

GhostNetZero: AI for Detecting Marine Ghost Nets

Zhongqi Miao^{1+*}, Gabriele Dederer^{2+*}, Mareen Lee², Emrah Birsin³, Crayton Fenn⁴, Kai Krutzke⁵, Theodoros Stougiannis⁵, Eva Borges⁵, Lisa Stempel⁵, Rahul Dodhia¹, Juan Lavista Ferres¹⁺

1. AI for Good Lab, Microsoft, USA
2. WWF, Germany
3. Opitz Consulting Deutschland GmbH, Germany
4. Fenn Enterprises, USA
5. Accenture, Germany

* Represents equal contributions

+ Corresponding authors: Zhongqi Miao (zhongqimiao@microsoft.com), Gabriele Deberer (Gabriele.Dederer@wwf.de), Juan Lavista Ferres (jlavista@microsoft.com)

Abstract

Abandoned, lost, or otherwise discarded fishing gears (ALDFG), commonly referred to as ghost nets, pose a persistent global threat to marine biodiversity. Constructed from durable synthetic polymers, ghost nets remain intact for decades, continuing to entangle and kill marine organisms while damaging habitats and imposing economic burdens on fisheries and coastal communities. Despite their ecological significance, ghost nets are notoriously difficult to detect due to oceanic dispersal, submersion, and burial in sediment. Side-scan sonar has emerged as a powerful detection tool, but its high cost and limited spatial coverage constrain its large-scale application.

In this study, we evaluate the feasibility of applying modern computer vision and AI techniques to sonar-derived imagery for automated ghost net detection. In our experiments, we achieved an approximately 90% ghost net detection rate in data collected from the Baltic Sea and the Puget Sound regions. To operationalize this approach, we developed GhostNetZero, a human-in-the-loop web platform that integrates AI predictions with expert review, streamlining validation workflows and enabling iterative model refinement.

Our results highlight the promise of AI-assisted sonar analysis in scaling ghost net detection, complementing costly manual surveys and supporting targeted removal operations. By advancing automated detection methods, this study contributes to global efforts to mitigate the impacts of ghost gear and safeguard marine biodiversity.

1. Introduction and Background

1.1 Ghost Nets and Their Impacts on Marine Biodiversity

Ghost nets, formally referred to as abandoned, lost, or otherwise discarded fishing gear (ALDFG), are a pervasive and persistent source of marine debris. Once adrift, these nets do not merely entangle marine organisms—they become self-filling traps. Fish and other animals that die within them act as bait, attracting new prey and predators into the net, perpetuating a cycle of capture and mortality known as ghost fishing (Matsuoka et. al, 2005). This process undermines marine biodiversity, damages habitats, and imposes economic costs on fisheries and coastal communities. This phenomenon undermines marine biodiversity, damages habitats, and imposes economic costs on fisheries and coastal communities. Historically, fishing gear was manufactured from natural fibers that degraded rapidly, posing limited long-term risks. In contrast, contemporary gear is predominantly synthetic—constructed from durable polymers such as nylon, polyethylene, and polypropylene—that can remain intact in the marine environment for decades (Richardson et al., 2022).

Quantifying the scale of ALDFG is challenging, but recent global estimates suggest that approximately 2% of all fishing gear deployed annually is lost, equivalent to 80,000 square kilometers of nets, millions of traps, and billions of long line hooks entering the ocean each year (Richardson et al., 2019; Richardson et al., 2022). Expressing losses in such gear-specific units is increasingly recognized as a more reliable indicator of ecological risk than earlier mass-based figures.

The ecological consequences of ghost nets are well documented in various studies (Perroca et al., 2024), confirming that hundreds of species—including fish, invertebrates, seabirds, marine mammals, and sea turtles—become entangled or are killed by derelict fishing gear (Kühn and van Franeker, 2020). Entanglement can cause drowning, starvation, or severe injury, while gear resting on the seabed damages benthic communities such as corals and seagrasses. Although precise population-level impacts are difficult to quantify due to inconsistent reporting and variable survey effort (e.g., uneven monitoring across regions), the scientific consensus is that ghost nets represent a chronic, global threat to marine biodiversity (Global Ghost Gear Initiative, 2023).

1.2 Detecting and Scaling Ghost Net Removal

Despite their persistence and ecological significance, ghost nets are notoriously difficult to locate and quantify once lost at sea. Ocean currents and storms can transport nets far from their point of origin, dispersing them across wide areas or depositing them in remote, hard-to-access regions. Many nets sink below the surface, becoming entangled on the seabed or buried in sediment, rendering them invisible to surface observation. Others drift in the water column, moving unpredictably with tides and currents. These dynamics, combined with the vastness of the marine environment and limited survey capacity, mean that only a fraction of ghost nets are ever detected or recovered (Global Ghost Gear Initiative, 2023).

Side-scan sonar has proven to be one of the most effective tools for overcoming these challenges. By emitting sound waves that reflect off objects on the seafloor, side-scan sonar produces high-resolution images capable of revealing derelict fishing gear, even when nets are partially buried or obscured by benthic habitats. This technology enables systematic seabed mapping over large areas, improving both the efficiency and accuracy of ghost net detection (Global Ghost Gear Initiative, 2023). However, side-scan sonar operations are costly and require substantial logistics, including vessels, trained crew, and sonar experts. Furthermore, a single vessel with crew and equipment can typically survey only around 130 hectares per day. Compared to the size of the ocean (~36 billion hectares), the spatial coverage of such survey can be limited. To gain broader insights into ghost net distribution, it is necessary to rethink the methodology by incorporating existing datasets, which can be analyzed to identify likely hotspots and guide future removal efforts.

To address this need, in this paper, we explore the feasibility of modern AI methods to identify areas with a high likelihood of ghost net presence. Our initial models achieve detection rates of around 90% on datasets from the Baltic Sea and Puget Sound. We also introduce GhostNetZero, a human-in-the-loop web interface that integrates AI with human expertise to operationalize ghost net detection. Finally, we discuss potential improvements and future directions for this project.

2. Data Background

In this study, we have used two datasets. One is from the Baltic Sea region, Germany, and the other is from the Puget Sound region, US.

2.1 Baltic Sea:

The Baltic Sea dataset was collected under the WWF Germany Ghostnet Program, an initiative aimed at locating and documenting derelict fishing gear in the Baltic Sea. After a 2018 training workshop with sonar specialist Crayton Fenn, the WWF field team adopted side-scan sonar as a core survey method. From 2019 through 2023, data were acquired using a Marine Sonic Mark II system operating at a 600 kHz frequency.

Survey sites were identified in collaboration with local fishermen to focus on areas with a high probability of gear loss. Transects spanned near-coastal waters at depths of 5 m to 18 m, reflecting typical regional fishing grounds. From the cumulative sonar detections, 239 representative images were selected for model training and evaluation.

Since 2022, state-funded programs in the German federal states of Mecklenburg-Vorpommern and Schleswig-Holstein have broadened the scope of the work. These initiatives provide financial and logistical support for fishermen, sonar technicians, and retrieval teams, increasing the mapped inventory of ghost-net locations and supplying a growing stream of annotated sonar data to enhance future AI model training.

2.2 Puget Sound:

The Puget Sound data was collected by Fenn Enterprises in collaboration with the Northwest Straits Foundation and Natural Resources Consultants, Inc. Data collection was carried out in 2014 across selected areas of Puget Sound, Seattle, USA. Puget Sound is an extensive inland estuarine system encompassing roughly 1,000 square miles of marine waters and complex shoreline habitats.

The survey targeted nearshore to mid-depth zones, spanning from the intertidal shoreline to approximately 36 m (120 ft) water depth. Within these transects, sonar imagery was examined for indications of lost fishing gear. From the full set of detections, 173 discrete targets were ultimately selected to train the detection model used in this study. The dataset analyzed in this study was acquired using a Sea Scan PC side-scan sonar system (Marine Sonic) operating at a frequency of 600 kHz.

Puget Sound has a well-documented history of intensive commercial and recreational fishing. Government records over the past four decades provide detailed information on fishing activity, which informed the spatial design of the survey. The search area was strategically aligned with known salmon migration corridors and historic fishing grounds, reflecting regions of sustained fishing effort and higher likelihood of derelict gear accumulation.

3. Methods

3.1 Model Selection

To automatically detect ghost nets from sonar-derived imagery, we selected DeepLabV3 (Chen et al., 2017) with ResNet50 (He et. al, 2015) as the backbone computer vision model for semantic segmentation. This choice was driven by two key considerations. First, unlike image-based categorical classification (Russakovsky et. al, 2015) or bounding box-based object detection methods (Girshick 2015), semantic segmentation methods such as DeepLabV3 provide pixel-level classification, enabling more accurate localization of ghost nets. This is particularly important for irregularly shaped ghost nets, such as those in

string-like configurations, as it allows divers to accurately target retrieval efforts. Second, DeepLabV3 is a widely adopted and thoroughly studied semantic segmentation model, known for its straightforward architecture and ease of implementation, making it suitable for this application (Chen et al., 2017).

3.2 Data Preparation

3.2.1 Image conversion




Figure 1. Raw uncut sonar imagery is pre-cut into small image segments and then annotated into binary semantic segmentation masks for ghost nets

The sonar data were preprocessed to generate imagery suitable for computer vision analysis. Raw sonar signals, which represent acoustic reflections from the seafloor, were first converted into grayscale images. Then, these sonar images were separated into left and right views to account for the dual-channel nature of side-scan sonar data. Depending on the resolution of the generated images, they were split into patches of either 2000×500 or 1000×250 pixels. All patches were then resized to a uniform dimension of 2000×500 pixels to standardize input for the DeepLabV3 model. Both datasets used in this study—from the Baltic Sea and Puget Sound regions—were manually annotated and verified by marine science experts using the Azure Machine Learning Studio Labeling tool. As the task was to segment ghost net regions, annotations were converted into binary masks (see Figure 1). In total, the Baltic Sea dataset comprised 239 annotated image segments, while the Puget Sound dataset included 173 annotated segments.

3.2.2 Train, Validation, and Test Split

In this study, due to the focus on evaluating the feasibility of AI-based ghost net detection and the limited size of the datasets, the 239 Baltic Sea image segments were divided into 205 for training and 34 for validation. Similarly, the 173 Puget Sound image segments were split into 145 for training and 28 for validation. No separate test set was allocated. Hyperparameters, including learning rate, batch size, and optimizer settings, were optimized using the validation subsets to ensure robust model performance.

3.2.3 Data Augmentation

To enhance the generalizability of the DeepLabV3 model and mitigate overfitting, we applied random data augmentation techniques to the input imagery during training. Including transformations such as rotation, flipping, and intensity adjustments. Each augmentation was applied with a predefined probability to balance diversity and stability in the training process.

4. Results

In this study, we trained and evaluated three models to assess the performance of AI-based ghost net detection method: one was trained on Baltic Sea (BS) dataset, one was trained on Puget Sound (PS) dataset, and one combining both BS and PS datasets. Each model was evaluated on both the BS and PS validation sets to assess its generalizability across regions.

Model performance was measured using two metrics: mean Intersection over Union (mean IoU, Everingham et. al, 2010), a standard semantic segmentation evaluation metric, and centroid detection rate (CD) with buffer ranges of 3px, 5px, 10px, and 20px. The centroid detection rate, a metric commonly used in remote sensing and medical imaging (Stereńczak et. al, 2012, Xing et. al, 2016), considers a predicted polygon a true positive if its centroid falls within or near the target polygon. If multiple predictions are within the range of the same annotated polygons, they are considered one single true positive. This approach accommodates the potential inaccuracies in our annotations and the discontinuity of polygon predictions. The rationale for selecting this metric alongside mean IoU is detailed in Section 5.1.

The results are presented in Table 1. The BS-Only model, trained on Baltic Sea data, achieved a mean IoU of 0.740 for BS and 0.547 for PS, indicating strong performance on its training region but limited generalization to Puget Sound. Its centroid detection rates improved with larger buffers, ranging from 0.761 (3px, BS) to 0.891 (20px, BS) for BS, and remained constant at 0.607 across all buffer sizes for PS, suggesting poor adaptability to the PS dataset. The PS-Only model, trained on Puget Sound data, recorded a mean IoU of 0.611 for BS and 0.683 for PS, showing better performance on its training region but weaker transferability to BS. Its centroid detection rates ranged from 0.543 (3px, BS) to 0.674 (20px, BS) for BS, and from 0.889 (3px–10px, PS) to 0.940 (20px, PS) for PS, with the highest rate reflecting optimal performance on PS data with the largest buffer. These results also suggest that even though the mean IoU of PS-Only model is relatively lower, the majority of the predictions are very close to the annotated ghost nets.

The combined BS+PS model, trained on both datasets, exhibited the most robust performance, with a mean IoU of 0.739 for BS and 0.685 for PS, demonstrating enhanced generalizability. Its centroid detection rates increased with buffer size, ranging from 0.761 (3px, BS) to 0.891 (20px, BS) for BS, and from 0.821 (3px, PS) to 0.929 (20px, PS) for PS. Notably, the BS+PS model achieved centroid detection rates of 0.891–0.929 for the 20px buffer across both regions. These findings underscore the combined model’s effectiveness in handling regional variations and improving detection accuracy with larger spatial tolerances.

Table 1: Model performance on the Baltic Sea and Puget Sound datasets

Model	mean IOU		CD (3px)		CD (5px)		CD (10px)		CD (20px)	
	BS	PS	BS	PS	BS	PS	BS	PS	BS	PS
BS-Only	0.740	0.547	0.761	0.607	0.783	0.607	0.804	0.607	0.891	0.607
PS-Only	0.611	0.683	0.543	0.889	0.587	0.889	0.652	0.889	0.674	0.940
BS+PS	0.739	0.685	0.761	0.821	0.783	0.893	0.848	0.893	0.891	0.929

5. Discussion

5.1 Mean IoU v.s. centroid detection rate

As introduced in Section 3, the centroid detection rate considers a predicted polygon a true positive if its centroid falls within or near the target polygon. Compared to conventional semantic segmentation evaluations like mean IoU, the centroid detection rate offers a more relaxed evaluation criterion, which is particularly relevant for real-world applications of ghost net detection for several reasons. First, the quality of manual annotations can vary and is not always reliable due to factors such as human error or inconsistent sonar imagery, making precise overlap metrics less robust. Second, for practical purposes, identifying the relative locations of ghost nets is sufficient to guide divers to retrieval sites, reducing the need for exact boundary delineation.

This distinction is illustrated in the example in Figure 2, where the model successfully detects both annotated ghost nets in the image, providing actionable location information for divers. However, traditional mean IoU might yield a lower score due to imperfect boundary alignment, potentially underrepresenting the model's practical utility. In this scenario, using the centroid metric, the detection would be scored as two true positives and one false positive, reflecting a more operationally relevant assessment. This flexibility in the centroid approach ensures that the evaluation better aligns with the needs of practitioners, offering a more instructive metric for deployment.

Additionally, the centroid detection rate accounts for disconnected predictions that represent a single ghost net object. For example, in Figure 2 (b), the three predicted polygons in the top right corner collectively correspond to a single annotated ghost net and are counted as one true positive. This capability is particularly valuable for ghost nets with fragmented or irregular shapes. By aggregating such predictions, the centroid metric provides a more holistic evaluation of detection performance, complementing the limitations of mean IoU and enhancing the model's applicability in real-world efforts.



(a) Annotation mask



Figure 2. Example of a predicted ghost net mask and its corresponding annotation mask.

5.2 GhostNetZero – A human-in-the-loop ghost net detection platform

We have developed GhostNetZero.ai (in collaboration with Accenture), a human-in-the-loop platform designed to enhance ghost net detection using AI and expert oversight. This platform allows users to upload raw sonar files, which are automatically processed by the AI models trained in this study. The models generate preliminary predictions, identifying segments likely containing ghost nets. This automation streamlines the workflow by reducing the need for manual review of entire datasets.

Experts utilizing GhostNetZero.ai are only required to review and verify sonar segments flagged with ghost net predictions, rather than the entire sonar dataset, significantly reducing the time and effort involved in the validation process. This human-in-the-loop approach complements the AI models with expert oversight, enhancing the reliability of predictions. Feedback from expert reviews can also be leveraged to fine-tune and iteratively refine the AI models, improving their performance. As model performance increases, the need for human review will decrease, further boosting the efficiency of the ghost net detection process. The overall pipeline of this platform, from data upload to expert validation and model retraining, is illustrated in Figure 3.

5.3 Next steps

As mentioned in the previous sections, the dataset used in this study is relatively small, comprising only 239 annotated image segments from the Baltic Sea and 173 from Puget Sound. This limited scale can constrain the performance of AI models, particularly for a non-traditional task such as segmenting ghost nets. One of our primary next steps is to largely increase the number of annotations to enhance model accuracy and robustness.

In addition to expanding the absolute number of annotations within single regions, we aim to broaden the geographical scope of our dataset. Currently, our annotated data are confined to the Baltic Sea and Puget

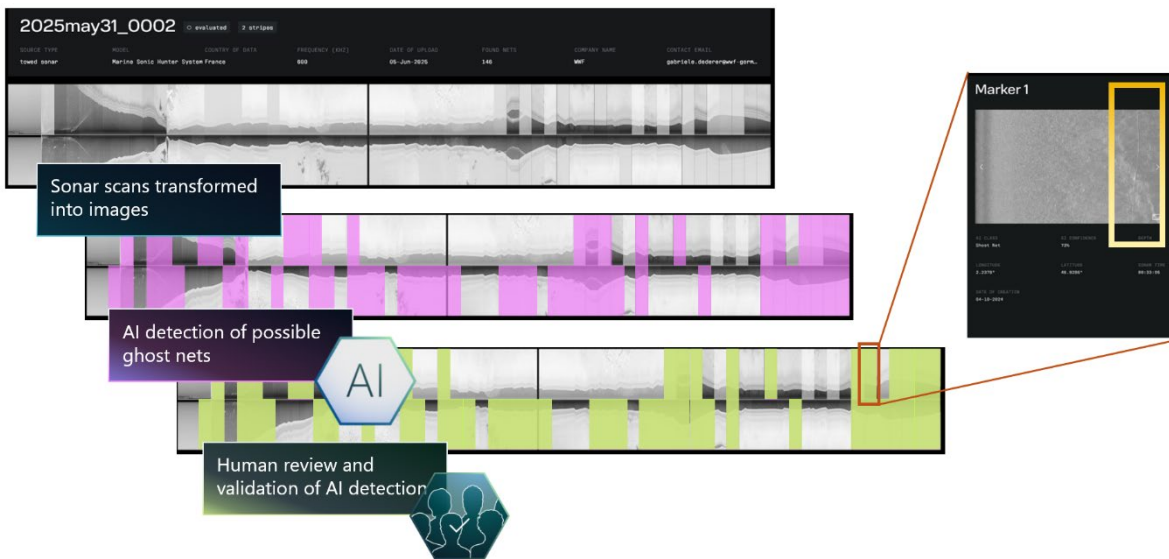


Figure 3. Pipeline of GhostNetZero.ai.

Sound areas. To extend the impact of our ghost net detection efforts to a global audience, we plan to incorporate annotated datasets from other regions worldwide. This expansion will facilitate the development of a versatile AI model capable of supporting diverse groups across different marine environments.

Beyond ghost nets, other types of lost fishing gear, such as crab pots, also pose threats to marine life and warrant attention. As part of our future work, we will expand the annotation categories to include these additional lost gears, enhancing the model's utility for groups with varied operational needs.

By addressing these next steps—scaling up annotations, diversifying geographical coverage, and broadening gear categories—this study lays the groundwork for a more comprehensive and globally applicable solution to mitigate the ecological impacts of marine debris. Through continued innovation and collaboration, we aim to advance the fight against ghost gear, safeguarding marine biodiversity for future generations.

References

- MATSUOKA, T., NAKASHIMA, T. and NAGASAWA, N. (2005), A review of ghost fishing: scientific approaches to evaluation and solutions. *Fisheries Science*, 71: 691-702. <https://doi.org/10.1111/j.1444-2906.2005.01019.x>
- Global Ghost Gear Initiative (2023) *End plastic pollution: Towards an international legally binding instrument*. Washington, DC: Ocean Conservancy. Available at: https://oceanconservancy.org/wp-content/uploads/2023/11/GGGIALDFGILBIWhitePaper_Web.pdf (Accessed: 8 September 2025).
- Kühn, S. and van Franeker, J.A. (2020) 'Quantitative overview of marine debris ingested by and entangling marine megafauna', *Marine Pollution Bulletin*, 151, p. 110858. doi:10.1016/j.marpolbul.2019.110858.
- Richardson, K., Hardesty, B.D. and Wilcox, C. (2019) 'Estimates of fishing gear loss rates at a global scale: A literature review and meta-analysis', *Fish and Fisheries*, 20(6), pp. 1218–1231. doi:10.1111/faf.12407.
- Richardson, K., Wilcox, C., Vince, J., Hardesty, B.D. and Campbell, M.L. (2022) 'A global assessment of the amount of fishing gear lost to the ocean each year', *Science Advances*, 8(7), eabj7912. doi:10.1126/sciadv.abj7912.
- Perroca, Júlia Fernandes, Tommaso Giarrizzo, Ernesto Azzurro, Jorge Luiz Rodrigues-Filho, Carolina V. Silva, Marlene S. Arcifa, and Valter M. Azevedo-Santos. "Negative effects of ghost nets on Mediterranean biodiversity." *Aquatic Ecology* 58, no. 1 (2024): 131-137.
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., & Yuille, A. L. (2017). DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4), 834–848
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Deep Residual Learning for Image Recognition. arXiv preprint arXiv:1512.03385.
- Russakovsky, Olga, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang et al. "Imagenet large scale visual recognition challenge." *International journal of computer vision* 115, no. 3 (2015): 211-252.
- Girshick, Ross. "Fast r-cnn." In *Proceedings of the IEEE international conference on computer vision*, pp. 1440-1448. 2015.
- Everingham, Mark, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. "The pascal visual object classes (voc) challenge." *International journal of computer vision* 88, no. 2 (2010): 303-338.
- Stereńczak, K., and S. Miścicki. "Crown delineation influence on standing volume calculations in protected area." *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 39 (2012): 441-445.
- Xing, Fuyong, and Lin Yang. "Robust nucleus/cell detection and segmentation in digital pathology and microscopy images: a comprehensive review." *IEEE reviews in biomedical engineering* 9 (2016): 234-263.