

Real-Time Dynamic Pose Estimation Systems in Space: Lessons Learned for System Design and Performance Evaluation

Chad ENGLISH, Galina OKOUNEVA, Pierre SAINT-CYR, Aradhana CHOUDHURI, Tim LUU

Abstract- Real-time pose estimation systems are fundamentally difficult to evaluate. Pose estimation performance is a function of the object geometry, environmental conditions, line of sight, sensor characteristics, algorithm characteristics, and computational platform. This paper presents lessons learned from Neptec's two operational space systems for real-time dynamic pose estimation: the Space Vision System used by NASA for assembling the International Space Station, and the TriDAR Autonomous Rendezvous & Docking System used for spacecraft docking to the International Space Station and automated satellite servicing. Lessons learned are discussed in the context of design, trade-offs, challenges, and performance evaluation. This paper also discusses applying the lessons towards adapting these technologies for other applications.

These lessons result in a proposed generalized framework for performance evaluation of pose estimation systems. Such a task is not without challenges given the reliance of pose estimation on so many factors external to the system. Shape-based, targetless systems are particularly difficult to evaluate given the potential range of object shapes and views. However, predictable performance is critical for the success of space operations and NASA's strict requirements have provided many valuable lessons towards performance predictability. Collaborative research is presented on metrics for quantifying shape with respect to pose estimation, providing a potential means for evaluating and predicting system performance with respect to shape properties.

Index Terms—Pose estimation, ICP, lidar, performance evaluation, smart sensing, Principal Component Analysis

1. INTRODUCTION

Dynamic operations in space that involve imaging, modeling, or assembling generally require real-time 6 degree-of-freedom (DOF) pose estimation to track the coordinate transformations between the sensor and objects or between multiple objects seen by the sensor. In machine vision, pose estimation techniques generally come in two flavours: those that use fiducial targets placed on the object to be tracked and those that use natural features of the object. Fiducial targets are more invasive but traditionally have provided more robust and accurate measurements than feature-based techniques.

Over the last 20 years Neptec has developed and deployed a variety of machine vision systems for pose estimation. The Space Vision System (SVS) is a camera-based system that uses fiducial targets for assembling the

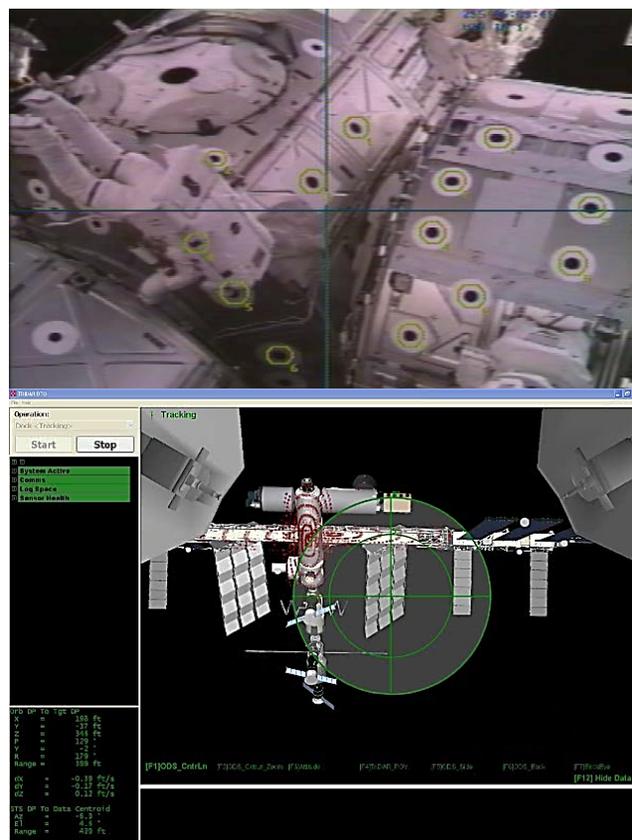


Figure 1: Displays from target-based Space Vision System (top) and Shape-based TriDAR AR&D system (bottom). The TriDAR display is shown here using 3D models updated by real-time pose estimation from a spiral TriDAR scan (red points on space station). Section 2.3 describes scan patterns.

International Space Station (ISS) and has flown on over 20 shuttle missions [26] (Figure 1, top). The TriDAR is a hybrid triangulation and pulsed time-of-flight (TOF) scanning laser-based object tracking system for autonomous rendezvous and docking (AR&D) ([35],[36],[37]) (Figure 1, bottom). Spin-offs have applied these machine vision technologies to space, defence, and industrial applications, including dynamic 3D imaging, 3D mapping, inspection, automatic target recognition, metrology, helicopter landing, and rover navigation.

This paper provides an overview of lessons learned in the development and evaluation of SVS and TriDAR real-time 6 DOF pose estimation systems. The overview is not intended to place these systems in the context of the entire pose estimation field, nor as a review of the latest image processing techniques. It is intended to provide a historical case study of challenges and trade-offs in the development

Manuscript received December 10, 2010; revised April 22, 2011. This paper is extended from "Smart Sensing for Real-Time Pose Estimation, Assembly, and Inspection Using 3D Laser Scanning Systems" published at the NIST *Performance Metrics for Intelligent Systems*, Baltimore, MD, USA, Sep., 2010.

C. English and T. Luu are with Neptec Design Group, 302 Legget Dr., Kanata, Ontario, Canada, K2K 1Y5, (e-mail: cenglish@neptec.com and tlou@neptec.com); P. Saint-Cyr, A. Choudhuri, and G. Okouneva and are with the Dept of Aerospace Engineering, Ryerson University, 350 Victoria St., Toronto, ON, Canada, M5B 2K3 (e-mail: psaintcy@ryerson.ca, aradhana.choudhuri@asc-csa.gc.ca, and gokounev@ryerson.ca).

of operational real-time pose estimation systems in a difficult environment where performance is critical. This paper illustrates how solutions to complex performance evaluation were developed, and how these lessons might apply to more general pose estimation evaluation.

The overview is divided into a series of steps. Section 2 describes the history of pose estimation in space leading up to available and emerging state-of-the-art sensing technologies for this type of application. Section 3 introduces the smart sensing approach used to achieve real-time 3D pose estimation, including functional parameters that affect and complicate pose estimation performance evaluation. This section also introduces spin-off applications of real-time pose estimation with varying goals that further complicate evaluation. Section 4 presents challenges in defining and evaluating metrics for pose estimation including logistical problems. Section 5 describes the verification and validation (V&V) process for SVS and TriDAR to demonstrate that the real-time pose estimation systems would work correctly in orbit, given the extreme environment and criticality of success. The V&V process divides the pose estimation into core components and tests them separately to make the evaluation tenable. Section 6 sets up a framework of shape-based metrics for separating intrinsic object characteristics from the system performance. These shape metrics are used to define an artifact optimized for pose estimation that has potential as an evaluation standard to compare pose estimation systems. Finally, Section 7 summarizes the work and provides conclusions for performance evaluation of pose estimation systems.

2. REAL-TIME POSE ESTIMATION IN SPACE

2.1 Evolution of Rendezvous & Docking Systems

For high-end, robust, mid- to long-range applications based on real-time 6 DOF pose estimation there are essentially two options: *target-based* systems or *shape-based*. Target-based systems use fiducial targets placed on the tracked object that are tracked using a camera, stereo vision, or laser tracking system. Shape-based systems generate 3D images of the tracked object and align it with a model of the object without adding targets to the object.

In its very basic form, pose estimation provides position (3 DOF) or position and orientation (6 DOF). Additional variables such as rates can be derived or measured. For space applications, the most basic and demanding use of pose estimation is autonomous rendezvous & docking (AR&D) of spacecraft. These craft can move relatively unconstrained in all 6 DOF, potentially at high relative rates under harsh conditions of lighting and temperature variation, and with very high risk associated with failure.

The state-of-the-art in AR&D technology has only matured in the last decade. The DARPA Orbital Express mission in 2007, sponsored by NASA and Boeing, saw the first successful AR&D system flown on-orbit [20]. This success was soon followed by the European Space

Agency's Jules Verne Automated Transfer Vehicle (ATV) successful mission of 2008 [16]. Prior to these successes, rendezvous, docking, and berthing systems required hands-on control by astronauts. Such systems date back to the optical cross-hairs of the 1966 Gemini 8 test docking with the unmanned Agena Target Vehicle up to target-based Space Vision System (SVS) for assembling the ISS. The Russian Kurs system uses radio telemetry from an array of transmitting and receiving antennae on both the chase and target vehicle to determine range, bearing, and attitude. The Kurs docking system has been in operation since 1986 aboard Soyuz TM-2, and its precursor, Iгла, since 1967. Despite this long heritage and proven reliability, Kurs is largely precluded from being used in modern AR&D systems because of its prohibitive total mass of 85 kg for the chase vehicle and 80 kg for the target vehicle, and power consumption of 270 W for the chase vehicle and 250 W for the target vehicle. It also relies on vacuum tubes with questionable lifetime.

It is only recently that fully automated systems have been attempted. The NASA Demonstration of Autonomous Rendezvous Technology (DART) mission in April 2005 was the first attempt to perform an automated rendezvous, but with no capture or docking, without a human in the loop [38]. The mission failed when DART prematurely placed itself in a retirement phase after missing some cues and colliding with the target vehicle. The automated rendezvous system used on DART was NASA's Advanced Video Guidance Sensor (AVGS). AVGS was later successfully used on the Orbital Express mission in 2007. This system uses cooperative retroreflective targets on the target vehicle and applies photogrammetric algorithms to the imaged target locations in a camera, much like SVS. AVGS improves upon SVS by using two active flood lasers to image the retro targets and filters out other light frequencies, making it more immune to varying lighting conditions. AVGS shows great promise but suffers from limitations of all target-based systems: the requirement for placement of targets on the vehicle and limited views over which pose can be measured.

The aforementioned Orbital Express mission used four sensor systems to accommodate AR&D operations: a laser rangefinder for target range information, a visible light sensor for daytime pointing of the rangefinder, an infrared (IR) sensor for nighttime pointing, and the AVGS system for close proximity final approach and capture [25].

The Jules Verne ATV mission also used a series of sensors and instruments to successfully rendezvous and dock with the ISS using GPS, a videometer that bounced pulsed laser beams off of passive retroreflectors, and image pattern analysis [15]. Redundant back-up lidar systems used pulsed TOF lasers on the retroreflectors to confirm range, but not orientation. The ATV system is limited to close range operations, similar to AVGS, with the additional constraint of GPS for longer range operations.

Stereo-vision has been used for tele-operation during the ETS-VII mission [31]. This system successfully demonstrated rendezvous and docking but was operated by

ground operators. Similar systems were proposed for use on the TECSAS and ConeXpress (now Smart-OLEV) missions. Stereo cameras are feasible as AR&D pose estimation sensors, but are susceptible to the same problems as other camera-based systems, in particular to varying lighting conditions, lack of features and low contrast, and errors in camera calibration.

There have been several successful laser-based systems for AR&D. The RELAVIS system developed by Optech uses 3D lidar data to provide pose estimation of target vehicles [2][21]. RELAVIS is a shape-based lidar pose estimation system similar to Neptec's TriDAR. The main differences between these systems involve the hardware design, the fitting algorithms, and the subsequent accuracy and speed performance. The RELAVIS system was successfully tested on the U.S. Air Force XSS-11 mission in 2005 to provide range and bearing information. Neptec's TriDAR system successfully flew on STS-128 and STS-131 for both docking and undocking of the Space Shuttle to the ISS and has been selected as baseline for several future satellite servicing and docking missions.

2.3 Autosynchronous Scanning

Single-point flying spot scanners allow precise control and measurement. The traditional disadvantages of such scanners include the mass, volume, and power required for moving parts and comparatively slow data acquisition rate.

The data rate can be mitigated by smart sensing techniques as described in Section 3. Speed and precision can also be optimized using an *autosynchronous scanning* design. First developed by the Canadian National Research Council (NRC), Neptec licensed and adapted the autosynchronous design as the Laser Camera System (LCS) for tracking SVS targets and imaging in space [42]. After the Columbia shuttle accident in 2003, the LCS was adapted for imaging and inspecting of the Space Shuttle's Thermal Protection System (TPS) on orbit due to its low measurement noise [11]. The autosynchronous design allows for a very narrow instantaneous field-of-view of the detector that can follow the laser spot by using the two sides of the same mirror. The low noise at standoff distances of meters makes the autosynchronous design ideal for metrology applications, and Neptec has since adapted it for high-precision manufacturing applications in the Laser Metrology System (LMS) [51].

The same autosynchronous design was expanded to the aforementioned TriDAR AR&D sensor for long range pose estimation. The TriDAR multiplexes a TOF pulsed lidar in the same optical path as the triangulation-based continuous wave (CW) laser of the LCS, thereby optimizing the complementary nature of triangulation and TOF lidar [14]. Pose estimation is accomplished using real-time Iterative Closest Point (ICP) software originally developed for automatic target recognition [13][34].

The use of a TOF lidar in the autosynchronous triangulation design also greatly increases the dynamic range of the lidar by having the laser spot "walk off" the

detector at closer ranges. Dynamic range defines the ability to maintain sufficient signal to detect distant, dark objects without causing the detector to saturate on close, bright objects. The dynamic range also allows the lidar to better scan through obscurants and is the basis for the Obscurant Penetrating Autosynchronous Lidar (OPAL) [52]. OPAL is essentially a TriDAR with the CW triangulation laser removed and with other modifications to enhance the obscurant penetration. It is currently used in a forward kinematic fashion using GPS and inertial information to geo-reference OPAL data to aid in landing helicopters in brownout or whiteout conditions [53]. If necessary, OPAL can still perform the same pose estimation as the TriDAR.

2.4 Flash Technology

Flash lidar typically refers to simultaneously acquiring all pixels of a 3D range image using TOF. For the last decade, flash imaging has been perceived by many as the holy grail of 3D sensing because of its potential size and data acquisition rate. There are many ways to accomplish flash 3D imaging: Continuous wave signal modulation, avalanche photodiode (APD) arrays, and range gating. However, they all share the similar concept of flood illumination and detector array. A flash lidar acquires all of the data at once without moving parts; in doing so, it requires many times more instantaneous power than flying spot scanners because it divides the power across the many detectors of the pixel array. This division of power limits combinations of range, field-of-view (FOV), and resolution. Some flash lidars can zoom the FOV, but resolution is in a fixed relationship to FOV.

Conversely, scanning sensors can often adjust scan size, resolution, and pattern which gives them an almost infinite zoom capability (limited by laser spot size) over a larger field of view and much further range for a given laser power. By using smart sensing techniques (Section 3), a scanning lidar can be directed to gather only needed data, thereby minimizing scan time and making more efficient use of the available laser power. The smart sensing capability comes at the price of moving parts and increased power consumption compared to flash lidars. The dynamic range of flash lidars tends to be much lower than flying spot scanners, particularly compared to autosynchronous geometry, because all of the range pixels in a flash lidar are made simultaneously from a single pulse or signal.

It is now clear that flash imaging will not completely replace scanning lidars, but rather will trade scan volume and dynamic range for data rate, making them better suited for different applications such as 3D video and tracking faster objects over a smaller volume.

3. SMART SENSING

Real-time dynamic pose estimation opens the door to a variety of applications, each with different performance requirements. This section presents the technical approach Neptec uses for real-time pose estimation and introduces a

variety of applications that have been developed as a result, ranging from proof-of-concept to operational systems. These applications are introduced in the context of their different uses of the real-time pose estimation that lead to different requirements for pose estimation performance. These differences become important in the context of standardized pose estimation performance evaluations discussed in subsequent sections.

3.1 The MILD Approach

Development of the TriDAR AR&D system has provided key lessons learned in building real-time systems. A key lesson is maximizing whole system efficiency in the collection and processing of data to extract the information of interest. Efficiency is achieved through an approach to 3D sensing that Neptec refers to as *smart sensing*. This approach applies a paradigm of “more information, less data” (MILD) with two key elements. First, intelligent algorithms process the raw 3D data directly at the sensor head. Second, the output of the processing algorithms drives the 3D scanner to gather only the data necessary to compute the next set of information. The system hence quickly converges on the collection of minimal datasets providing maximal information.

The TriDAR AR&D system represents a basic smart sensing system. The autosynchronous scanning design allows highly flexible scan patterns. Figure 2 shows three typical scan patterns for the TriDAR with representative laser spot spacing at uniform time intervals. The Lissajous pattern (left) has high point density in the corners and low in the center making it useful for monitoring transitions at the edges. The Rosette pattern (center) has highest point density at the center and lowest at the periphery, much like the human eye nerves. This makes it ideal for high resolution imaging of central features while maintaining width and height in less important regions for alignment stability and monitoring purposes. The Spiral pattern (right) is uniform in angular and radial spacing and hence ideal for surface digitization, inspection, or for uniform distribution of features.

Uniformly spaced raster scan patterns (not shown) are also typical for relatively static imaging such as LCS inspection of the Space Shuttles. Arbitrary patterns can also be used, limited only by the inertia of the scanning mirrors and galvo current. The three scan patterns in Figure 2 make optimal use of the natural sinusoidal inertias to cover an area in the fastest time. A single pattern is referred to as a frame of data. Frame rates depend on the chosen pattern, pattern parameters (Lissajous nodes, Rosette petals, Spiral turns), number of points, and size. Full FOV frame rates of 5-10 Hz are typically achievable.

Scan patterns are key to smart sensing real-time pose estimation. The patterns dynamically adapt to the pose estimate. The pose estimate and initialization algorithms are designed for unorganized point clouds and are not reliant on the scan pattern structure or a regular grid as with correlation operator techniques. Pose estimation is

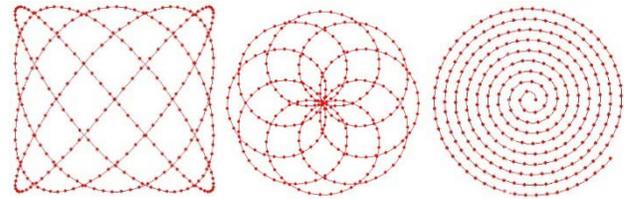


Figure 2: Dual-axes scanning patterns: 4x5 Lissajous (left), 8-petal rosette (center), and a 9-turn spiral (right).

initialized using a Polygonal Aspect Hashing technique [36]. The hashing algorithm creates N-point polygons from the unorganized scan points and compares them to a reference model hash table. The cumulative matched polygons provide a limited set of pose candidates that are tested by selecting the one that best fits the scan to the database model. The key to optimum efficiency and performance is a sparse reference database, storing only the dominant, discriminating geometric features. This provides a sufficiently accurate initial pose estimate while keeping the number of matches manageable.

The initial pose estimate is then used in a proprietary ICP algorithm to finely align the scan data to the model. The 6 DOF tracking uses the output pose estimate from one frame as the initializing pose estimate for the next frame using ICP [35]. The pose estimate is also used to adjust the size of the scan pattern to only gather data on the object or region as well as the position of the scan pattern within the sensor’s FOV. The size and position of the object is re-projected into the sensor FOV to define the new scan boundaries. Adaptively scanning in this manner collects only data necessary for the pose estimate. Adaptability is extended by changing the scan pattern and region as the target vehicle gets closer and fills the FOV. Scan patterns are optimized via pre-flight testing and pre-programmed for specific ranges. Future work includes pattern selection automation from real-time metrics.

Smart sensing mitigates the slower data rate and improves efficiency of scanning sensors. Flash-based sensors can gather large datasets very quickly but do so in a brute force manner where the laser power is distributed over the whole FOV, much of which does not fall on the target region of interest. With the autosynchronous scanner, all laser power can go towards generating useful, and ultimately used, data points. Furthermore, computational power is often a limiting factor and smart sensing can generate enough useful points per frame to max out the processing power while providing accurate pose estimates. The increased data rate of other systems means it often has to be sub-sampled to accommodate computational limitations for real-time processing. Combining these factors with the added data precision, autosynchronous scanning and the resulting real-time pose estimate can be of equal or better quality over a wider range of operations given the same computing hardware.

Real-time 6 DOF pose estimation is an application itself for a class of operations that includes assembly, positioning, and docking. Target-based systems, like SVS and AVGS, can be highly accurate and reliable under

controlled conditions and range of orientations where targets can be used. Targetless tracking systems such as TriDAR and RELAVIS extend the capability and reliability to less controlled conditions, such as space and subsea [27], but are limited to shapes that have definitive geometric features [35]. The following sections discuss additional smart sensing applications that use real-time, dynamic 6 DOF pose estimation as an intermediate step towards extracting additional information. Such systems put different requirements on the pose estimation

3.2 Dynamic Metrology

Metrology describes the measurement of parts. Metrology traditionally involves static measurements, often through full digitization of sampled parts via a Coordinate Measurement Machine (CMM). Dynamic pose estimation opens up the potential for more robust and efficient measurement in situ for a class of metrology applications. As a proof-of-concept, Neptec developed a demonstration of measuring car door gap and flush while in motion on an assembly line, tracking the door to determine the location of the gap.

A lesson learned in this process is that the pose estimate need not be precise. It only needs to be sufficient to aim the scanner near the gap. The gap and flush measurements themselves need to be precise in real time, and this can be accomplished via scanner precision and the ability to make many statistical measurements dynamically.

3.3 Object Recognition

A superset of real-time pose estimation is first recognizing the object to be tracked. Neptec's pose estimation system was first derived from real-time 3D Automatic Target Recognition (ATR) work [13][34]. A key feature discovered in early ATR work is that salient features of 3D shapes can be sufficiently generated by sparse, but smart, scanning. Rapid scan patterns can acquire data and recognize the object in less than a second. Another lesson is that continual tracking and scanning can vastly improve the recognition confidence [7]. As with the metrology, the pose estimate need not be precise when the extracted information is the identity of the object, unless pose is also important for additional operations.

3.4 Inspection and Digitizing

Smart sensing can also be applied when the data itself is the application's product. One class of applications that fits this circumstance is the digitization and inspections of parts or surfaces. Typically, this is performed under highly controlled conditions in which 3D data is gathered from multiple views of a part and stitched together to make a full 3D model. The controlled conditions are necessary to know the coordinate transformations between the views. The model is then assembled using the forward kinematics of known sensor and part positions.

A less controlled approach uses unknown sensor positions but tracks fiducial targets to calculate relative motion between views. An even less controlled alternative that makes use of the 3D shape of the part to align multiple views. These approaches use inverse-kinematic solution in which the relative sensor positions are inferred from the multiple dataset alignments [6][39]. A lesson learned from LCS imaging of the Space Shuttle TPS is that the accuracy of the sensor pose estimate is essentially irrelevant. What is important is that the data alignment is accurate [11]. Inverse kinematic solutions will generally provide inaccurate pose estimates unless the sensor is near the centroid of the scan data as in 360° scanning. Section 4.4 discusses this detail specifically.

The MILD paradigm can play a particularly important role in 3D dynamic imaging. Real-time pose estimation comes from aligning the same features in multiple scans, hence redundancy is inherently required. Redundancy adds unnecessary data to the 3D model. When aligned, excess points can be eliminated or, more preferably, improve the model by extracting better surface measurements. More information is encoded using less data. The efficiency from such a framework allows the model to be built with linear complexity and real-time surface reconstruction [47].

3.5 Dynamic mapping, navigation, and Geographic Information Services (GIS)

A second class of applications where the 3D data itself is the product is 3D terrain mapping. In principle, mapping is identical to object digitization, just with a very big object. In practice there are many differences. There is much less redundancy, use of multiple data sources (airborne and street-level lidar), and the need for synchronized high-end GPS and IMU navigation equipment to provide a forward-kinematic solution with accuracies comparable to the sensor accuracy. Small angle errors can contribute significant errors to the geo-referenced location of a point. Compounding this problem, ground vehicles will often lose some or all GPS signal in dense urban environments, relying solely on the inertial systems which inherently suffer from drift errors.

The inverse-kinematic solution can help to close the kinematic chain and improve the data by providing redundancy in both the sensor pose and the data pose. The redundant estimates (forward using GPS and IMU, inverse using data alignment) are not equal in weight. Generally speaking, the local pose estimate will be much better than the one acting through a moment arm. The pose estimates can use a weighted combination, such as in an Extended Kalman Filter (EKF). Such a system provides improved robustness to keep the EKF from drifting when navigation data is unavailable for short durations. Using tracking pose as the source data, this principle was successfully tested by the Canadian Space Agency (CSA) [1].

The approach is also the basis for the Optical IMU prototype used with Terrapoint's truck-based TITAN lidar collection system [19]. The inverse pose estimate was

maintained by using redundancy from the multiple TITAN sensors and hence was still prone to drift error, but it is independent of the inertial drift error. Maintaining accurate sensor pose is not irrelevant, as in object digitization, because the data must be accurately geo-referenced.

Closing the kinematic loop is also the basis for Simultaneous Localization and Mapping (SLAM). In the case of SLAM, absolute knowledge of location is absent. Only the inertia systems and scan data are available for generating both the map and locating the sensor (and hence the vehicle). SLAM is typically used for robotic rover navigation in unknown environments such as lunar and Martian exploration. The similarity between the requirements for the GIS mapping and rover SLAM means that similar real-time pose estimation techniques can be applied. Closing the kinematic loop for rovers in this manner has been investigated with Carleton University [33], Carnegie Mellon University [50], and with CSA for possible implementation on the Juno rover [23].

A lesson in closing the loop has been the difficulty in defining performance. It can be described independently for forward and inverse solutions, or combined. Consistency of the data is also a measure, as with object digitization, but ultimately absolute accuracy of the geo-referenced data is the goal, complicating the evaluation.

4. PERFORMANCE EVALUATION CHALLENGES

Perhaps the largest challenge for smart sensing is the quantification of pose estimation performance in any meaningful way. It seems simple in principle: measure the pose of objects and compare to truth data. Neptec has done this with SVS and TriDAR for space mission qualifications.

The problem with the approach is that pose estimation is highly situational. It is reliant on the quality of sensor data, processing algorithms, viewpoint, environmental conditions, and, perhaps most critically, on the object shape itself. (In a target-based system the shape is the arrangement of targets.) A pose estimation system might perform well when tracking one object in one setting, but perform poorly with a different object or different setting.

Furthermore, the lessons learned in the previous section showed that the same system can be differently optimized for different applications with dramatically different requirements. One cannot generically evaluate such a system without first identifying both the specific conditions and application it will be optimized for.

The following sections describe some of the components of the pose estimation chain and how they complicate performance measurement.

4.1 Sensor Data

3D sensor performance metrics typically describe the accuracy, precision, resolution, FOV, and raw data rates. Better performance in some of these metrics does not necessarily translate into improved pose estimation performance. A faster data rate does not provide improved

pose estimation if most of the points do not fall on the target or if the computation system is saturated.

Summary sensor metrics based on statistical properties do not always indicate key sensor behaviours that can affect pose estimation such as edge artifacts, intensity artifacts (black/white edges), dynamic range, saturation behavior, thermal and vibration stability, and other characteristics that can distort or affect the data.

The pattern of collection is also of critical importance. A standard scanning lidar may collect 50,000 points per second but does so in a fixed raster scanning pattern over the FOV that takes seconds to scan. Object and sensor motion during scanning can grossly distort the data. The pose estimate from such a scan is ambiguous at best since it doesn't represent a single fixed relative pose.

Conversely, an asynchronous scanning lidar with the same quality specifications can scan in programmable patterns, producing much lower distortion distributed evenly across the pattern. All of the scanner data may be on target and useable if set up in a smart sensing control loop, and hence the "used data rate" might be higher and less distorted than other sensors with higher raw data rates. Sensor specifications are insufficient for pose evaluation.

Flash lidars collect data simultaneously across all of the pixels, including close and far objects. While this can produce much unused data, a bigger concern may be the dynamic range. Simultaneous range measurements within a single scene cannot be too large, particularly if dark and shiny objects are in the scene, or else signal return will either be below some threshold or saturate the detector. Flash lidars typically achieve dynamic range on the order of a few hundred to one. Conversely, TriDAR is a scanning lidar that makes independent measurements with a bi-static optical design that can result in an effective dynamic range on the order of millions to one.

Edge effects can also be of critical importance. TOF lidars tend to extend edges outward as the edge of the beam hits the surface even after the centerline has moved off. Triangulation laser systems tend to distort the peak shape and shift the edge toward or away from the sensor, depending on the edge orientation [51]. Phased lidar systems have similar edge effects that produce distorted edges that are not random noise. Edge distortions affect the fitting of the scan data and can degraded pose estimates when the edges are important to the shape.

The above factors make the data, and resulting pose estimate, highly dependent on specific operational circumstances, complicating generic evaluation.

4.2 Object Geometry

Pose performance is highly dependent on object geometry, particularly weak geometry and degrees of freedom. The typical example for targetless pose is a sphere or a cylinder. Even with objects that have a defined pose, generalized quantification of performance is difficult. A coffee mug has a well-defined orientation in 6 DOF, but if the sensor cannot see the handle it has the same problem

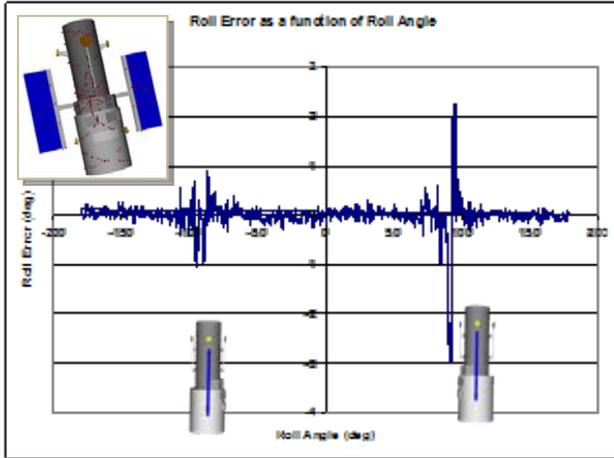


Figure 3: Hubble telescope (model inset) roll error as a function of roll angle, showing large error when viewing solar arrays edge on at ± 90 degrees.

as a cylinder. The accuracy is highly dependent on whether the sensor can see the handle or not. This is the real problem Neptec had in defining performance for the Hubble Telescope, shown in Figure 3 [35]. There is no single accuracy value that can describe this performance.

The difficulty is compounded by operational circumstances that affect whether or not key features are measured or missed, including scan patterns and zoom. Generic pose estimate metrics are meaningless when accuracy depends on the object, view, range, and many scanning parameters.

4.3 Algorithm Design

It may seem trivial to state that the pose estimation algorithm affects performance. The problem is that the definitions of “better” and “worse” are often highly dependent on the specific operation. Tuning parameters may trade one feature for another, such as robustness for speed or accuracy.

A common algorithm for pose estimation is the Iterative Closest Point (ICP), developed simultaneously by Besl and MacKay [4] and Chen and Medioni [9]. Many variants have been proposed and tested to improve ICP for specific optimizations such as speed or robustness. These improvements act on various components of the ICP algorithm itself: point matching, point rejection, weighting, minimization, convergence, and iteration criteria. As with parameter tuning, variations in the algorithms can also optimize for certain behaviors over others, complicating the evaluation process. When evaluating performance, the behavior that was optimized needs to be considered.

4.4 Data Location Relative to Reference Frame

Perhaps one of the most overlooked components of pose accuracy is the location of the reference point for which the pose is estimated. ICP and similar pose estimation

algorithms align the centroid of the collected (and used) target points with the centroid of the matched points on the reference object, with orientation computed to minimize a distance error between matched points. The centroid of the collected data is likely not at the reference point on the object, and the centroid changes as the scan data changes.

If only changes in object pose are of interest, the reference point may not be of great importance. However, if the reference point is meaningful to the operation, such as a docking port, the distance between this point and the collected data is problematic because it provides a moment arm to magnify orientation error.

Let P be the $3 \times n$ set of n points on the object in the sensor coordinate frame with no measurement error. Let Q be the same set of points in the object reference frame defined at the point of interest (such as a docking port). Let \bar{P} and \bar{Q} be centroids of the two datasets. The true pose of the object is defined by the 3×1 translation vector T and 3×3 rotation matrix R such that

$$P = RQ + TN \quad (1)$$

Here N is a $1 \times n$ vector of ones to complete the format. Likewise, the centroids map through the same pose,

$$\bar{P} = R\bar{Q} + T \quad (2)$$

Define E as a $3 \times n$ set of sensor measurement errors for each coordinate of each point in P . The sensor measurements of these points then become

$$P' = P + E \quad (3)$$

and

$$\bar{P}' = \bar{P} + \bar{E} \quad (4)$$

Note that E has not been assumed to follow any model and so can include noise, artifacts, and point matching errors.

The pose estimate is denoted (T', R') and solves the best fit between P' and Q . The rotation estimate, R' , is typically found via pseudo inverse techniques. The translation estimate is then found via the data centroids by

$$T' = \bar{P}' - R'\bar{Q} \quad (5)$$

Define the pose rotation error to be

$$R_{err} = R[R']^{-1} = RR'^T \quad (6)$$

or

$$R' = R_{err}^T R \quad (7)$$

This definition sets the rotation error as the rotation matrix that moves the estimated orientation to the true orientation. Using (2), (5) and (7), the pose translation error is

$$T_{err} = T' - T = [\bar{P}' - R'\bar{Q}] - [\bar{P} - R\bar{Q}]$$

$$= (\bar{P}' - \bar{P}) - (R' - R)\bar{Q}$$

$$\text{or } T_{err} = \bar{E} - (R_{err}^T - I_3)R\bar{Q} \quad (8)$$

Here I_3 is the 3 x 3 identity matrix. Hence the pose translation error is a combination of the mean error of point measurement, \bar{E} , and the rotation error acting through moment arm \bar{Q} .

It is important to reiterate what \bar{Q} is. It is the centroid of the physical points on the object that are measured by the sensor, given in the object coordinate system located at the reference point on the object for which the position is being estimated. A large \bar{Q} means the centroid of the collected data is far from the reference point. Conversely, the position error of the reference point can be minimized if it is at the centroid of the sensor data.

This feature of pose estimation is important in at least three ways. First, it is exactly the problem of inverse kinematics outlined in Section 3.5. For navigation, it is far more accurate to have the sensor gather data all around it such that the sensor is at or near the data centroid.

Second, it means that even for a given object, environment, pose system, and view, the pose accuracy cannot be generically described. It matters where the chosen reference point for tracking is located. For absolute pose estimation applications, such as docking or grasping the object, the reference point is generally not a variable. It is the point of interest for the operation.

Third, this moment-arm relationship inherently implies a trade-off when tracking an object. Gathering data surrounding the object reference point generally reduces positional error, but this must be traded off against the strength of the object shape for alignment. If the reference point is in the middle of a flat plane or on a smooth, spherical surface, perhaps a nearby detailed feature could provide a better pose estimate even with the added moment arm. Conversely, the strongest shape features might be far from the reference point and hence rejected for closer, though locally weaker, feature scans.

The dependence of pose accuracy on reference point complicates the concept of generic performance evaluation methods. Specific applications with fixed operational parameters simplify the process. The next section looks at case studies of pose evaluation where the operations are limited but not fixed, and the resulting complexity.

4.5 Verification and Validation Process

The best solution to the evaluation problem may be through synthetic simulation. This seems counter-intuitive, especially since many computer vision journals do not accept purely simulated results. However, a simulator allows a battery of statistical tests over a wide range of testing parameters that would be impractically expensive to perform in live testing. If the simulator is sufficiently developed and validated, it can arguably provide better generalized performance metrics than live testing by covering a wider range of statistical variations.

This is the paradigm NASA and Neptec have used for pose estimate system verification and validation (V&V) and certification processes. The process breaks down as in the diagram of Figure 4.

This error propagation process describes testing for both target-based SVS and targetless, shape-based TriDAR. For SVS, NASA and Neptec used a battery of validated mission simulators. Initially, half-scale models were used in the lab to simulate space operations. After the initial operational flights of SVS, this process was replaced by general certification of the system. Data collection errors were modeled as target centroid errors based on the expected size of the target in the image plane. The sizes of targets for a mission were modeled using the Target Image Model (TIM) simulator. Certification was accomplished using a standard target array for a matrix of lab testing scenarios and comparing tracking and centroid error results with the specifications.

For SVS, object definition errors consisted of errors in the target positions within the array as well as coordinate transformations between arrays and reference points. These positions could be affected by survey errors and thermal and pressure changes. The SVS Accuracy Analysis Program (SAAP) performed Monte Carlo simulations of SVS operations using these input errors propagated through the pose algorithms to evaluate performance.

The SVS Accuracy Model (SAM) provided pose sensitivity to centroid measurement. An important output from SAM was which targets contributed the most towards pose accuracy and hence were the most sensitive to measurement error. These targets became the focus of real-time flight analysis. The equivalent lesson for shape-based pose estimation is identifying key features that contribute to the data alignment and are hence the most sensitive to measurement errors and artifacts.

Characterization of the input error sources and validation with full scale mission simulations in the lab were key factors to the acceptance of the above tools and process. Even with this analysis, the effect of shadowing on the targets could not be adequately simulated even though it was a primary cause of poor performance. Instead, the limitations of shadowing drove the operational timelines for SVS. Shadowing was indeed simulated, given the expected positions and orientations of the ISS and shuttle during a given flight. However, these simulations were used to pick operational timelines rather than to drive expected performance.

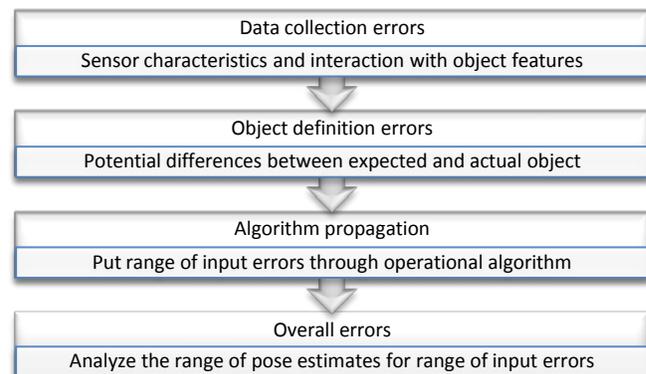


Figure 4: Error propagation modeling for V&V.

A similar process was used for TriDAR performance evaluation. For missions STS-128 and STS-131, a full TriDAR simulator was developed for the purposes of mission planning [36]. TriDAR sensor characteristics were measured and incorporated in the simulation, including Gaussian noise levels, outlier characteristics, beam divergence, and material reflectivity. Lab testing using scale models of ISS modules validated the simulator performance.

To ensure the simulation was of the highest fidelity, the communications interface was designed to be identical to that of the actual sensor. Therefore, the pose algorithms and hardware in the loop could not distinguish between the connections or data sources, whether from the simulator or the sensor. In addition to simulation fidelity, this enabled robust contingency planning by being able to simulate any hardware failure and plan according to the results. Finally, for simulation, the fidelity of the target model was just as important as the fidelity of the sensor characteristics. Accurate ISS models were provided by NASA in the correct configuration expected during shuttle rendezvous. These models allowed the pose estimation algorithms to test output over many solar panel and radiator configurations, which may move during rendezvous, robotic platform configurations, which are unpredictable, and Russian/European vehicle configuration, which may change due to launch schedule. A variety of approach trajectories were also tested to ensure robustness to different approach speeds, limited geometry in sensor field of view, and trajectories that are out of the approach corridor or off nominal. The evaluation method applied lessons learned from SVS and fully integrated the testing capabilities into the TriDAR operational system.

This simulation capability was critical for success. Given all of the above factors, no existing AR&D facility can accommodate such testing fidelity and versatility. Because the TriDAR pose estimation system was exposed to this type of simulation sophistication, it was able to perform 6 DOF real-time targetless tracking of the ISS the first time it acquired data from the target, and tracked the ISS throughout a complete revolution in pitch via fly-around during the second TriDAR mission on STS-131.

Although it is clear that there are too many variables to simply define or generalize pose performance, some variables may already be set in stone. For SVS, while some input was accepted for the positioning of SVS targets, optimal positioning could not always be met due to ISS constraints such as extra-vehicular activity (EVA) pathways and cabling. Another operational constraint was orbital day/night frequency causing changes in lighting environment. For TriDAR missions to the ISS, the addition of geometric targets was not permitted and the movable configuration of the ISS was not pre-determined. The shuttle rendezvous, approach, and docking trajectory was pre-defined up to an approach corridor. The sensor installation site on the shuttle could only be chosen from a select few positions. The laser on the sensor was qualified to be eye-safe, which reduced emitted power and thus

affected range performance. Finally, sensor operation was constrained such that it did not interfere with the shuttle crew during rendezvous operations.

As detailed above, these pose solution aspects are interdependent. The operational constraints changed the sensor, in terms of emitted laser power and the way it operated. The SVS sensor necessitated a specific design of SVS targets. In optimizing pose performance then, all aspects in the chain must be optimized together, rather than independently, as each aspect affects the others. In this context, standardization in an absolute context may be limited to characteristics of a specific component that has utility across a broad spectrum of applications.

Despite the challenges in quantifying performance, clearly there are measurable differences between pose estimation systems and many differences will cover a wide range of operations. Improvements in algorithms or sensor designs must improve something measurable or else the claim is moot. In this context, standardization may be useful for relative comparison of systems across a range of tests that span selected areas of the operating space. Neither the absolute nor relative performance in such tests may be generalizable, but they can highlight the strengths and weakness of various systems. The following section develops a framework for such tests.

5. EVALUATION FRAMEWORK

The premise for the proposed evaluation framework is similar to that of the performance evaluation processes for SVS and TriDAR. The framework allows for two types of testing: relative performance comparisons over a range of tests that span useful operational space, and absolute performance evaluation for specific operations in synthetic simulation using a model of the pose estimation system that has been certified using the same set of tests.

Clearly, the absolute evaluation is more involved and expensive, but modeling and validation only needs to be done once and can then be inexpensively applied to a variety of operational circumstances. Figure 5 shows a rough outline for the evaluation framework. The following sections discuss the framework components.

5.1 Pose Estimation System Categorization

The first step in evaluating a pose estimation system is to define the category of the system by its operational scenario. Target-based systems need to be evaluated separately from shape-based systems. Within target-based systems there are many sub-categories. Some target-based systems work with arbitrarily placed target arrays with no a priori knowledge of the target locations. Others use fixed target arrays to be attached to the object, or can be placed arbitrarily but require surveying of the targets to define an a priori array. (SVS falls into this category.)

There are also hybrid target/targetless systems. In one version, a 3D sensor images portions of the object and a second fixed sensor tracks targets on the first sensor. The

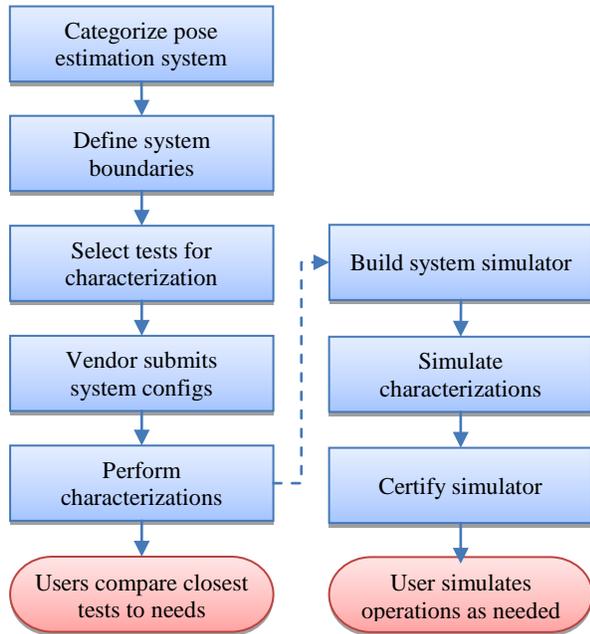


Figure 5: Proposed performance evaluation framework.

pose estimate of the object is then a function of both kinematic chains. Another type of hybrid system operates nominally as a targetless system for easy to track shapes, but uses fiducial targets added to difficult shapes to be tracked as reference points across multiple datasets.

Section 3 demonstrated that the purpose of the pose output is also critical in defining the category. Assembling parts requires accurate relative pose estimates. Docking, grasping, and navigation require accurate absolute pose estimates. (TriDAR falls in this category.) 3D imaging for inspection requires accuracy in the final model, not the pose estimate. (LCS falls in this category.) Dynamic tracking for in situ measurements require only approximate absolute pose estimates. (LMS falls in this category.)

Multi-use systems complicate the categorization step. TriDAR can perform imaging and inspection operations similar to LCS and LMS. This is not to say that categorizing TriDAR is impossible, but it would be re-categorized to be evaluated for each type of operation to ensure the proper metrics are evaluated.

It is outside the scope of this paper to exactly define the operational categories. It is also outside the scope to define the important performance metrics to quantify. Those will tend to follow the category. Generally speaking, pose estimation systems provide either 3 DOF or 6 DOF poses. The accuracy, precision, and resolution of the position and orientation are obvious metrics, as is the rate of pose output. Many other metrics may be defined. Some systems provide velocity estimates as well. Sensitivity of the pose estimation system to various conditions might be of interest, such as lighting, surface reflectivity, shape or target placement errors, and so on. This paper focuses on the evaluation process rather than specific metrics.

5.2 Pose Estimation System Boundaries

Within a category of operation, evaluation requires definitive boundaries between the system being tested and the test parameters. In some cases this may be fixed by the category. Pose estimation systems using a fixed target array can only be tested using that target array. Systems using arbitrary target placement should be tested across many array configurations with varying strength. Shape-based targetless system should be tested against a variety of shapes and surface types and reflectivity.

Sections 3 and 4 discussed in detail the range of potential parameters within a given system, including data collection parameters, pose algorithm parameters, and computational platform. Target tracking parameters (window sizes, feature matching, centroiding) may also fall into this category for target-based systems.

On the other side of the boundary are the operational parameters, including environment (lighting, clutter, thermal), object properties (size, shape, materials, reference point), and operational scenario (object and sensor motion, sensor viewpoints). From this division it might be tempting to define a set of tests where the operational parameters are held constant and the pose system parameters are varied, and vice versa. There are two problems with this approach: the number of permutations of parameters is prohibitively large, and the system parameters tend to be tuned to the operational scenario and, therefore, they are not independent factors.

These two problems tend to cancel each other out somewhat. The vast majority of permutations of pose system parameters and operational scenarios can be eliminated because the combinations would never be used together. Rather, for a given set of operational scenarios there are a limited set of realistic system parameters. This doesn't eliminate the problem, however, but rather offloads the problem to a submission of optimized parameters for each operational scenario for testing. The problem then becomes optimizing the parameters for each test scenario which may involve running tests using multiple parameters to find optimum settings. Hence the problem is not avoided but is transferred from the evaluation tests to the planning stage.

This problem is minimal for systems with limited parameter settings or if fixed instructions exist for optimizing the system. For other cases the solution may be to perform parameter selection through rigorous testing in synthetic environments. This is a bit of a chicken and egg approach as the testing itself may be used to validate the models used for simulation, but if validation does occur it confirms both the model and the parameters. If not, the model will need improving.

The paradigm in this framework is to compare systems in appropriate scenarios using appropriate parameters. If one performs consistently better, it is reasonable to predict that it will perform better in similar operations. It also provides an estimate of absolute system performance.

5.3 Standardized Tests

Two types of tests are proposed that model the V&V and certification tests in the NASA evaluation process. The first set characterizes the performance over a range of conditions and are therefore best described as characterizations. The test scenarios should span some realistic conditions for the category of operation. The results of these tests merely provide relative performance of systems, perhaps with an estimate of absolute performance for similar operational scenarios.

The most difficult part of such a process is the ability to define “similar”. Given the boundaries from the previous section, it is the operational parameters that need to be varied here. These include, but are not limited to, range, viewpoint, environment, relative motion, and, perhaps most importantly, object shape.

Object shape is perhaps the most difficult to quantify in terms of influence on performance. Section 6 provides a framework for defining shapes in such terms. Further, an optimum shape artefact for pose estimation is defined using these shape metrics. Hence a series of characterization tests can vary the object shape from weak to strong as defined by these metrics. In this context, “similarity” to a specific operation can be defined in terms of the shape metrics whereby the intended operational usage is analyzed in terms of these shape metrics.

The second set of tests certifies that synthetic models accurately represent the system performance. The purpose of certification is to demonstrate that an analytical model of the pose system is representative in terms of performance, providing confidence that simulation results provide a reliable estimate of system performance in operations not physically tested. Certification can provide a higher level of confidence and reliability in expected performance for a given operation. Certification does not include testing the performance of a system for the intended operation. Rather, it certifies that the model adequately represents the system performance within some boundaries. Synthetic operations can be performed on a case by case basis as needed by the system user to evaluate expected performance, whether for selecting systems or for predicting operational outcomes.

Certification is more complex and therefore more expensive for the system vendor. The expense of developing a synthetic operation that mimics the user’s scenario also adds expense to the user. However, for critical operations, this process is much less expensive than developing a laboratory mock-up test for every operation, or fully implementing a system to test functionality. Certification involves several components. The pose estimation system needs to be modeled end-to-end, including sensor measurements and algorithms. Characterization tests described in the prior section also need to be modeled in the same synthetic environment. The characterization tests can then be re-run in simulation and results compared to the results of the real characterization tests. Performance metrics can then be compared. Further,

synthetic testing allows a much wider range of tests with minimal extra effort. Small variations can be used to define statistical properties to compare to the real characterization results, as well as provide sensitivity analysis. Monte Carlo simulation can also bound expected performance.

Standards for the simulator platform and pose system would need to be well defined. As outlined in Section 4.4, SVS certification involved a battery of laboratory tests to demonstrate fiducial target tracking accuracy and robustness. The simulation environments of SVS Accuracy Analysis Program, SVS Accuracy Model, and Target Image Model could then reliably use the estimates for analysis and prediction. Missions were also simulated in a full 3D synthetic environment fed through the SVS video processing system to measure performance.

The lessons learned from SVS were used in the TriDAR testing environment. Instead of numerical models, the simulator included a full 3D synthetic environment and sensor simulator. The testing capability is built into the operational system. It is only a matter of switching the data source from the real sensor to the sensor simulator viewing the synthetic environment. In this case, the synthetic environment was not an additional effort because the pose estimation system uses the same synthetic models to track the real objects during real operations.

The same approach can be used for standardized certification. A standardized 3D simulation platform can operate modularly, with the pose estimation system abstracted out as a plug-in model, and the modeled scene and operation using standard format objects.

6. QUANTIFYING SHAPE FOR POSE ESTIMATION

One approach to quantifying shapes has been addressed through constraint analysis, an application of Principal Component Analysis. Constraint analysis directly assesses the sensitivity of the pose error to differential variation in pose and does not use geometrical features of the object. The assessment is performed through the numerical indices which are formed of the eigenvalues of the covariance matrix. Constraint analysis was introduced by Simon in [45] for optimal selection of points used to accurately position the patient for radiation therapy. A continuous version of constraint analysis, the Continuum Shape Constraint Analysis (CSCA) was presented in [32] and [43]. In CSCA, the discrete summation over points in the covariance matrix formulation was replaced by integrals over triangles constituting the CAD model of the shape. This way, the indices derived from CSCA become independent of the scanning process and, if the mesh is fine enough, reflect the pure geometry of the shape. CSCA indices were used for selection of an optimal scanning direction and for selection of the “best” scanning area on the object’s surface. The following section give a brief insight into CSCA theory and how it was used to generate an optimized shape for pose estimation.

Principal Component Analysis is also widely used in computer vision for face and gesture recognition. The

initial application of it for facial recognition and detection was most likely introduced by Turk & Pentland in [48], and further developed in [3] and [12]. For the use of PCA for gesture recognition, see, for example, [8].

The detailed discussion on invariant features (see, for example, [17], [18], [41], [44], and [49]), shape descriptors ([5], [22] and [46]) and identification of salient features in range images ([24] and [29]) and their usage in shape registration is beyond the scope of this paper.

6.1 Constraint Analysis Indices

We consider a problem of alignment, or registration, of a shape with its misaligned copy. A point-based normalized cost function, which is minimized to obtain a small pose estimate, $\mathbf{p} = (\mathbf{d}, \boldsymbol{\theta}) = (d_x, d_y, d_z, \theta_x, \theta_y, \theta_z)$, of the misaligned shape, can be represented as

$$E = \frac{1}{2N} \sum_{i=1}^N \Delta \mathbf{r}_i^T \Delta \mathbf{r}_i \quad (9)$$

where $\Delta \mathbf{r}_i = \mathbf{r}'_i - \mathbf{r}_i = \mathbf{d}_i + (\mathbf{1} + \boldsymbol{\theta}^\times) \mathbf{r}_i$. The notation $\mathbf{c} = \mathbf{a}^\times \mathbf{b}$ is the matrix implementation of the vector cross-product $\mathbf{c} = \mathbf{a} \times \mathbf{b}$ with the skew symmetric construct operator defined by

$$\mathbf{a}^\times = \begin{pmatrix} a_x \\ a_y \\ a_z \end{pmatrix} = \begin{pmatrix} 0 & -a_z & a_y \\ a_z & 0 & -a_x \\ -a_y & a_x & 0 \end{pmatrix} \quad (10)$$

In point-to-plane registration, $\Delta \mathbf{r}_i$ is represented as $\Delta \mathbf{r}_i = \mathbf{n}_i \mathbf{n}_i^T (\mathbf{1} - \mathbf{r}_i^\times) \mathbf{p}$, and the cost function becomes

$$E = \frac{1}{2} \mathbf{p}^T \mathbf{E} \mathbf{p} \quad (11)$$

where the matrix \mathbf{E} is

$$\mathbf{E} = \frac{1}{N} \sum_i \mathbf{E}_i = \frac{1}{N} \sum_i \begin{pmatrix} \mathbf{n}_i \mathbf{n}_i^T & -\mathbf{n}_i \mathbf{n}_i^T \mathbf{r}_i^\times \\ \mathbf{r}_i^\times \mathbf{n}_i \mathbf{n}_i^T & -\mathbf{r}_i^\times \mathbf{n}_i \mathbf{n}_i^T \end{pmatrix} \quad (12)$$

and \mathbf{n}_i is the unit-length normal vector to the surface at the point \mathbf{r}_i . Using eigenvalues $\{\lambda_k\}$ and eigenvectors $\{\mathbf{q}_i\}$ of the matrix \mathbf{E} , the cost function can also be represented as

$$E = \frac{1}{2} \mathbf{p}^T \mathbf{E} \mathbf{p} = \frac{1}{2} \sum_1^6 \lambda_i (\mathbf{p}^T \mathbf{q}_i)^2 \quad (13)$$

The eigenvector \mathbf{q}_6 , which corresponds to the transformation with minimal error, is an important parameter: larger values of λ_6 entail larger residual errors, and the pose algorithm continues its iterations longer delivering better pose accuracy.

The constrained analysis indices are functions of the eigenvalues $\{\lambda_k\}$. The following relations between the norm of the small pose error $\|\boldsymbol{\varepsilon}_p\|$, the standard deviation $\sigma_{\boldsymbol{\varepsilon}_p}$ of the small pose error and the upper boundary ε residual error hold:

Minimum Eigenvalue Index:

$$I_{min} = \sqrt{\lambda_{min}} \quad (14)$$

$$\|\boldsymbol{\varepsilon}_p\| \leq \frac{1}{I_{min}} \varepsilon$$

Inverse Condition Number Index:

$$I_{invcond} = \sqrt{\frac{\lambda_{min}}{\lambda_{max}}} \quad (15)$$

$$\frac{\|\boldsymbol{\varepsilon}_p\|}{\|\mathbf{p}\|} \leq \frac{1}{I_{invcond}} \varepsilon$$

Noise Amplification Index (NAI) [30]:

$$I_{NAI} = \frac{\lambda_{min}}{\sqrt{\lambda_{max}}} \quad (16)$$

$$\|\boldsymbol{\varepsilon}_p\| \leq \frac{1}{I_{NAI}} \varepsilon$$

Expectivity Index [32]:

$$I_{exp} = \left(\sqrt{\sum_k \frac{1}{\lambda_k}} \right)^{-1} \quad (17)$$

$$\sigma_{\boldsymbol{\varepsilon}_p} = \frac{1}{I_{exp}} \frac{\sigma_{data}}{\sqrt{N}}$$

where σ_{data} is the upper boundary on the measurement noise on the data points.

The above formulas show how the indices can be used to assess an object's view for pose estimation: high values for each index indicate better pose performance in some aspect. (They generally cannot all be optimized simultaneously.) Maximizing I_{min} , we minimize the smallest semi-axis of the error ellipsoid $\mathbf{p}^T \mathbf{E} \mathbf{p} = 1$, and hence limit the largest pose DOF error. Maximizing I_{NAI} , we minimize the eccentricity of the error ellipsoid, meaning we spread the error as evenly as possible. Maximizing $I_{invcond}$, we minimize the range of pose errors across the DOFs. The Expectivity Index accounts for all eigenvalues and provides an exact estimation of the standard deviation of the pose error norm. Figure 6 illustrates the above relation and presents a graph of the NAI vs. norm of the pose error for the Space Shuttle. Note that pose error decreases with NAI index up to some value above which little statistical improvement is achieved.

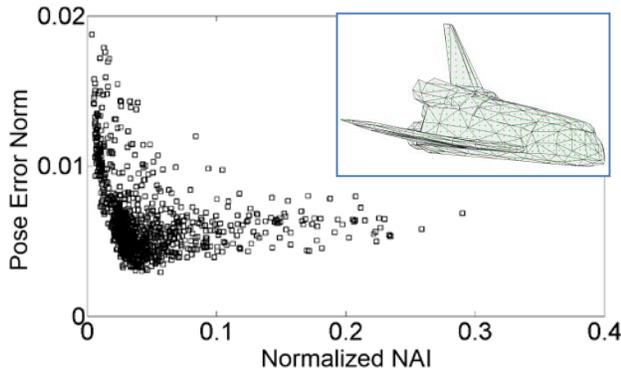


Figure 6: Shuttle model (inset) with NAI vs. pose error norm over 1000 pose views.

In [28], the basis of Continuous Shape Constraint Analysis (CSCA) was developed. The summation over a discrete set of points was replaced by integration over the triangles which form the CAD model of the shape. Moreover, a viewing direction, \mathbf{v} , was taken into account. The continuous version of the matrix \mathbf{E} is:

$$\mathbf{E} = \int_S \begin{pmatrix} n_i n_i^T & -n_i n_i^T r_i^x \\ r_i^x n_i n_i^T & -r_i^x n_i n_i^T \end{pmatrix} v dS \quad (18)$$

where the viewing factor v is

$$v = \begin{cases} \mathbf{v}^T \mathbf{n}, & \text{if } dS \text{ is unobstructed: } \mathbf{v}^T \mathbf{n} > 0 \\ 0, & \text{if } dS \text{ is obstructed: } \mathbf{v}^T \mathbf{n} < 0 \end{cases} \quad (19)$$

CSCA versions of the above indices can be graphically represented using a view-sphere. Scanning directions are mapped onto the sphere of a desired resolution to create a sphere-mesh as shown in Figure 7 (left). The index value for a given shape for a particular scanning direction is found by projecting this direction out from the origin through the sphere surface, and the radial distance of the intersection is then taken as the index value. Figure 7 (right) shows the spherical map of the Expectivity Index, along with the cube polyhedron and the view-sphere. To read the function value from any given scanning direction, one can simply read the radial distance to the surface of the function-map along the direction of the vector.

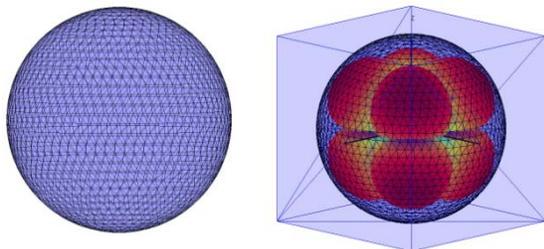


Figure 7: View-sphere (left) and graphical representation of Expectivity Index for a cube (right).

6.2 Optimized Shape for Pose Estimation

Constraint analysis was used to develop a methodology for design of a 3D target shape for accurate pose estimation. The shape selection was based on the Expectivity Index. The index was computed for different families of solids, and among 116 solids, the symmetric cuboctahedron was selected as a shape which has the highest minimum and maximum values of Expectivity Index over the 360 degrees scanning cone [10]. Figure 8 (left) shows the inhomogeneous index map for the dodecahedron compared to the steady high index value map for the cuboctahedron in Figure 8 (right).

The symmetric cuboctahedron delivers ambiguous poses. This means that registration errors are equally low for different poses and the cost function has multiple local minima. The symmetric cuboctahedron has 24 rotational symmetries and its cost function has 8 local minima. To overcome this ambiguity, the symmetric cuboctahedron was transformed into an asymmetric shape, named the Reduced Pose Ambiguity Cuboctahedron (RPAC). Each face of the RPAC has a unique set of dihedral angles $\theta_{kj} = \cos^{-1}(n_k, n_j)$ between this face and the adjacent faces. The deformation parameters were also chosen in such a way that the Expectivity Index for RPAC had high values from all views. The RPAC is shown in Figure 9.

To validate the suitability of the RPAC for pose estimation, a plastic model of the shape was scanned using Neptec's LCS. The graph of the Expectivity Index vs. pose error presented in Figure 10 shows that all RPAC's views deliver approximately equal pose errors.

The uniqueness of the RPAC's dihedral angles implies a possibility to use a look-up table (LUT) for pose. A LUT could therefore provide an independent pose measurement from that of the generalized shape pose algorithm, such as ICP, or initialization estimate to test the system. If used as a fiducial target, the RPAC and LUT can provide high-accuracy 3D pose estimation at high frequency since the LUT requires few surface measurements and no iterations. In this context, it can be attached to a moving object or act as a fixed object as viewed from a moving sensor, providing rapid and accurate 6 DOF pose.

A LUT was developed using a CAD model of the RPAC, and a representation of the stored reference data

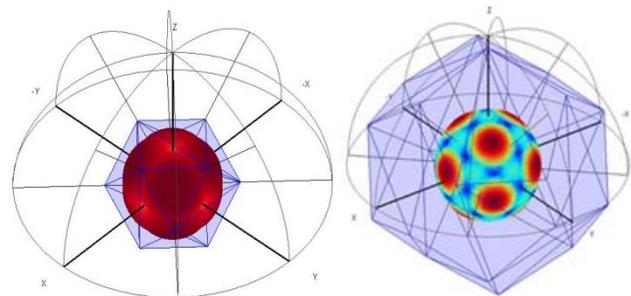


Figure 8: Expectivity Index maps for Cuboctahedron (left) and Dodecahedron (right)

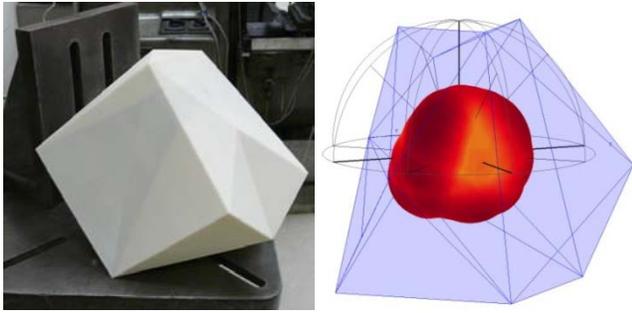


Figure 9: Reduced pose-ambiguity cuboctahedron (RPAC): plastic model (left) and Expectivity Index map (right)

can be seen in Table 1. The unique geometry is expressed in the LUT in terms of relationships between specific sets of faces, one set for each face of the RPAC. Each set refers to a central face (face 0) and the three neighbouring faces with which it shares an edge (faces A, B and C). The LUT incorporates: a) Normal vectors N_i and vertices $[V_1, V_2, V_3]_i$ for each face i of the RPAC, $1 \leq i \leq 20$; b) angles θ_A, θ_B and θ_C between the normal of face A, B or C and that of the central face 0. Within each set, all the data are sorted into that specific $[0,A,B,C]$ sequence. The sequencing criterion is different for different sets; c) Group numbers G_A, G_B and G_C of the angles in a particular set (all angles, according to their numerical values, are divided into groups); d) The sequencing criterion M used to put a given set of data into its $[0,A,B,C]$ sequence.

Table 1: Data representation in RPAC LUT.

Central Face (Face 0)	Face A	Face B	Face C	[A,B,C] Sequencing Criterion
	$[\theta_A]_i$	$[\theta_B]_i$	$[\theta_C]_i$	
	$[G_A]_i$	$[G_B]_i$	$[G_C]_i$	
$[N_0]_i$	$[N_A]_i$	$[N_B]_i$	$[N_C]_i$	
$[V_{0,1}]$ $[V_{0,2}]$ $[V_{0,3}]_i$	$[V_{A,1}]$ $[V_{A,2}]$ $[V_{A,3}]_i$	$[V_{B,1}]$ $[V_{B,2}]$ $[V_{B,3}]_i$	$[V_{C,1}]$ $[V_{C,2}]$ $[V_{C,3}]_i$	M_i

The pose algorithm using the above LUT data is called Face Normal Pose Estimation (FNPE) [40]. The steps of FNPE can briefly described as follows: 1) Identify the planes and normals in the point cloud data using RANSAC and then identify best-fit triangles; 2) Identify the required $[0,A,B,C]$ pattern in the fitted triangles and calculate the corresponding angles between faces; 3) Execute the table lookup algorithm (this procedure also determines the best-fit rotation between point cloud normals and reference normals); 4) Rotate the model to the point cloud using the previously found best-fit rotation, and determine the translational vector by simultaneously minimising the distances between each set of points in the point cloud and its corresponding reference plane (from the rotated model).

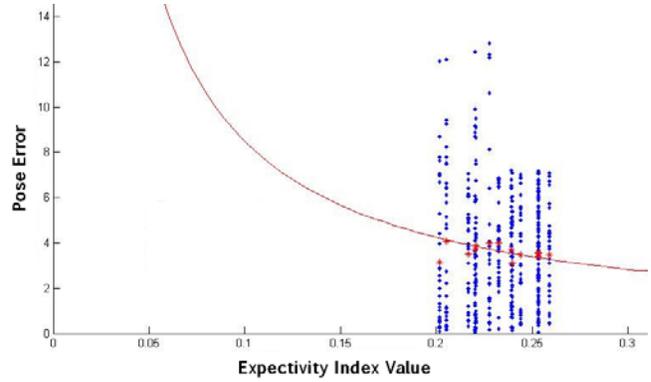


Figure 10: Experimental results for RPAC.

Figure 11 represents the result of an experiment in which 50 FNPEs were generated for each of 15 different lidar scans of the RPAC, scanned using Neptec’s LCS. Since no truth data were available for the poses of the RPAC in these scans, each FNPE was used to initialize the ICP, which was run using a “cleaned” version of the point cloud (to maximize the accuracy of the final ICP pose). For each different point cloud, FNPE error (or ICP correction) is the difference between an individual FNPE and the mean ICP pose estimate of that point cloud. The ICP orientation correction is described using ϕ , which is the magnitude (in degrees) of the angular component of the error quaternion between the FNPE and ICP poses. Similarly, the ICP displacement correction is the magnitude (in [mm]) of the displacement vector. The fit quality of the FNPE is assessed using the root-mean-squared distance between each point in the point cloud (also called root-mean-squared error, or RMSE), and the corresponding closest point on the model. The two curves shown in the figure are cumulative error density plots of the results, in terms of orientation error and RMSE.

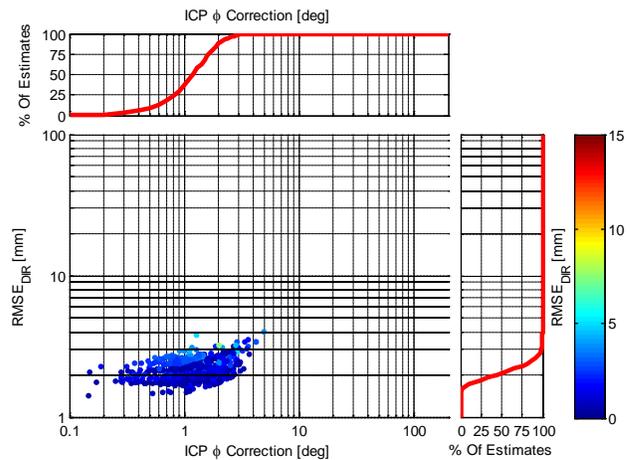


Figure 11: Error characteristics of 750 FNPE of the RPAC using scans from Neptec’s LCS. Colour represents ICP displacement correction in mm.

FNPE was also tested in simulations, using simulated lidar scans generated from 2000 views around the RPAC,

evenly distributed in view angle. Figure 12 (top) shows a sinusoidal projection of the positions of all 2000 viewpoints, coloured according to the RMSE (in mm) of the pose estimate generated from that view. For each viewpoint, the FNPE pose was used to initialise ICP (using the entire point cloud, not a “cleaned version”). Figure 12 (bottom) shows the same sinusoidal projection of viewpoints, coloured according to the RMSE of the final ICP pose. Empty portions of the plot represent views from which an FNPE could not be made. There are two possible reasons for this: the pattern of faces needed for pose estimation is simply not visible when scanning the shape from that perspective, or the pattern of faces is technically visible but could not be reconstructed from the available point cloud data. The lower left portion of the plot shows a region demonstrating the former case. Many of the failures in the right hand portion of the plot reflect the latter. Re-optimisation of the RPAC could possibly prevent the first scenario entirely, and reduce the degree of algorithmic optimisation necessary to prevent the second. The high error observed in the upper right hand portion of the plot, is the result of a combination of the (parametric and algorithmic) implementation of FNPE, point cloud noise, scan density, and basic RPAC geometry.

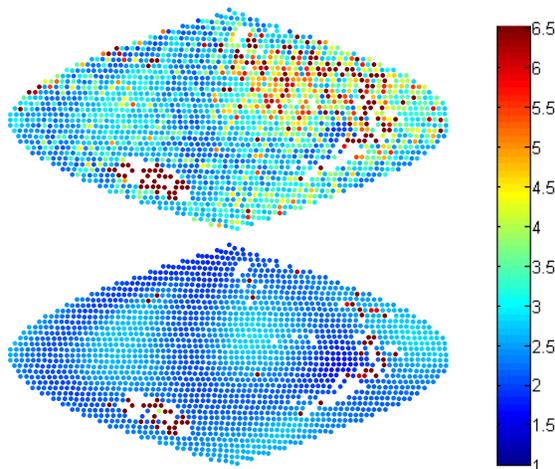


Figure 12: Root-Mean-Squared Error between point cloud and model corresponding to each of 2000 evenly distributed views around the RPAC: (top) LUT-based pose estimate, (bottom) after subsequent application of ICP using LUT-based estimate as initial guess. The scale is in mm.

When the FNPE technique was applied to real data from Neptec’s LCS, the average difference between the estimated orientation and an ICP-based orientation was 1.03 degrees. The average difference between the estimated position and ICP-based position was 1.08 mm. Although the FNPE technique requires more development, it is clear from the experimental and simulated results shown that this approach can be used with the RPAC for accurate pose estimation, and warrants further study.

6.3 Application to Performance Evaluation

The constraint analysis indices provide a set of metrics that define shapes in terms of their pose estimation strengths, and the Expectivity Index correlates with the norm of the pose error vector. A set of standard object artifacts can be developed covering a range of pose strengths and used in the characterization tests described in Section 5.3.1. The standard artifacts can then be described by their indices in tables. Additionally, a standard toolset would need to be developed to define the metrics for a particular user’s operational object(s) or scene.

These artifacts and tables then complete the characterization and certification tests. For simpler operations, a user would simply analyze their object to determine the shape metrics, and look up test results for the closest representative test, or interpolate between bounding tests. This would quickly narrow down the pose estimation systems that can meet their operational needs.

For more critical operations where synthetic simulation of the operation is necessary, system certification would be accomplished using the standard artifacts, both for real in the characterization tests and repeated in a synthetic environment using virtual artifacts.

Finally, the standard artifacts could be used by vendors of pose estimation systems for self-improvement over a range of object pose weaknesses and to optimize, and possibly automate, their own system parameter settings by evaluation of object shape metrics.

7. SUMMARY AND CONCLUSIONS

This paper presented many of Neptec’s lessons learned over 20 years of real-time, dynamic pose estimation system development, operations, and performance evaluation. Emerging sensor technologies are creating even higher data rates with smaller sensor footprint and better reliability. smart sensing techniques use information extracted from the data to drive the gathering of new data to optimize whole system efficiency, allowing very high dynamic range and significantly reducing wasted effort gathering and processing unnecessary data.

These capabilities have introduced challenges in the performance evaluation of pose estimation systems. More traditional systems use fiducial targets, often in fixed arrays that limit the operational parameters. The newer shape-based, flexible, smart sensing pose systems open up a wide variety of new possible operations, with a selection of examples provided for which Neptec has operational applications at various levels of development. Each application uses the pose estimate differently and so relies on different performance characteristics. Additionally, in being this flexible, these systems provide significantly more parameters for optimization and testing.

Certification of the target-based SVS for assembling the ISS on over 20 shuttle missions has provided many lessons learned for designing V&V and certification tests. Many of these lessons were incorporated into the TriDAR

performance evaluation process. A critical lesson for flexible pose estimation systems is that performance cannot be generalized due to heavy reliance of performance on object shape and views.

What can be generalized are the process steps to characterize a system over a range of operations and object shapes, as well as process steps to certify system models for synthetic testing of the intended operation. The former tests might be sufficient for simpler operations, whereas certified simulators and synthetic testing may be required for more critical operations where confidence in performance is necessary, thereby reducing risks and expenses of alternative means of testing. A generalization of these process steps have been turned into a proposed performance evaluation framework presented in this paper.

The reliance on shape potentially undermines the proposed approach, but this is solved by the use of metrics that describe objects and views in terms of their pose estimation strengths and weaknesses. These metrics have been used to design, build, and test an optimized shape (RPAC) for system performance evaluation by removing object weaknesses from the operation. A lookup table for the RPAC has also been developed for comparison and initialization of pose estimation systems, and can also be used as a high-fidelity, high-speed 3D target for many real-time 6 DOF pose estimation problems. The same design process can be used to vary the shape strengths to define a standard test set for characterization and certification.

Much of the proposed framework for performance evaluation is based on experience and completed works, including testing of the RPAC and building a LUT for rapid pose estimation. However, many details still need to be addressed such as classification systems, ranges of operations, ranges of shape metrics and artifacts, standardization of synthetic testing tools, and even the particular system metrics to test.

ACKNOWLEDGEMENT

The authors would like to thank additional contributors to the development of this work including Donald McTavish, John Enright, Gail Bouchette, and Stephane Ruel. We would also like to thank the support of the National Sciences and Engineering Research Council of Canada (NSERC) and the Canadian Space Agency.

REFERENCES

- [1] Aghili, F., Kuryllo, M., Okouneva, G., and English, C. 2011. Fault-Tolerant Position/Attitude Estimation of Free-Floating Space Objects Using a Laser Range Sensor. *IEEE Sensors Journal*, vol. 11, no. 1, pp. 176-185.
- [2] Allen, A., Mak, N., and Langley, C. 2005. Development of a Scaled Ground Testbed for Lidar-based Pose Estimation, *IEEE IROS 2005* (Edmonton, AB, Aug 2). 10-15.
- [3] Becker, B., Ortiz, E. 2009. Evaluation of face recognition techniques for application to Facebook. *8th IEEE Int Conf on Automatic Face & Gesture Recognition*.
- [4] Besl, P.J., and McKay, N.D., 1992. A Method for Registration of 3-D Shapes, *IEEE Trans Pattern Analysis and Machine Intelligence*, 14(2), 239-256.
- [5] Belongie, S., Malik, J., and Puzicha, J. 2002. Shape Matching and Object Recognition Using Shape Contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol.24, No.4.
- [6] Blais, F., Beraldin, J.A., El-hakim, S., Godin, G. 2003. New Development in 3D Laser Scanners: From Static to Dynamic Multi-Modal Systems, *6th Conf Opt 3-D Meas Tech*, (Zurich, Switzerland, Sep 22-26).
- [7] Bouchette, G., Iles, P., English, C., Labrie, M., Powaschuk, B., Church, J., Maheux, J. 2007. Rapid automatic target recognition using generic 3D sensor and shape-from-motion data, *SPIE DSS* (Orlando, FL, Apr. 9-13). 6566:19.
- [8] Braido, P., Zhang, X. 2004. Quantitative analysis of finger motion coordination in hand manipulative and gestic acts. *Human Movement Science*, Vol. 22, No. 6, pp. 661-678.
- [9] Chen, Y., and Medioni, G., 1992. Object modeling by registration of multiple range images, *Image Vis Comput*, 10(3), 145-155.
- [10] Choudhuri, A., Okouneva, G., McTavish, D., Bouchette, G. 2010. Design of Optimal Shapes for Space Docking using LIDAR-Based Vision Systems, *ASTRO 2010* (Toronto, ON, May 4-6).
- [11] Deslauriers, A., English, C., Bennett, C., Iles, P., Taylor, R., Montpool, A. 2006. 3D Inspection for the Shuttle Return to Flight, *SPIE DSS* (Orlando, FL, Apr 17-20). 6220.
- [12] Draper, B., Baek, K., Bartlett, M., Beveridge, J. 2003. Recognizing faces with PCA and ICA. *Computer Vision and Image Understanding*, Vol. 91, No. 1-2, pp. 115-137.
- [13] English, C., Ruel, S., Melo, L., Church, P., Maheux, J. 2004. Development of a Practical 3D Automatic Target Recognition & Pose Estimation Algorithm. *SPIE DSS* (Orlando, FL, Apr 12-16). 5426:112-123.
- [14] English, C., Zhu, S., Smith, C., Ruel, S., Christie, I. 2005. Tridar: A Hybrid Sensor for Exploiting the Complimentary Nature of Triangulation and LIDAR technologies, *i-SAIRAS* (Munich, Germany, Sep 5-8). ESA SP-603, p.79.1.
- [15] ESA, 2008. Rendezvous and Docking Technology, ATV Information Kit, February 2008.
- [16] ESA Media Relations Office, 2008. Europe's automated ship docks to the ISS. *European Space Agency News*, April 3.
- [17] Gelfand, N., Ikemoto, L., Rusinkiewicz, S., and Levoy, M. 2003. Geometrically Stable Sampling for the ICP Algorithm. *Proc. of 4th International Conference on 3-D Digital Imaging and Modeling (3DIM 2003)*, Stanford CA, October.
- [18] Gelfand, N., Mitra, N.J., Guibas, L.J., Rottmann, H. 2005. Robust Global Registration. *Eurographics Symposium on Geometry Processing (SGP 2005)*.
- [19] Harrison, J.W., Ferrie, F.P., Heffor, S.W., Samson, C., Kusevic, K., Mrstik, P., Iles, P. 2009. Finding Anomalies in High-Density LiDAR Point Clouds, *GEOMATICA*, 63(4):397-405.
- [20] Howard, R.T., Heaton, A., Pinson, R., Carrington, C., Lee, J., Bryan, T., Robertson, B., Spenser, S., Johnson, J. 2008. Orbital Express Advanced Video Guidance Sensor. In *2008 IEEE Aerospace Conference* (Big Sky, MT, March 1-8).
- [21] Jasiobedzki, P., Se, S., Pan, T., Umasuthan, M. Greenspan, M.. 2005. Autonomous Satellite Rendezvous and Docking Using LIDAR and Model Based Vision, *SPIE DSS* (Orlando, FL, Mar 28-31). 5788:54-65.
- [22] Johnson, A., and Herbert, M. 1999. Using Spin Images for Efficient Multiple Model Recognition in Cluttered 3-D Scenes, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol.21, No.5.

- [23] Kneisel, K., Gemme, S., Aghili, F., Okouneva, G., English, C. 2010. Towards Vision-Based SLAM for Planet Exploration, *ASTRO 2010* (Toronto, ON, May 4-6).
- [24] Lee, C.H., Varshney, A., Jacobs, D.W. 2005. Mesh Saliency. *Proc. of SIGGRAPH'05*, pp. 659-666.
- [25] Leinz, M., and Chen, C. 2002. Autonomous Rendezvous and Capture Sensor System (ARCSS), *2002 Core Technologies for Space Systems Conference* (Colorado Springs, CO, Nov)
- [26] MacLean, S., and Pinkney, L. 1993. Machine Vision in Space, *Can. Aeronaut. Space J.*, 39(2), 63-77.
- [27] Maurelli, F., Petillot, Y., Mallios, A., Krupinski, S., Haraksim, R., Sotiropoulos, P. 2009. Investigation of portability of space docking techniques for autonomous underwater docking, *OCEANS 2009-EUROPE* (Bremen, Germany, May 11-14).
- [28] McTavish, D.J., Okouneva, G., and English, C.. 2010. Continuum Shape Constraint Analysis: A class of Integral Shape Properties Applicable for LIDAR/ICP-based Pose Estimation. *Accepted to International Journal of Shape Modeling*.
- [29] Mortara, M., Spagnuolo, M. 2009. Semantics-driven Best View of 3D shapes. *Computers & Graphics*, No. 33, pp. 280-290.
- [30] Nahvi, A., Hollerbach, J.M. 1996. The Noise Amplification Index for Optimal Pose Selection in Robot Calibration. *Proc. IEEE Intl. Conf. Robotics and Automation*, 22-28.
- [31] Ohkami, Y., and Oda, M. 1999. NASDA's activities in space robotics, *i-SAIRAS '99* (Noordwijk, The Netherlands June 1-3). 11-18.
- [32] Okouneva, G., McTavish, D., Choudhuri, A. 2010. Principal Component Analysis for Pose Estimation Using Range Data. *Computers and Simulation in Modern Science*, v. IV, pp. 111-124, WSEAS Press.
- [33] Piechocinski, S., and Sasiadek, J. 2005. Experimental determination of relative motion measurement accuracy for an auto-synchronous triangulation scanning laser camera, *SPIE DSS* (Orlando, FL, Mar 30-Apr 2). 5791:208-217.
- [34] Ruel, S., English, C., Melo, L., Berube, A., Aikman, D., Deslauriers, A., Church, P., Maheux, J. 2004. Field testing of a 3D automatic target recognition and pose estimation algorithm, *SPIE DSS* (Orlando, FL, Apr 12-16). 5426:102-111.
- [35] Ruel, S., English, C., Anctil, M., Church, P. 2005. 3DLASSO: Real-time pose estimation from 3D data for autonomous satellite servicing, *i-SAIRAS*, (Munich, Germany, Sep 5-8). ESA SP-603.
- [36] Ruel, S., Luu, T., Anctil, M., Gagnon, S. 2008. Target Localization from 3D data for On-Orbit Autonomous Rendezvous & Docking, *IEEE Aerospace Conference* (Big Sky, MT, March 1-8). 1-11.
- [37] Ruel, S., and Luu, T. 2010. STS-128 on-orbit demonstration of the TriDAR targetless rendezvous and docking sensor, *IEEE Aerospace Conference* (Big Sky, MT, Mar 6-13).
- [38] Rumford, T. 2002. Demonstration of Autonomous Rendezvous Technology (DART) Project Summary, *2002 Core Technologies for Space Systems Conference*, Nov.
- [39] Rusinkiewicz, S., Hall-Holt, O., and Levoy, M., 2002, Real-Time 3D Model Acquisition, *SIGGRAPH 2002*, (San Antonio, TX, Jul 21). 21(3):438-446.
- [40] Saint-Cyr, P., 2010. A Framework for Non-ICP LIDAR-based Pose Estimation Using an Optimally Constrained 3D Target. MASC. Thesis, Ryerson University.
- [41] Salvi, J., Matabosch, C., Fofi, D., and Forest, J. 2007. A review of recent range image registration methods with accuracy evaluation. *Image and Vision Computing*, vol. 25, no. 5, pp. 578-596.
- [42] Samson, C., English, C., Deslauriers, A., Christie, I., Blais, F., Ferrie, F. 2004. Neptec 3D Laser Camera System: From Space Mission STS-105 to Terrestrial Applications, *Canadian Aeronautics and Space Journal*, 50(2) 115-123.
- [43] Shahid, K., and Okouneva, G. 2007. Intelligent LIDAR Scanning Region Selection for Satellite Pose Estimation. *Computer Vision and Image Understanding*, 203-209.
- [44] Sharp, G., Lee, S., and Wehe, D. 2002. ICP registration using invariant features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 1, pp. 90-102, Jan.
- [45] Simon, D.A. 1996. "Fast and Accurate Shape-Based Registration", Ph.D. Thesis, Report CMU-RI-TR-96-45, Carnegie Mellon University.
- [46] Taati, B., and Greenspan, M. 2011. Local shape descriptor selection for object recognition in range data. To appear in *Computer Vision and Image Understanding*, Vol. 115, No. 5, 681-694.
- [47] Tubic, D., Hebert, P., Deschenes, J.-D., Laurendeau, D. 2004. A unified representation for interactive 3D modeling, *3DPVT 2004* (Thessaloniki, Greece, Sep 6-9). 175-182.
- [48] Turk, M., Pentland, A. 1991. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, Vol. 3, No. 1, pp. 71-86.
- [49] Tuytelaars, T., and Mikolajczyk, K. 2007. Local Invariant Feature Detectors: A Survey. *Foundations and Trends in Computer Graphics and Vision*, vol. 3, no. 3, pp. 177-280.
- [50] Wettergreen, D., Jonak, D., Kohanbash, D., Moreland, S., Spiker, S. 2009. Design and Experimentation of a Rover Concept for Lunar Crater Resource Survey, *47th AIAA Aerospace Sciences Meeting* (Orlando, FL, Jan 5-8).
- [51] Zhu, X., Miller, S., Kwan, M., and Smith, I.C. 2005. A high resolution 3D laser camera for 3D object digitization, *SPIE Defense and Security Symposium* (Orlando, FL, Mar 30 - Apr 1). 5791:120-127.
- [52] Zhu, X., Church, P. 2008. Lidar for obstacle detection during helicopter landing, *SPIE Defense and Security Symposium* (Orlando, FL, Apr). 6950:1-8.
- [53] Zhu, X., Church, P., and Labrie, M. 2010. AVS LIDAR for Detecting Obstacles Inside Aerosol, in *Vision and Displays for Military and Security Applications*, Springer Science + Business Media, Chapter 13, 177-187.

Dr. Chad English is Director of R&D at Neptec Design Group in Ottawa, Canada. He received a B.Eng (1994) jointly from Acadia University, and the Technical University of Nova Scotia, Halifax, Canada, and an M.Eng. (1996) and Ph.D. (2000) from Carleton University, Ottawa. He was awarded senate medal for outstanding achievement at the doctoral level for his research in robotics and prosthetics. Dr. English joined Neptec Design Group in 1999 as an engineering analyst for the SVS and LCS which he supported for many flights. In this role he was awarded a NASA GEM award and a Group Achievement Award.

Dr. Galina Okouneva received her M.A.Sc. in Mechanics (1982) and Ph.D. degree in Physics and Mathematics (1987) from Moscow State University, Russia. From 1989 to 1992 she was an Associate Professor of the Dept of Computer Science, Tver State University, Russia. In 1993 she joined the University of Aizu, Japan, as a visiting researcher. From 1998 to 2003, Dr. Okouneva was a Research Scientist at Neptec Design Group, Ottawa. Since 2003, she is Associate Professor in the Dept. of Aerospace Engineering of Ryerson University, Toronto, Canada. Her research interests include computer vision, computational geometry, dynamics and robotics.

Pierre Saint-Cyr completed his B.Eng. in Aerospace Engineering at Ryerson University, Toronto, Canada, in 2008. His interest in the research at Ryerson in computer vision and computational geometry prompted him to remain for his M.A.Sc. He received his M.A.Sc. in Aerospace Engineering from Ryerson University in 2011.

Aradhana Choudhuri received her B.Eng. in Aerospace Engineering (2007) and M.A.Sc (2009) at Ryerson University, Toronto, Canada. Her research interests lie in computer vision, robotics, space and satellite mechanics. Currently Aradhana Choudhuri is a junior engineer at the Canadian Space Agency.

Tim Luu is a Vision Systems Engineer at Neptec Design Group, Ottawa. He received a B.A.Sc. from the University of British Columbia in Engineering Physics and an M.A.Sc from Carleton University in Mechanical Engineering. At Neptec, he has helped develop the TriDAR vision system for AR&D of satellites. He has also filled the roles of lead analyst, test lead, and operations engineer in support of shuttle missions STS-128 and STS-131.