Lessons from Building a Semantic 3D Gaussian Splatting Pipeline for Multi-Year Coral Reef Monitoring

Alzayat Saleh*, Kang Han†, Renata Ferrari‡

*AIMS@JCU, James Cook University, Townsville, Australia

†La Trobe University, Melbourne, Australia

‡Australian Institute of Marine Science (AIMS), Townsville, Australia

***Abstract*—We report practical lessons from *SemanticReef*, an end-to-end pipeline that turns a routine single-camera underwater survey into a photorealistic 3D Gaussian Splatting (3DGS) reconstruction carrying *per-Gaussian* ecological labels. It is built on the multi-year AIMS *EcoRRAP* large-area-image datasets from the Great Barrier Reef (GBR) and Torres Strait. Rather than a benchmark, we share what worked, what failed in practice, and where the open problems lie. Three lessons stand out. First, classical structure-from-motion (SfM) with global bundle adjustment is still the most reliable camera-pose source underwater: the five feed-forward and neural SfM methods we tried degraded reconstruction or exceeded memory and disk budgets. Second, medium-capacity "underwater 3DGS" backbones do not yet capture reef-plot detail. Third, a training-free, alpha-weighted per-Gaussian label lift, paired with an adaptive water-removal filter and a domain fine-tune, converts off-the-shelf 2D segmentation into usable 3D ecological labels (hard-coral IoU ≈0.72). We close with the field-deployment problems that still gate reliable monitoring: metric-scale calibration, temporal co-registration, and domain shift in turbid water.**

I. Motivation

The Australian Institute of Marine Science (AIMS) *EcoRRAP* program (Ecological Intelligence for Reef Restoration and Adaptation) [1] has imaged 416 permanent reef plots across 104 sites and six regions in the Great Barrier Reef and Torres Strait over the past six years (2021–2026). Existing analytical pipelines (ReefCloud point cover, RapidBenthos [2], and AI-assisted manual segmentation in TagLab) work on 2D imagery, so they miss the vertical relief and overhangs of a reef that is really a 3D landscape (Fig. 1). 3DGS offers a photorealistic, real-time-renderable representation, but turning it into a *monitoring* tool requires per-Gaussian ecological labels, metric scale, and cross-year consistency. We summarise lessons from building this across hundreds of square metres of reef (∼2,385–4,092 images per 72 m^2 plot) within a multi-year monitoring program.

II. Lesson 1: Classical SfM is non-negotiable underwater

We experimented with replacing COLMAP [3] with faster feed-forward and neural SfM, and found two distinct failure modes. Two methods produced camera poses too inaccurate for splatfacto to lock onto: MERG3R (−8 dB PSNR vs. the COLMAP baseline) and GIM-DKM with GLOMAP (−9.23 dB). Two others exceeded our compute budget even on a 48 GB L40S HPC node, not just a consumer GPU: MASt3R-SfM on cached-pointmap memory and MapAnything on VRAM at ∼2,000 views (newer variants only partly close this gap). Separately, a learned reef-specific ego-motion SfM *reports* 7.84% mean relative pose error versus COLMAP's 1.85% on the same transect [4], about 1 m of drift on a 12 m plot, enough to break per-colony tracking. *Takeaway:* classical SfM with global bundle adjustment is still the most reliable camera-pose source underwater, and the faster alternatives we tried fail on pose accuracy or on a compute cost that bigger hardware does not lift, so we keep COLMAP.

Fig. 1. Two real-time oblique views of the same TSAU_FL1S_P4 reconstruction (photorealistic 3DGS, water column removed), freely navigable from any angle. Tabular, massive and branching colonies stand off the substrate, giving the vertical and overhanging 3D structure that 2D-projected benthic-cover metrics systematically under-sample. Producing a representation like this from a routine single-camera survey, then making it *measurable*, is the task this paper reports lessons from.

III. Lesson 2: Medium-model "underwater 3DGS" does not yet capture reef-plot detail

Underwater 3DGS is a fast-moving area (SeaSplat [5], WaterSplatting [6], R-Splatting [7], UW-GS [8], ReefMapGS [9]).

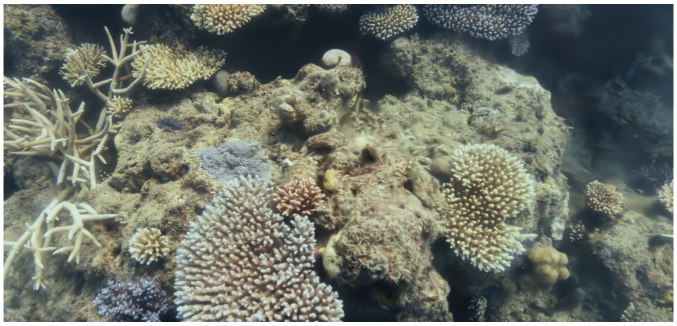

(a) photorealistic reconstruction

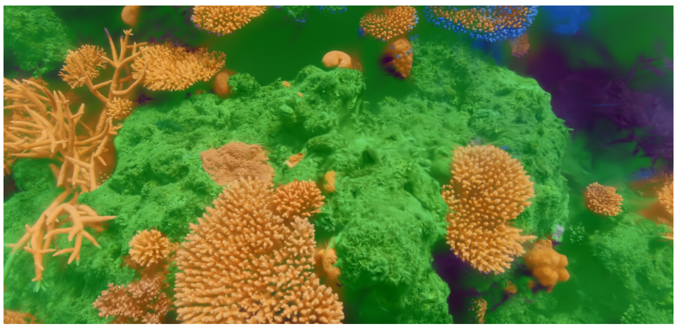

(b) per-Gaussian ecological class colour

Fig. 2. Per-Gaussian semantic labelling on the fully-pipelined ONLI_BA1S_P1 plot (DeepReefMap production head). (a) the photorealistic 3DGS reconstruction; (b) the same Gaussians 50/50-blended with their ecological class colour via the alpha-weighted lift (coral orange, algae green, water/background blue/violet), exposing class structure (coral morphotypes, algae, substrate) while preserving texture.

These optimise image clarity, but in our hands the medium-capacity decomposition models traded reef-structure fidelity for water modelling: a tuned WaterSplatting reconstruction looked visibly blurry despite a competitive 21.94 dB, while standard splatfacto with unsupervised enhancement (UDnet [10]) preserves fine coral texture and reaches 24–27 dB on our cleaner plots (indicative, not a matched-plot comparison). *Takeaway:* for monitoring, we keep a high-capacity general 3DGS backbone (splatfacto) plus a separate enhancement front-end, and treat underwater-specific backbones as a reconstruction-stage option to revisit, not a default; in our tests they mainly suppressed floaters rather than adding detail. None of these methods attaches ecological semantics or measurement, which is the layer we add next.

IV. Lesson 3: A training-free semantic lift turns 2D labels into 3D ecology

We assign ecological classes to every Gaussian by lifting a frozen 2D segmentation head across views; our production head is the 37-class DeepReefMap head of Sauder et al. [4], which we fine-tune to our domain (below). A naive multi-view majority vote leaks foreground labels onto occluded back-facing Gaussians (22.9% mislabelled as water on our gold-standard plot). An alpha-weighted lift adapted from FlashSplat [11], which accumulates each Gaussian's opacity-transmittance contribution to same-class pixels, cuts that to

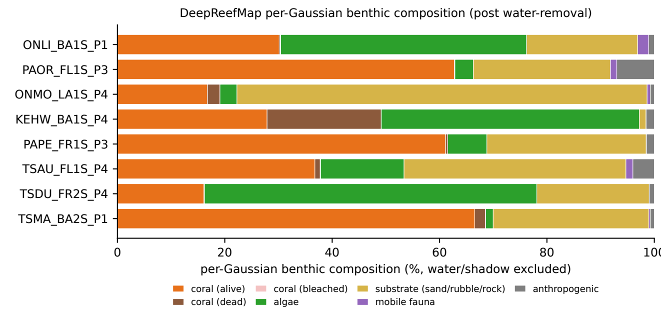

Fig. 3. Per-Gaussian benthic class composition across the eight pipelined plots (DeepReefMap head, alpha-weighted lift), after adaptive water removal and renormalisation over labelled benthic Gaussians. Coral, algae and substrate dominate every plot (e.g. PAOR_FL1S_P3 reaches ∼63% coral, consistent with its high RapidBenthos hard-coral cover), whereas the out-of-domain six-class SegFormer head over-predicts water to 85–90% here.

2.4% with no retraining of the splat (Fig. 2). Unlike identity-learning approaches (Gaussian Grouping [12], SAGA [13]), we keep a training-free, taxonomy-preserving lift suited to cover and volume measurement.

The dominant failure was *domain shift*: an out-of-domain head over-predicts "water" on shallow back-reef plots in turbid water (up to 89.9% on one plot). Two fixes helped. First, a per-plot *adaptive water-removal threshold* (tightening confidence as predicted water rises) preserves reef while removing true water. Second, and decisively, we *fine-tune* the head on RapidBenthos-derived GBR label tiles (3,597 tiles, 8 plots) under a leave-cluster-out protocol. This reaches a mean held-out hard-coral IoU of 0.72 over six folds (from 0.87 on the best plot down to 0.49 on the algae-dominated worst plot). For scale (though not directly comparable), RapidBenthos reports a segment-overlap IoU of 0.79 (0.87 on overlapping segments) [2], a class-agnostic metric distinct from our per-class semantic IoU. *Takeaway:* the training-free lift is how labels reach 3D, but domain robustness must be *baked into the 2D head* via fine-tuning, not assumed; the algae-dominated worst-case plot (0.49) is the honest floor.

V. Future challenges for field-ready reef mapping

Several challenges remain before these labelled reconstructions become a dependable monitoring instrument. **Metric scale.** We calibrate from physical fiducials (a 30 cm dumbbell), but automatic detection across plots is unsolved. **Temporal co-registration.** Comparing a colony across years needs each year's reconstruction aligned to the same fixed markers in full 3D pose; our 2-marker similarity transform is not yet accurate enough to trust per-colony change. **Semantically-resolved 3D metrics.** The labelling lets us measure true per-colony surface area and volume, but isolating individual colonies is still manual [14]; that, not PSNR, is the bar for operational use.

Acknowledgements

We thank the AIMS 3D modelling team (Program 1: Ecology and Monitoring) and the AIMS *EcoRRAP* (Ecological

Intelligence for Reef Restoration and Adaptation) program for the survey data, 3D reconstructions, and validation ground truth, and we acknowledge the RapidBenthos pipeline [2] as our benthic-cover reference. The Reef Restoration and Adaptation Program (RRAP) was funded by the partnership between the Australian Government's Reef Trust and the Great Barrier Reef Foundation, partners include: the Australian Institute of Marine Science, CSIRO, the Great Barrier Reef Foundation, Southern Cross University, the University of Queensland, Queensland University of Technology and James Cook University. Pipeline code and processed per-plot outputs are available from the authors on request; raw *EcoRRAP* imagery follows the AIMS data-sharing policy.

REFERENCES

[1] R. Ferrari, S. Gordon, M. Toor, S. Noonan, M. Lechene, M. Becker, A. Piccaluga, J. Harianto, E. Aston, A. Guillame, M. Drake, D. Burn, M. Alvarez-Noriega, M. Nielsen, E. Brinkner, W. Figueira, A. Hoey, D. Lennon, and K. Fabricius, "Reef restoration and adaptation program – integrated field-testing and sub-program management (ECO-01) final report," Australian Institute of Marine Science, Tech. Rep., 2025, 28 pp.

[2] T. Remmers, N. Boutros, M. Wyatt, S. Gordon, M. Toor, C. Roelfsema, K. Fabricius, A. Grech, M. Lechene, and R. Ferrari, "RapidBenthos: Automated segmentation and multi-view classification of coral reef communities from photogrammetric reconstruction," *Methods in Ecology and Evolution*, vol. 16, pp. 427–441, 2025.

[3] J. L. Schönberger and J.-M. Frahm, "Structure-from-motion revisited," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[4] J. Sauder, G. Banc-Prandi, A. Meibom, and D. Tuia, "Scalable semantic 3D mapping of coral reefs with deep learning," *Methods in Ecology and Evolution*, vol. 15, no. 5, pp. 916–934, 2024.

[5] D. Yang, J. J. Leonard, and Y. Girdhar, "Seasplat: Representing underwater scenes with 3d gaussian splatting and a physically grounded image formation model," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2025, arXiv:2409.17345.

[6] H. Li, W. Song, T. Xu, A. Elsig, and J. Kulhanek, "WaterSplatting: Fast underwater 3D scene reconstruction using gaussian splatting," in *International Conference on 3D Vision (3DV)*, 2025. [Online]. Available: https://arxiv.org/abs/2408.08206

[7] G. Huang, H. Wang, Z. Qi, W. Lu, D. Bull, and N. Anantrasirichai, "From restoration to reconstruction: Rethinking 3d gaussian splatting for underwater scenes," *arXiv preprint arXiv:2509.17789*, 2025.

[8] H. Wang, N. Anantrasirichai, F. Zhang, and D. Bull, "UW-GS: Distractor-aware 3D Gaussian splatting for enhanced underwater scene reconstruction," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2025, arXiv:2410.01517.

[9] D. Yang, J. Hong, J. J. Leonard, and Y. Girdhar, "Reefmapgs: Enabling large-scale underwater reconstruction by closing the loop between multimodal slam and gaussian splatting," 2026. [Online]. Available: https://arxiv.org/abs/2604.11992

[10] A. Saleh, M. Sheaves, D. Jerry, and M. R. Azghadi, "Adaptive deep learning framework for robust unsupervised underwater image enhancement (UDnet)," *Expert Systems with Applications*, vol. 268, p. 126314, 2025.

[11] Q. Shen, X. Yang, and X. Wang, "FlashSplat: 2D to 3D Gaussian splatting segmentation solved optimally," in *European Conference on Computer Vision (ECCV)*, 2024.

[12] M. Ye, M. Danelljan, F. Yu, and L. Ke, "Gaussian grouping: Segment and edit anything in 3d scenes," in *European Conference on Computer Vision (ECCV)*, 2024.

[13] J. Cen, J. Fang, C. Yang, L. Xie, X. Zhang, W. Shen, and Q. Tian, "Segment any 3d gaussians," 2025. [Online]. Available: https://arxiv.org/abs/2312.00860

[14] E. A. Aston, S. Duce, A. S. Hoey, and R. Ferrari, "A protocol for extracting structural metrics from 3d reconstructions of corals," *Frontiers in Marine Science*, vol. 9, p. 854395, 2022.