Bresenham's algorithm and its application in straight-line interpolation," *J. of Engineering Manufacture*, vol. 216, no. 3, pp. 459–463, Mar. 2002.
- [31] P. Blondel, *The Handbook of Sidescan Sonar*. Springer Berlin Heidelberg, New York, NY, USA, May 2010.
- [32] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, Cambridge, MA, USA, Aug. 2005.