3D Recon Overview

Zupt, LLC

Introduction

Zupt, LLC has been using Inertial Navigation Systems (INS) either as standalone sensors or coupled with other aiding sensors in the Oil and Gas industry for over 15 years. Since 2015, Zupt has invested significant resources into integrating inertial technology with imaging sensors for various offshore survey tasks.

Over the last 6 years, we have developed several "turn-key" systems, tightly coupling Inertial Measurement Units (IMUs) with imaging sensors. These systems are intended to replace or augment "conventional" survey methods or, in the case of 3D Recon, generate a totally new deliverable for the integrity management sector. This unsolicited document specifically discusses Zupt's '3D Recon' and '3D Recon Mini' System:

3D Recon is a standalone 3D modeling sensor that incorporates three machine vision cameras and a very capable MEMS IMU within a single, compact 4000-meter rated subsea housing. 3D Recon delivers a real-time sparse point cloud that doubles as a relative-to-structure navigation system AND near real-time high-density 3D models (point cloud or mesh) of what is seen subsea. While initially developed for Integrity Management (IM) and Inspection, Repair, and Maintenance (IRM) purposes, 3D Recon has also developed significant interest for mooring chain link inspection, hull inspection, jumper/spool metrology, pipeline Out-Of-Straightness (OSS) surveys, pre-install/as-built, and many more applications.

3D Recon – How it works

3D Recon combines 3 high resolution, global shutter machine vision cameras with a very capable MEMS IMU in a single 4,000m rated subsea housing. Additionally, unique lights are deployed with 3D Recon to ensure the area to be mapped has the required light intensity.

3D Recon utilizes techniques derived from computer vision and inertial navigation to generate low resolution (sparse) and high resolution (dense) point clouds. The process begins with a short alignment period, where acceleration and angular velocity from 3D Recon's inertial measurement unit are used to determine heading pitch and roll. Next, the three cameras gather images from the scene. The intensity variance are used to find features; these features occur in areas where there is a large gradient (change in pixel intensity from pixel to pixel) in the horizontal and vertical directions. The regions around the detected features are then used to compute descriptor vectors for individual features. These descriptor vectors are then used to match features between the three cameras. Once the features are matched between cameras, known camera intrinsics (focal length, distortion, and optical center) and extrinsics (camera location and orientation) are used to triangulate the features, generating a sparse point cloud in the left camera frame. The 2D and 3D points, along with feature descriptors are then passed into the navigation engine.

Once 3D recon has been aligned, the navigation engine integrates incoming acceleration and angular velocity at 200 Hz to update the navigation states (position, velocity, attitude, and deterministic sensor errors) and their corresponding matrix of uncertainty, or covariance. When camera measurements are available, 3D Recon performs global matching where new features are matched to existing features in the map. The search is limited to previously seen features that should exist in the field of view. After limiting the search region, features are then matched using their descriptors. If their descriptors match within a set tolerance, the current position and attitude is used to project the features from the camera frame to the global frame. This involves applying rotational and linear lever arms to the 3D point set to put it in the IMU frame of reference followed by a second transformation that uses the current position and orientation to project the points into the global frame. Once the descriptor has matched, the global position of the current globally matched candidate is compared. After all features have been matched, they go through a further test to assure that the current pose with respect to the global map is in agreement with all features.

This test is completed by using random sampling consensus (RANSAC) and an solving a nonlinear optimization problem at minimizing the 2D reprojection error. RANSAC is used to determine an optimal model and point set in the presence of a lot of outliers. It is 3D Recon's last step before accepting a matched feature to use for navigation. If the feature has been matched, the navigation filter performs an update, which updates the entire state and covariance. If the feature has not been matched, it is projected into the map and added as a new feature. In parallel to the main navigation engine, 3D Recon also collects data to optimize the full path and map. This involves logging uncertainty, states, and measurement at important locations along the route. These measurements and states are known as keyframes, and are triggered based on distance traveled, number of matches, or loop closure events. The optimization thread is computationally intensive and runs at a much lower rate than the main navigation engine. After running the optimization, an updated map and pose is fed back into the main navigation engine. The dense point clouds are generated using the position and orientation computed from the optimization. For generating the dense point clouds, a match is attempted for every pixel based on intensity rather than feature descriptors.

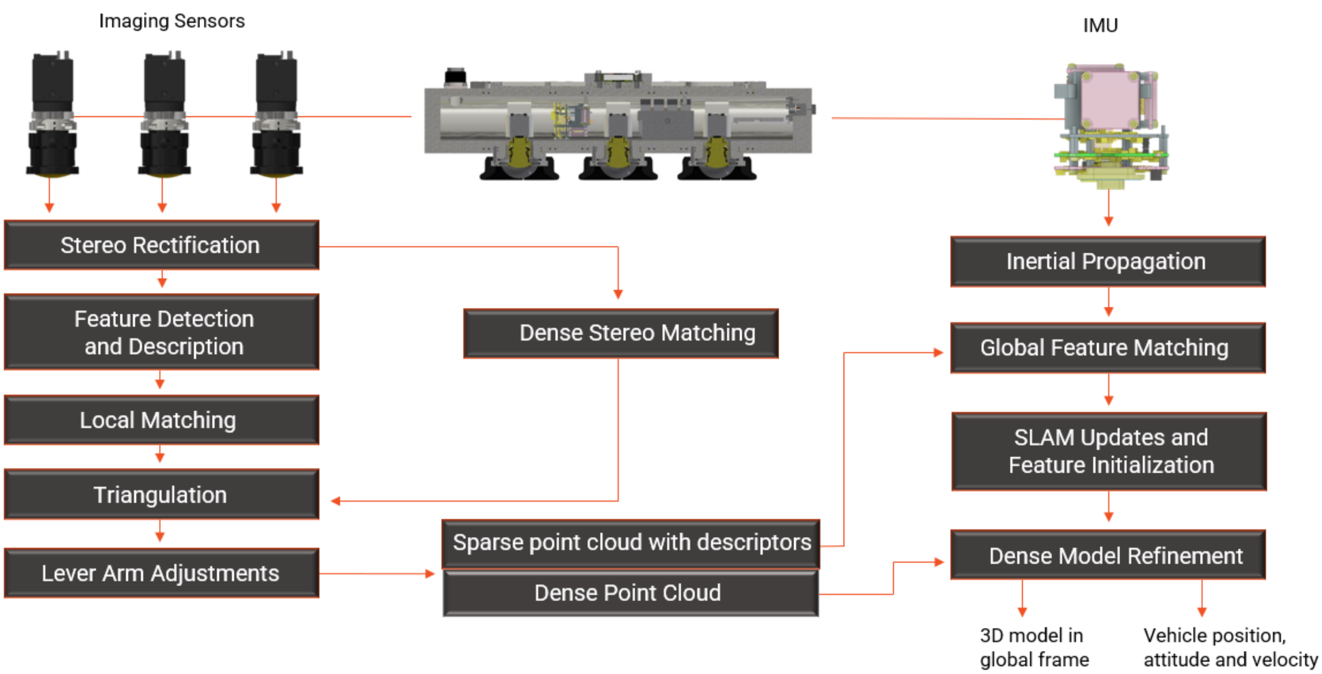

reliable point cloud generation. 3D Recon captures three images from the three cameras within the instrument at a frame rate of 5Hz. Feature descriptors are classically used for image-to-image feature matching within the same epoch to allow for the point cloud generation. In this case feature descriptors from camera to camera (across track) as well as from frame to frame to frame (along track) will be used to ensure any movement of the target structure in the Earth reference frame does not impact the accuracy of the model (point cloud) being generated. The same will be true for the SLAM solution, descriptors will be used to ensure that the model remains correctly scaled even in the presence of motion as data is being acquired.

Data Acquisition and ROV Requirements

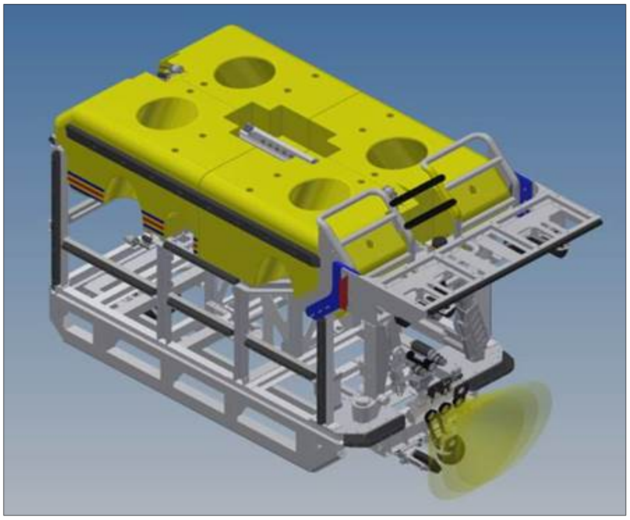

3D Recon has a diagonal field of view of apx. 85° x 65° and as such at a 2m 3D Recon has a "swath" width of approximately 3.75 meters. Multiple passes can be combined into the single point cloud as described above. 3D Recon requires a standard survey mux (RTS Gen5 or Innova Mux) to be installed onto the work class ROV and 3D Recon uses port 12b GB Ethernet to the surface. A detailed interface document may be provided delivered to outline ROV interfacing based on a project specific equipment suite. While 3D Recon can navigate as a stand-alone solution (navigating relative to the model that is being built), it is also possible to take in external aiding information to improve the positioning solution and navigate in a global reference frame. Without aiding information, 3D Recon navigates and builds a model with respect to a given starting coordinate. When aiding is included, the inputs are used to update the inertial navigation states in the form of a Kalman filter update.

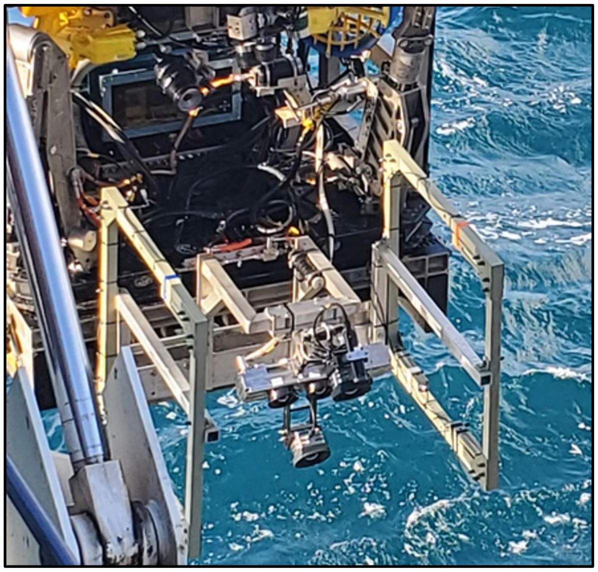

A single RS232 is required for the management and configuration of 3D Recon. A GB Ethernet port is required for the image data. All Power, serial communications and the imaging data is connected via a single PBOF cable to the ROV (RTS Mux or Innova Mux). The GB Ethernet connection is available as copper. Power requirements are 24Vdc 150W.

Through various testing environments, it has found that it is crucial to have well balanced illumination over the area being imaged, for this reason it is important that 3D Recon be mounted/held at a location that will allow the provided lights to illuminate the area without blockages or shadows from other ROV mounted equipment.

The surface work (office) space required for data acquisition and data processing will be the size of one or two conventional desk. A single desk workspace being apx. 1.5m or 60” wide by 0.75m or 30” deep. One desk space will be needed for the 3D Recon data acquisition hardware and processing computer (large mutli-GPU processor).

The 3D recon workspace should be no more than 6m or 20ft from the output of the Mux topside as we utilize 230400bps RS232 comms through the mux and the maximum (copper) cable length to sustain this comms speed is this distance. Either from the mux or the standard ROV survey port at the surface can be up to 30m or 100ft.

Linear Distance Accuracy

The linear scaling accuracy of 3D Recon is ~1:1000. The depth (of field) accuracy of the model built by 3D Recon while flying 2m above a target is sub pixel. When flying closer to the target the accuracy will be better than this. A pixel at 2m altitude is ~0.9mm. As long as the ROV can reliably fly along the area of a structure in question, an offset of 2m or less the model will be built with this level (sub pixel) of accuracy. If applicable, Zupt may be able to use magnetically attached checker boards as well as some unique targets to ensure the data can easily and quickly be quality controlled during both data acquisition and data processing. Standard point cloud quality metrics (SD or 1 sigma) PSNR methods or generalized Hausdorff distance methods will be used to define the quality of the delivered point cloud.

Deliverable and Examples

The “deliverable” from 3D Recon is a high resolution, spatially correct model (point cloud) as well as the position of the ROV with respect to the model. The images below are the color palletized point cloud images of millions of points.

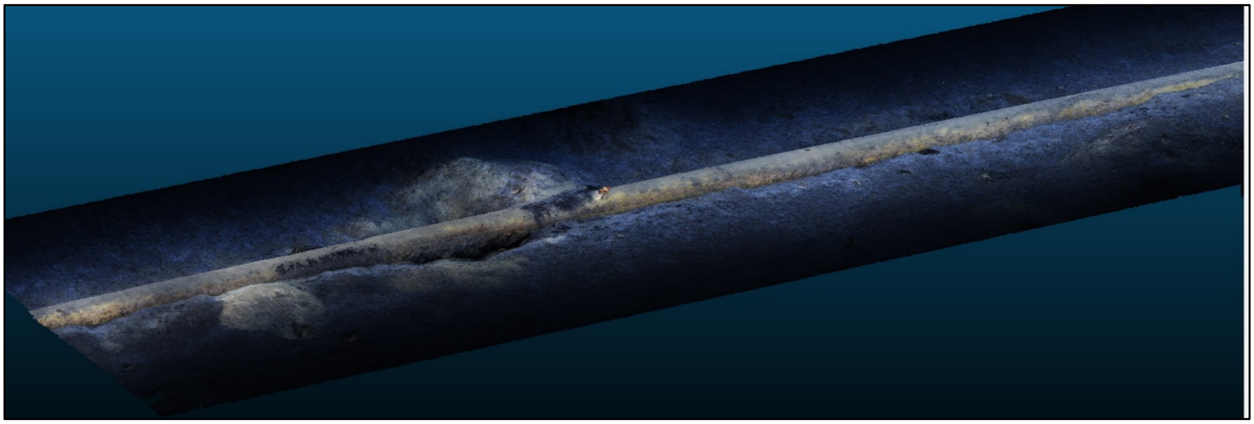

What differentiates 3D Recon from similar subsea imaging systems is the tight integration of a high accuracy MEMS IMU contained within the same subsea housing as the three-camera baseline. The integration of inertial technology provides the sensor with relative depth to the features it uses to build the model. This allows for continuous data collection through short periods where relative navigation from images alone will fail (kicked up clouds of mud, lack of features within the image, large fish shoals, etc.)

Due to the relative accuracy of the positioning solution available from the sparse point cloud 3D Recon has applications where precise positioning/modeling requirements exist for such applications as precise CP surveys of subsea assets or ultrasonic thickness inspections of deep SPAR hulls.

3D Recon Mini Concept – Possibly of Interest

3D Recon Mini a small, custom derivative (in concept form) of our 4,000m rated 3D Recon (a high-resolution 3D modeling system that integrated 3 machine vision cameras and a MEMS IMU within a single subsea housing) to be used to provide high resolution 3D models of the imaged chain mooring links.

Zupt's 3D Recon solution is a standalone subsea modeling tool that includes three 2048 x 2448 imaging sensors, a capable MEMS IMU, and PCBs within a single, 4000-meter rated subsea housing. 3D Recon differs from "conventional" photogrammetry solutions due to the tight coupling of an Inertial Measurement Unit into the imaging solution. The IMU integration provides a real time relative-to-structure navigation solution, eliminates the need to mobilize an external INS (ROVINS, SPRINT, etc.), and leads to faster processing of dense point cloud deliverables as well as more tolerance of particulate matter within the water column as data is being collected.

While 3D Recon's current configuration is physically too large and too heavy for certain scopes involving observation class ROVs, the software and processing solution from 3D Recon can be utilized in a variety of derivatives, such as the custom solution that will be referred to as "3D Recon Mini".

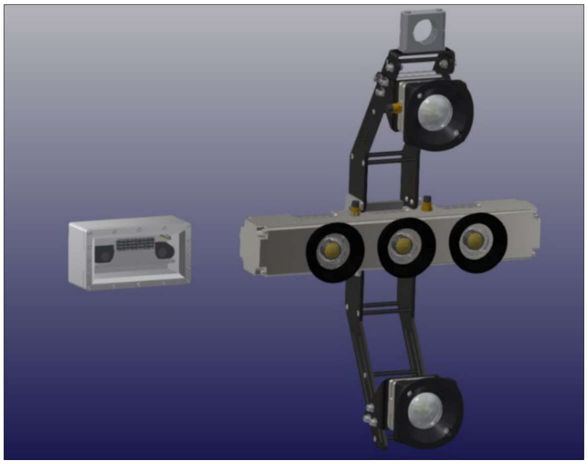

3D Recon Mini (left) size compared to 3D Recon (right)

3D Recon Mini maintains the same core components as 3D Recon but is a much more compact, lighter form (Dimensions: 220mm x 120 mm x 95mm, Weight in Air: ~4Kg, weight in water 0.5Kg), specifically designed to produce higher density 3D models at a smaller structure-to-sensor offset than 3D Recon. The sensor is specifically designed to deliver models with a 3D pixel size less than 0.5 millimeters. The same software and processing techniques will be used with 3D Recon Mini as typically used with 3D Recon.

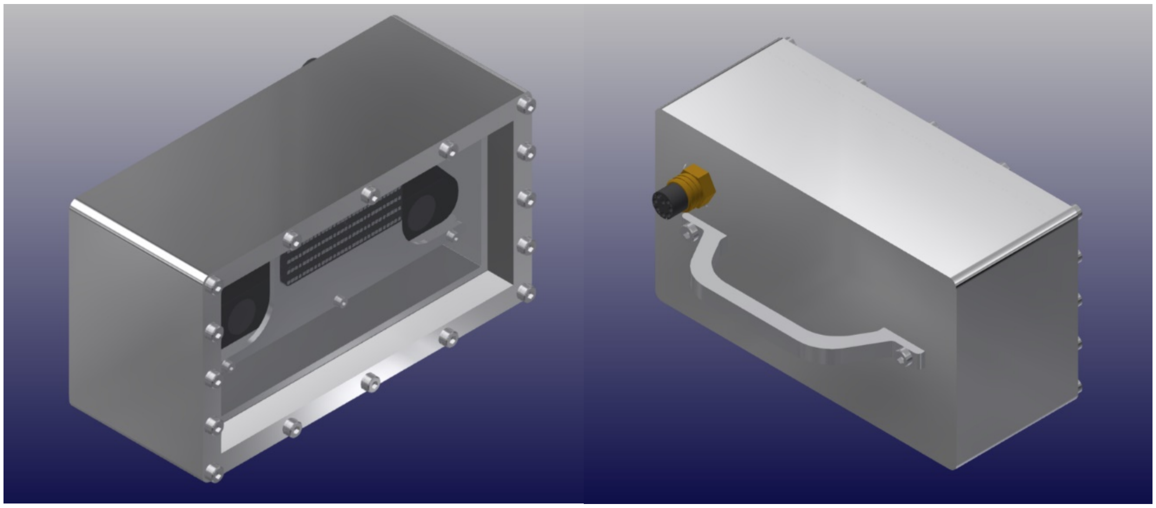

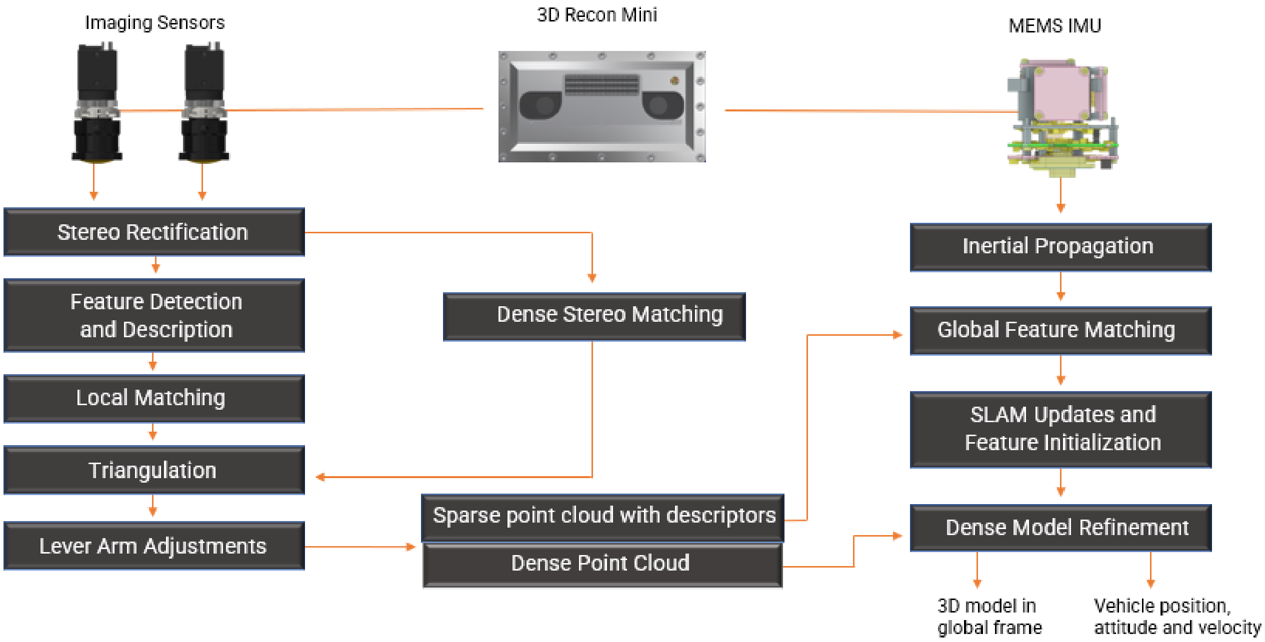

The custom solution includes 2 high-definition cameras (~2200 x 1200-pixel resolution for the initial delivery, followed later by ~2,600 x 2,600 resolution if needed) separated on a 125 millimeters baseline. The solution maintains the inclusion of a MEMS IMU fixed within the camera baseline to provide a precise pose estimate between image data acquisition and reduce the possibility of shifts/vibrations between the camera baseline and IMU. 3D Recon Mini is contained within an anodized aluminum subsea housing rated for 300-meter water depths. Zupt has included a simple bolted clamp to attach 3D Recon Mini to the ROV frame. We can also supply a handle if the ROV has a basic manipulator to allow the vehicle to carry the system.

Due to size limitations, Zupt would be processing the imagery and inertial data topside. For this reason, 3D Recon Mini will require a single fiber channel from the ROV to provide power (5 V, <2 A) and communications to the surface. 3D Recon Mini's subsea housing will include a bulkhead fiber connector that will allow for direct interfacing to the ROV or via a Zupt provided interfacing whip.

In water depths of 30 meters (or less) there will be significant light from the surface. Because this natural lighting may lead to undesirable shadows and uneven scene illumination, 3D Recon Mini includes some low power LED lights within the subsea housing for proper, diffused illumination of the scene being imaged.