Accurate Underwater 3D Reconstruction of Coral Structures

An in-house photogrammetry pipeline for converting underwater video into measurable, textured three-dimensional coral models.

Coral reefs are among the most geometrically complex ecosystems found underwater, and that geometry, not just its appearance, is what shapes habitat and biodiversity. Diver observation and still photography, the conventional survey methods, record what a reef looks like but not its true three-dimensional form, which makes measuring change over time difficult. This paper sets out that problem, the environmental conditions that make it hard to solve underwater, and a video-based reconstruction method that addresses both.

The problem: geometry that conventional survey cannot capture

A coral formation cannot be adequately documented by a photograph or a diver's notes. Its value, structurally and ecologically, lies in its three-dimensional form: branching structures, cavities, surface irregularity, and fine morphological detail that a flat image cannot represent and a manual measurement cannot practically cover in full. Without a true three-dimensional record, two surveys of the same reef taken months apart cannot be directly compared. There is no way to measure growth, degradation, or physical damage with any precision, only a general visual impression that the reef looks different.

What is needed is a model, not a photograph: one that can be measured, viewed from any angle, and compared directly against a model of the same structure captured at a different time.

The gap: why underwater reconstruction is difficult

Reconstructing three-dimensional geometry from video is an established technique on land. Underwater, the same technique runs into conditions that have no equivalent in air, and each has to be actively managed rather than assumed away.

- Light attenuation and scattering reduce contrast and distort colour, making visual features harder to detect reliably.
- Suspended particles shift position between frames and can generate false matches between images that were never really looking at the same point.
- Motion blur, from vehicle movement or current, weakens both feature matching and camera position estimation.
- Insufficient overlap between frames leaves the software without enough common ground to reconstruct a connected model.
- Low-texture regions, such as open water or uniform surfaces, offer too little visual information to reconstruct at all.

A reconstruction method that ignores these conditions will fail silently, producing a distorted or incomplete model without any obvious warning that something has gone wrong. Any dependable system has to be built around these constraints from the start.

The solution: a video-based reconstruction pipeline

The system takes underwater video as input and produces a textured three-dimensional model as output, through a defined sequence of stages: frame extraction, image quality assessment, feature detection and matching, camera pose estimation, sparse reconstruction, depth estimation, dense reconstruction, surface generation, texture mapping, and accuracy validation, before the final model is exported in GLB format.

Two design choices in this pipeline directly address the conditions described above.

The first is adaptive frame extraction. Rather than processing every frame of a video at a fixed rate, the system adjusts extraction between three and ten frames per second based on scene complexity and camera motion, using a higher rate where overlap is limited or the camera is moving quickly, and a lower rate in feature-rich scenes where consecutive frames would otherwise be redundant. This keeps overlap sufficient for reconstruction without wasting processing on near-duplicate frames.

The second is active quality control. Frames affected by blur, poor focus, or heavy particulate interference are identified and excluded before they can degrade feature matching, and the reconstruction concentrates on regions of the scene with stable, identifiable features rather than forcing a result from regions, such as open water, that cannot support one.

Camera pose estimation and triangulated feature matching produce an initial sparse point cloud of the coral. Multi-view depth estimation then builds this into a dense reconstruction that preserves branching structures, cavities, and fine surface detail, which is converted into a continuous surface and textured using a mosaic of the original underwater imagery. The result is a model that carries both accurate geometry and the coral's true visual condition at the time of capture.

Accuracy is validated by comparing known physical distances on the coral or on reference objects in the scene against the equivalent measurements taken from the finished model, a standard and repeatable step rather than an optional check.

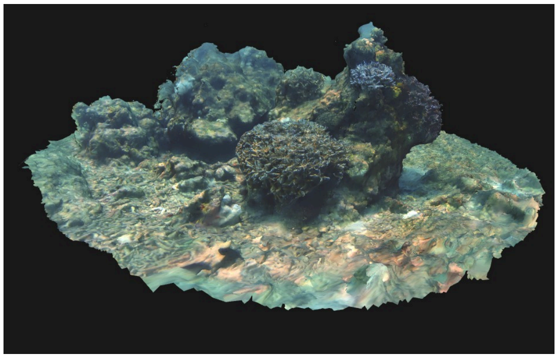
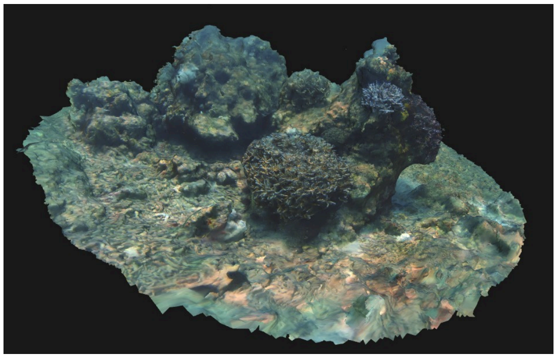

Output and application

The finished model exports in GLB format, a compact, portable format usable in web-based viewers, digital twin platforms, virtual reality applications, and marine monitoring systems. In practice, this turns reconstruction into a repeatable survey method: the same coral structure can be reconstructed on a schedule, with each model directly comparable to the last, enabling systematic monitoring of growth, degradation, or damage rather than a one-off documentation exercise. The same models support direct geometric analysis, including surface area, volume, and structural complexity, and the workflow can be paired with ROV and AUV platforms to generate three-dimensional models as part of a routine survey mission.

Conclusion

Coral geometry is difficult to document by conventional means and difficult to reconstruct underwater by any means that does not account for the environment's specific failure modes. The pipeline described here addresses both problems directly, using adaptive frame extraction and active quality control to work within the constraints of underwater video, and a defined multi-stage reconstruction process to convert that video into a measurable, textured three-dimensional model. The approach demonstrates that underwater video-based photogrammetry, applied through a purpose-built pipeline, is a practical and scalable method for turning underwater footage into a dependable long-term record of coral condition.

About EyeROV

EyeROV (IROV Technologies Private Limited) is headquartered at Maker Village, Kerala Technology Innovation Zone, Kochi, India, with a regional office in Doha, Qatar. The company designs, manufactures, and deploys Remotely Operated Vehicles, an Unmanned Surface Vessel, and the EVAP analytics platform, and conducts in-house research and development in underwater imaging, reconstruction, and analytics.

www.eyerov.com | sales@eyerov.com
