INTEGRITY DETERMINATION FOR IMAGE RENDERING VISION NAVIGATION

DISSERTATION

Sean M. Calhoun, Contractor

AFIT-ENG-DS-16-M-251

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

The views expressed in this dissertation are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the United States Government.

AFIT-ENG-DS-16-M-251

INTEGRITY DETERMINATION FOR IMAGE RENDERING VISION NAVIGATION

DISSERTATION

Presented to the Faculty

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

Sean M. Calhoun, B.S., B.S.E.E., M.S.E.E

March 2016

AFIT-ENG-DS-16-M-251

INTEGRITY DETERMINATION FOR IMAGE RENDERING VISION NAVIGATION

Sean M. Calhoun, B.S., B.S.E.E., M.S.E.E

Committee Membership:

John Raquet, PhD
Chair

Gilbert Peterson, PhD
Member

Jonathan Black, PhD
Member

ADEDJI B. BADIRU, PhD
Dean, Graduate School of Engineering and Management

Abstract

This research addresses the lack of quantitative integrity approaches for vision navigation, relying on the use of image or image rendering techniques. The ability to provide quantifiable integrity is a critical aspect for utilization of vision systems as a viable means of precision navigation. This research describes the development of two unique approaches for determining uncertainty and integrity for a vision based, precision, relative navigation system, and is based on the concept of using a single camera vision system, such as an electro-optical (EO) or infrared imaging (IR) sensor, to monitor for unacceptably large and potentially unsafe relative navigation errors.

The first approach formulates the integrity solution by means of discrete detection methods, for which the systems monitors for conditions when the platform is outside of a defined operational area, thus preventing hazardously misleading information (HMI). The second approach utilizes a generalized Bayesian inference approach, in which a full pdf determination of the estimated navigation state is realized.

These integrity approaches are demonstrated, in the context of an aerial refueling application, to provide extremely high levels (10^{-6}) of navigation integrity. Additionally, various sensitivities analyzes show the robustness of these integrity approaches to various vision sensor effects and sensor trade-offs.

To my family

Acknowledgments

I would like to thank the Air Force Research Laboratory for the opportunity to participate in research on automated aerial refueling. I would like to thank my advisor, Dr. John Raquet, and my entire research committee for their prolonged participation and support in this pursuit.

Most importantly, I would like to thank my wife and kids for all of their support, understanding, and unconditional love as I went through this process, without which, I would not have succeeded.

Table of Contents

Page

List of Figures

Page

List of Tables

Page

INTEGRITY DETERMINATION FOR IMAGE RENDERING VISION NAVIGATION

1 Introduction

Recently, there has been an increased recognition of Global Navigation Satellite System (GNSS) limitations in terms of robustness, availability, and interference. As a result of this recognition there has been renewed interest in developing non-GNSS based navigation systems to augment system capability. This has become particularly important with the trend toward autonomous systems, where required navigation performance (RNP) metrics, such as accuracy, integrity, continuity, and availability become operational drivers. Because of this trend, there is renewed interest in gaining navigational diversity utilizing imaging or vision aided navigation approaches.

Prior to the wide proliferation of Global Positioning System (GPS)-based navigation systems, much of the aviation navigation research revolved around developing methods for augmenting, and more importantly, bounding the error growth from Inertial Navigation Systems (INS). Radar systems were a major focus of this research. Systems such as TERCOM (terrain contour matching) [1] and SITAN (Sandia inertial terrain-aided navigation) [1][2] utilized radar measurements and their correlation to terrain databases to remove long term drift associated with INS. Recently, due to recognition of the vulnerabilities of GPS, there has been renewed interest in this line of research to develop non-GPS augmentation systems to aid INS.

Vision systems are now being heavily researched to provide navigation updates to aviation systems. Like the radar systems, vision systems can use 3-D terrain databases and imaging systems to provide periodic position updates [3]. Later research took the integration of vision sensors with INS systems further, where not only did the vision systems aid the estimation of the inertial errors, stochastic projections using the INS were used to aid in the image feature tracking process [4][5]. This approach is a very similar approach to more recent research focusing on multi-function uses of imaging and vision systems, such as SLAM (Simultaneous Localization and Mapping) [6][7][8][9], in which generation of the database and navigation are performed concurrently.

In addition to the use of vision systems in traditional aviation navigation, for applications such as formation flying, aerial refueling, rendezvous and docking systems, and even precision landing, there has been a significant body of research in the area of precise relative navigation [10][11][12][13]. In [10], Weaver presented a vision-based relative navigation solution for aerial refueling with the use of an *a priori* 3-D tanker model. His research helped provide motivation for this proposed vision navigation research, having shown the viability for the image rendering navigation approach.

For precision relative navigation applications such as formation flying, aerial refueling, rendezvous and docking systems, and even precision landing, there has been a significant body of research for the use of vision navigation systems [10][11][12][14]. In [10], Weaver presented a vision-based relative navigation solution for aerial refueling with the use of an *a priori* 3-D tanker model. A similar approach was demonstrated in [15]; results from these flight test showed that image rendering relative navigation is a

viable precision navigation technique for close formation flight, specifically aerial refueling. Flight test results from [15] demonstrated 95% relative navigation accuracies on the order of 35cm within the operational envelope.

As the body of vision-aided navigation research continues to grow, which to date has demonstrated the potential accuracy performance for these systems, consideration of other required navigation performance (RNP) metrics is required. Particularly, when considering safety-critical navigation applications, ensuring that systems are providing safe information and maintaining a high level of *integrity* is paramount, but is largely neglected in current vision navigation research.

1.1 Navigation Integrity for Flight Systems

The concept of integrity, particularly for navigation systems, refers to the level of trust that can be placed in a navigation system in terms of detecting gross errors and divergences [16]. The Radio Technical Commission for Aeronautics (RTCA), federal advisory commission for the Federal Aviation Administration (FAA) defines navigation integrity as the following [17]:

"The ability of a system to provide timely warnings to users when the system should not be used for navigation."

Integrity, as it relates to navigation systems, is a major consideration in the development and deployment of new systems, particularly for flight systems applications. Integrity for navigation systems equates to safety, and flight safety is often a major driver of system performance.

Many navigation applications have adopted the utilization of protection levels (PL), which are real-time navigation system outputs that bound the navigation system errors (NSE) to the required probability of integrity risk. In this approach to ensuring integrity, the NSE is bounded by the real-time PL, and as the long as the PL is below the alert limit (AL), which are determined by system designers and safety considerations, the system can continue its operation. Loss of integrity is defined by the case when the NSE > AL without an alert or, in other words, when NSE > AL and PL ≤ AL, when PL's are utilized.

The integrity risk requirement is defined as the probability that the navigation error is larger than the alert limit without timely warning. Typical ranges of integrity risk requirements for precision navigation, depending on the system applications, can vary anywhere from $1x10^{-4}$ to $1x10^{-9}$, depending on the application and the associated safety criticality. There are several factors that can lead to navigation system integrity events in which the navigation error exceeds an alert limit without warning. These events include hardware failures, software failures, environmental anomalies, and algorithm failures.

One of the richest sources of information for how integrity can be handled for precision relative navigation systems can be found with the Local Area Augmentation System (LAAS) [18], a GPS based civil landing system. The LAAS system focused on providing integrity under fault-free and single ground reference receiver failure conditions. The computation of real-time protection levels for the LAAS system utilizes error bounding estimates of the signal-in-space and projects those onto the position domain for horizontal and vertical navigation frames [16][19]. A summary of the associated integrity requirements for the LAAS precision navigation systems are show in

Table 1-1. This table illustrates how integrity requirements can have significant safety implications for the navigation system.

Table 1-1: Integrity Requirements for the LAAS Precision Navigation Systems [16]

System	Coordinate Frame	Integrity Requirement	Alert Limit
LAAS	Lateral	$1x10^{-9}$ in any 30 sec.	17m
	Vertical	$1x10^{-9}$ in any 15 sec.	10m

Additionally, the LAAS system employs several other quality monitors to detect potential signal anomalies and failures in the satellite vehicle (SV). One example is receiver autonomous integrity monitoring (RAIM). RAIM utilizes redundancy in satellite measurements to detect possible single satellite failures and compensate for the largest possible undetected biases in the computation of the protection levels.

An illustration of LAAS's approach to integrity and the PL concept is illustrated in Figure 1. In this case the vertical navigation system error (NSE) is shown relative to the reported Vertical Protection Level (VPL) and the Vertical Alert Limit (VAL). The protection levels are real-time outputs from the navigation system that bound the navigation error to a defined probability and serves as the timely warning mechanism. Under conditions when the VPL exceeds the VAL (i.e., VPL > VAL), the system is considered unsafe for that operation / application, and such a case usually results in the ceasing of operations. The ceasing of operations is a result of unacceptably high probability of misleading information (MI). However, even under this condition, this is not considered an integrity event, since the system properly alerted the user or aircraft system. Under this scenario the only condition that is considered an integrity event is

when vertical NSE > VAL, but VPL < VAL. Thus the system doesn't properly bound the true error and fails to report an unsafe condition.

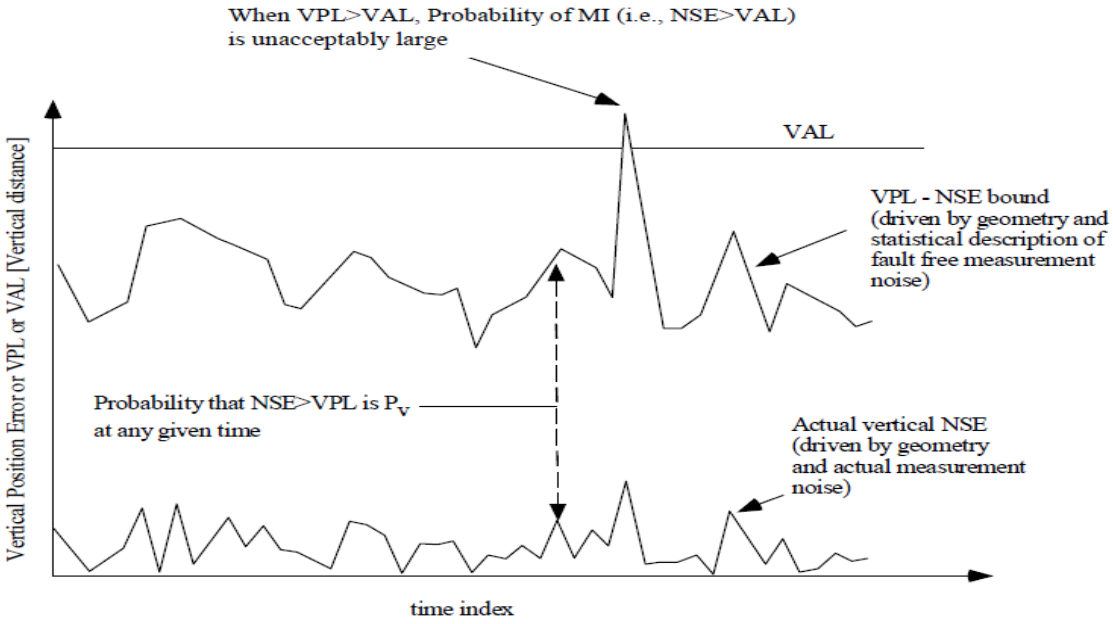

Figure 1: Depiction of Navigation Integrity Concepts (LAAS MASPS, D0-245B [18])

The computation of the protection level varies depending on the application and the specific implementation; however, an assumption often made is that the navigation errors can be modeled as a Gaussian process. Under this assumption, the computation of the protection level under nominal, no-fault conditions can be performed quite easily. Since the assumption is nominal, fault-free conditions, the only remaining considerations are process noise and measurement noise, which are handled using standard statistical analysis for Gaussian distributions. To illustrate this consider the probability density function (pdf) of a standard Gaussian distribution *f(x)*,

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(x-\mu)^2}{2\sigma^2}} \tag{1.1}$$

where σ is the standard deviation and μ is the mean. Like any pdf, to determine the probability that a sample drawn from this distribution falls within a given range, one must only integrate across that range. This is further illustrated in Figure 2, showing the probability that a sample of *x* falls within a region defined by the various increments of the standard deviation. Under the Gaussian assumption, this computation is easily handled using the Gaussian error function (erf) [20],

$$erf(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt \tag{1.2}$$

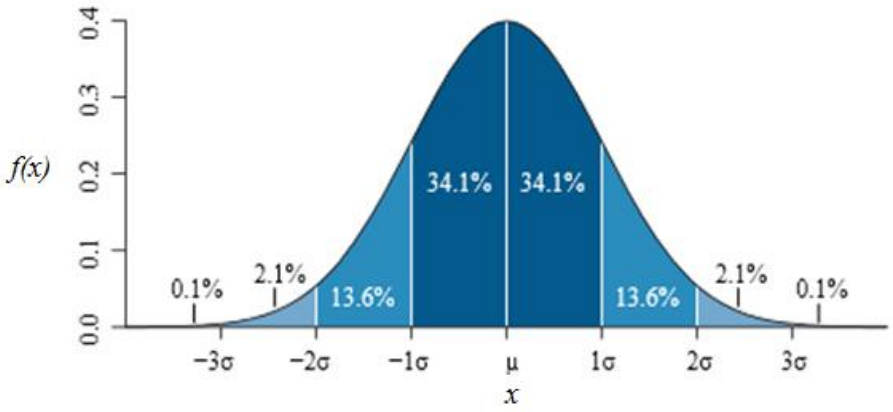

Figure 2: Standard Gaussian Probability Density Function [21]

When evaluated at $x/\sigma\sqrt{2}$, *erf(x)* gives the probability a sample from the Gaussian distribution with a standard deviation of σ is between lies within [μ - *x*, μ+ *x*].

Since most applications of navigation systems rely on Gaussian assumptions, a natural by-product is that the estimation process lends itself well to using a Kalman filter. An inherent output of the Kalman filter estimation process is the computation of a covariance matrix (P), defined as [22]

$$P \triangleq E[(x-\mu)(x-\mu)^T] \tag{1.3}$$

where E[·] is the expected value operator and diagonal terms are the squares of the standard deviation (σ^2). Therefore, to provide a bound on the estimate to a defined probability, one must apply the appropriate scale factor of the standard deviation. Consider Figure 2 and that we want to provide an uncertainty bound on the state estimate μ, such that the probability that estimation error is larger than the bound is less than approximately 31.7%, which will be referred to as P_{PL}. The uncertainty bound or protection level (PL) is computed as

$$PL = \sqrt{2}\sigma \cdot erfinv(1-0.317) = \sigma \tag{1.4}$$

where *erfinv(·)* is the inverse of the error function defined by Equation 1.2.

Using the above approach provides a simple approach for bounding an instance of the estimation error; however, this bound is just for that particular instance or sample draw. To ensure the integrity requirement is met, one must consider possibly many instances or samples across the associated time-frame of the defined integrity risk requirement. For instance, consider an integrity requirement defined across a period of 30 seconds, such as the LAAS horizontal requirement. To determine the risk of an error exceeding a protection level as computed above over that entire 30 seconds, the number

of independent samples, or estimations, that occur across that time-period must be considered. Considering this, the entire probability of an integrity event, still assuming nominal and no-fault conditions, becomes

$$Integrity\ Risk = \ (1 - P_{PL})^N \tag{1.5}$$

Where P_{PL} is the probability that the error exceeds the protection level for any given independent sample or state estimate, and N is the number of independent samples across the entire time span defined by the integrity requirement. Using LAAS as an example, the integrity requirement is defined across a 30 second time period. If we assume the time correlation of the errors is 10 seconds, the number of independent samples would be N = 3. For extremely small probabilities, this can be approximated as

$$Integrity\ Risk = \ N \cdot P_{PL} \tag{1.6}$$

The above methodology to provide integrity is a common approach utilized for navigation applications. This is most likely due to the simplicity in the approach and the fact that it's a natural extension of the Kalman filter, which is widely used for navigation purposes. However, the challenge arises for providing integrity for real-time systems when estimation techniques do not lend themselves well to the Kalman filter construct, a fact we find ourselves in very often with vision applications, a highly non-linear and non-Gaussian process.

1.2 Vision Navigation

Current vision systems are building on the continual advances in imaging sensors such as electro-optical, infrared and a wide variety of other spectral imaging systems to

reach a new level of confidence and reliability. In addition to improvements in the imaging sensors, there is also much work in the area of image processing. A number of techniques and approaches are being developed for processing the data from these sensors for everything from mapping, localization, and other navigation-related applications to medical imaging. Regardless of the specific application, many of these image-analysis approaches share common processing steps to achieve their desired goal. These steps are illustrated in Figure 3 and consist of transformation, correspondence, and state estimation.

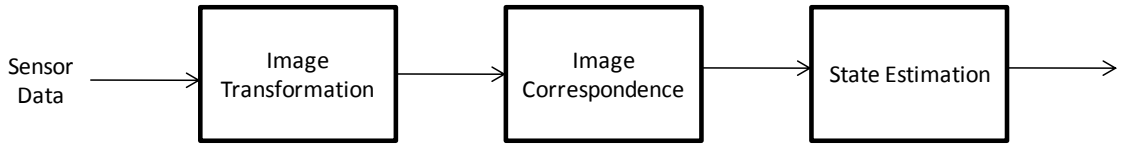

Figure 3: Generic Image Analysis Processing Steps

The first step, Image Transformation describes the process of transforming the raw sensor data to features that have desirable characteristics for one's application. As with most electronics sensors, the raw data from these image sensors often must be transformed to a domain where tractable analytic techniques can be applied to them, whether this data takes the form of light intensity edges from a camera or distance measurements from a laser scanning system. Many techniques exist for this transformation process, techniques such as Scale-Invariant Feature Tracker (SIFT) [23], Speeded Up Robust Features (SURF) [24], edge detection [25][26], corner detection[27][28], eigenspace windowing [29], Hough transformations [30], 3-D shape representations [31] or a host of any other approaches that can be found in the literature, each with different inherent advantages and limitations. The second step, Image

Correspondence, refers to the concept of indentifying and/or tracking specific features found in the data in the transformed domain. Examples of this include searching for and identifying an object within a scene, tracking a SIFT feature between frames or matching images and/or features for stereo vision applications. This research will focus on applying the correspondence to an image rendering approach, where real sensor images and rendered or synthesized images are compared. In this case correspondence is used as a metric of how well the real-image and rendered image relate to each other. Additional details will be provided in subsequent chapters further discussing this approach. Lastly, the State Estimation step describes the process of using this correspondence data to determine state information of the sensor, the world around it, or some relation between the two. This could include extracting navigation information via localization, motion observation, or determining location and pose information for an object of interest.

1.3 Motivation

Much of the vision-aided navigation research to date has focused more on system and algorithmic robustness, rather than quantitative and verifiable integrity, particularly for feature based processing. Along that line, in [5], Veth introduces the concept of regional bounding for feature correspondence between time sequenced image frames, and he includes some feature unique criteria that can provide some protection from feature correspondence errors. Although these approaches do yield some robustness for the algorithms, no quantitative integrity characterization was developed. One exception is work by Larson [14], who introduces a truly quantitative integrity monitor for failures in the mapping of features to pixels, particularly in the presence of a bias. Larson's

approach predicts the largest possible position error in the presence of one such bias due to feature mismatch using a GPS RAIM [32] type approach.

The current state of research addressing integrity for vision navigation, utilizing a image rendering or template matching approach, is even less mature. In fact, no instance of integrity specific work for image rendering vision navigation was identified. This is not surprising; given that most of the research in vision navigation has focused on feasibility studies and/or mission critical type applications where integrity isn't necessarily a driver.

In order to take advantage of the advances in image sensors and vision navigation techniques, particularly for safety critical operations such as close-formation flight, terrain following, or landing systems, rigorous techniques for providing integrity must be developed. For all the work going on in the area of vision navigation, there is an apparent gap in work being done in the area of integrity. However, this is not that unusual, given that most technology advancements in the area of navigation focus on accuracy, availability, and susceptibility to denial. Normally, robustness and safety aspects aren't addressed until systems are more mature. This is particularly true for vision systems, which historically have been used primarily for mission applications that do not have any safety implications, such as intelligence gathering. However, when the system has safety critical implications, as is the case with an unmanned aircraft navigating in close proximity to human operators, such as Automated Aerial Refueling (AAR) [10] or near airport terminal areas, there are very stringent safety requirements, thus requiring much more encompassing considerations of even the smallest integrity aspects of the system.

This research will address the lack of quantitative integrity approaches for vision navigation, specifically relying on the use of image rendering techniques. Given that system integrity includes many aspects, this research will focus strictly on algorithmic aspects under the nominal and no-fault (hardware) condition. Given that specific aspect of integrity, the thrust of the analysis will target the image correspondence process and the resulting relationship with state estimation.

1.4 Research Contributions

The goal of this research is to provide integrity service in a compatible format with existing systems, either in the form of a real-time computed protection level bounds on the relative state estimate, or that of an integrity monitor indicating when the relative state estimation error exceeds the alert limit. The basic definitional form of integrity risk that is adopted by this research is described as follows:

> *The probability the relative navigation position error exceeds the **ALERT LIMIT** without warning, shall not exceed the **INTEGRITY REQUIREMENT.***

Like existing navigation systems, such as GNSS, both the alert limit and integrity requirements are application specific and defined uniquely for that application. To give the research a specific set of real world conditions to evaluate the proposed approaches, the research focuses on the aerial refueling environment, where the operational goal is to maintain a relative position state within the refueling envelope while having a low probability of an undetected safety boundary violation. Although aerial refueling is used as the operation example, the developed integrity approaches should be applicable to any

image based navigation challenge, such as formation flying, landing operations, and even generalized indoor or terrestrial navigation.

The first contribution demonstrates the feasibility of a discrete vision-aided integrity monitor. The developed monitor has a 10^{-3} to 10^{-5} level of integrity monitoring performance during close formation flight operations. The integrity monitor approach generalizes the concept of integrity in terms of operating and alerting regions. Systems that utilize navigation systems generally have objective operating regions that require a certain navigation performance, whether this be around a glide-slope, a formation flight position, or even a flight path clearance. Navigation integrity becomes critical because large divergences from these operating regions, without an alert, can become safety risks. The AL is the instantiation of this concept, in which the AL is the threshold or measure of how much undetected divergence from the operating region can be tolerated without inducing unacceptably large safety risks. This level of independent monitoring could provide tremendous relief to a GPS based precision relative navigation system from a system safety and certification perspective. Additionally, the robustness of the integrity monitor demonstrated against several degrading imaging effects, including lens distortions, and reductions in pixel resolutions

The second contribution is a methodology for determining integrity for a vision systems based on a Bayesian inference approach. Using Bayesian inference for image-based navigation integrity, a configurable integrity risk level, on the order of 10^{-1} to 10^{-6} integrity risk is achieved. Additionally, the research investigates the algorithmic trade-space between safety (i.e., integrity risk) and operational alerts (continuity risk). Various

factors explored within this trade-space include integrity risk threshold, number of measurements, and informative versus non-informative prior information.

This chapter described the underlying motivation for this research and explored some of the historical approaches taken for navigation integrity. Chapter 2 describes the previous research related to vision navigation and explore some of the gaps these approaches have in terms of providing a quantifiable and rigorous integrity assurances. Chapter 3 provides a detailed description of a new discrete integrity monitor for vision navigation systems, including derivation of the underlying approach and various simulated analyses. Chapter 4 includes the second innovation of this research, that of the Bayesian inference approach, and includes a detailed derivation and simulation results. Finally, Chapter 5 provides some additional considerations, based on the research completed, including some future research aspects, and concludes with a summary of the results.

2 Background and Previous Research

In order to further understand the motivation of this research a more in-depth coverage of the previous related research is presented. This Chapter focuses on three main topics: (1) Image Correspondence, (2) Feature-Based Integrity Monitoring for Vision Navigation, and (3) pdf Estimation. Section 2.1 describes many of the most common image correspondence techniques, including both feature and image-rendering approaches. Section 2.2 is a more thorough discussion of Larson's integrity monitoring work. Finally, Section 2.3 introduces the concept of pdf estimation, a concept that will play a important role in the proposed research.

2.1 Image Correspondence

The ability to identify and/or track features within an image or series of images reliably and accurately, referred to as the image correspondence process, is one of the most critical steps in vision navigation. The image correspondence features can have a wide range of forms depending on the application and choice of image transformation.

The imaging correspondence process consists of two main categories or approaches, that of feature-based correspondence and image rendering correspondence. The distinction between these approaches is really driven by the fact that only image rendering or template matching approaches rely on generating synthesized images using *a priori* models. Feature-based correspondence usually deals with a set of somewhat arbitrary image descriptors defined at the pixel level, and are ideally based on image properties that are invariant to scale, orientation, and other image conditions. The feature-based correspondence process is most often handled on a feature-by-feature basis. This is

quite different than the approach taken for image rendering techniques, where the process relies on matching sets of pixels or features that represent whole objects or sections of objects. As the name indicates, the image rendering techniques utilize a priori models to be searched and/or tracked within a given image. Therefore, both the image and the model must be transformed into the domain where the correspondence search is to be performed.

2.1.1 *Feature Based Correspondence Applications*

This section is a brief review of some of the feature based vision navigation applications, touching upon some of the correspondence approaches used. Since this research is focused on the image rendering, a detailed exploration of these feature based techniques is not covered. Rather, a look at the overall approach while highlighting some of the noted challenges associated with the feature based techniques is provided.

2.1.1.1 SIFT Feature Tracking and Inertial Fusion

One of the more common approaches taken in the area of vision navigation is the fusion of vision and inertial sensors. Veth [5] developed an approach for integrating the concept of image-based feature detection and tracking into the classical Extended Kalman Filter (EKF) form. This approach enabled a deeper integration of the two sensors than had been achieved before.

Veth [5] used SIFT to construct a collection of features for a given image, then used the propagated state estimate from the Kalman Filter to predict the feature locations in the next image, a key contribution brought on by his work. The predicted features were then

matched with the actual features once the next image became available at each epoch. Once this matching occurred, the errors between the predicted and actual feature locations could be used in the correction of the navigation state and predicted feature location estimates. Figure 4 illustrates this concept.

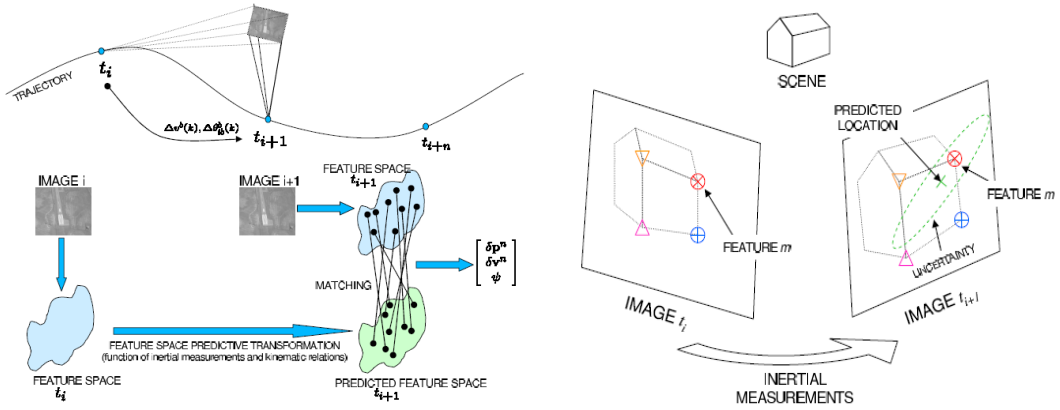

Figure 4: Image-Aided Inertial Navigation [5]

As depicted in Figure 4, the correspondence search for the feature matching can be constrained using the covariance information of the *a priori* Kalman filter estimate. To handle the matching of the features, Veth used the Mahalanobis distance between the predicted and measured features at any given time (t_i), weighting the distances with the feature vector covariance:

$$D_n(t_i+1) = [z_n^*(t_i+1) - \hat{z}_n^*(t_i+1)]^T P_{Z^*Z^*}(t_i+1)[z_n^*(t_i+1) - \hat{z}_n^*(t_i+1)] \quad (2.1)$$

where z_n^* is the feature descriptor vector and $P_{Z^*Z^*}$ is the covariance of the feature vector. The feature descriptor vector is composed of the pixel location (z), scale value (σ), rotation (θ), and the feature descriptor (z_d). Veth noted that the SIFT routine does not

provide any covariance information on the feature vector components for $\mathbf{P}_{\sigma\sigma}$,$\mathbf{P}_{\theta\theta}$, or $\mathbf{P}_{zdzd}$. Hence, the scale and rotation covariance elements were given no weight (i.e., $\mathbf{P}_{\sigma\sigma}$,$\mathbf{P}_{\theta\theta} \rightarrow \infty\mathbf{I}$), whereas the feature descriptor component was given unity uncertainty (i.e., $\mathbf{P}_{zdzd} \rightarrow \mathbf{I}$). However, there was valid uncertainty information available for the location of these features provided by the pose estimation from the Kalman filter, as described above. This pose estimate is used to select the subset of features that fall within a statistical uncertainty bound (e.g., 2σ). Hence a successful match is declared when the distance between the predicted and measured feature descriptor is below a pre-defined threshold. A uniqueness filter is used to ensure that only features that are sufficiently distinct in feature space are used in the state estimation process.

A similar approach is applied to handle the additional correspondence between stereo-image feature pairs [5]. Constraints are applied to restrict the search space between stereo-image pairs to be near the epipolar line. The Mahalanobis distance, Equation 2.1 described above is again used to match image features.

Although Veth [5] applied a rigorous approach for integrating the feature tracking and correspondence matching, including use of covariance information for feature isolation, there is no account for the possibility of matching to an incorrect feature and the subsequent navigation uncertainty of this condition in the state estimation covariance. To account for the uncertainty in the feature matching and resulting uncertainty in the state estimate reflected in the state covariance would need to be made. Specifically, an explicit relationship would need to be made describing how the uniqueness filter criteria, which Veth applies to ensure only sufficiently distinct features are used, influences the

Mahalanobis distance for feature matching, Equation 2.1. Additionally, a probabilistic formulation for determining when feature pairs are declared based on the Mahalanobis distance and the resulting navigation state uncertainty would be required. Without the accounting for these uncertainty considerations in this processing, no rigorous and quantifiable level of integrity can be claimed.

2.1.1.2 Simultaneous Localization and Mapping (SLAM)

Another application that is gaining much attention in the navigation community deals with feature tracking and correspondence is SLAM. The concept of SLAM is to not rely on predefined maps or infrastructure, but to build the reference while simultaneously utilizing this generated information to correct for navigation errors [6].

One of the key aspects in SLAM is the ability to associate real-time imaging data with data stored in the generated map. This process is of critical importance in SLAM algorithms, since false associations can cause errors in both the mapping and localization, thus leading to divergences. In [6][33] a feature set is generated using SIFT and the correspondence is performed using a nearest neighbor search. In the nearest neighbor search process, image features are compared to the n-dimensional database utilizing a variety of different search methods such as k-d tree searches [34] and Best-Bin-First [35] to find the best set of feature vectors. Although some ad-hoc screening techniques are usually applied to ensure that correct matching is achieved, such as simple ratio tests between the nearest and the second nearest match, there is nothing that accounts for mismatches from a sound statistical perspective that can be captured in state estimation

process. This lack of accounting for correspondence uncertainty again leads to the inability to quantitatively characterize any type of integrity performance.

2.1.1.3 RANSAC

The use of RANdom Sample Consensus (RANSAC) has been successfully utilized on image processing and correspondence problems to reject outliers [36][37][38][39], in a sense, providing a certain level of protection against data outliers. The concept of RANSAC is to utilize small random sets of feature data to find an initial set of model parameters consistent with the data. Then the algorithm compares this model to the entire data set looking for a definable level of consistency. If this threshold is achieved, the algorithm then utilizes all of the consistent data points to provide an updated state estimate, while rejecting inconsistent data as outliers. The RANSAC approach relies on three parameters that must be specified by the user: (1) the error tolerance, which dictates whether a data point is compatible with the model, (2) the number of random subset to evaluate, and (3) a threshold that determines the minimum number of compatible data points used to imply the model has been found [36].

Many feature based architectures utilize RANSAC to screen data outliers; however it generally does not provide explicit integrity information. The exception is under very simplistic conditions, where the model is a simple function of the data, it is theoretically possible to utilize the error tolerance parameter to establish analytical bounds [36]. However, this is rarely the case and no explicit relationship is possible, requiring tolerances be established through experimentation, thus losing any quantifiable integrity relationship.

2.1.2 *Template Matching and Image Rendering Correspondence*

This section is a review of some of the various approaches used for image correspondence utilizing image rendering or template matching. Image rendering correspondence is a different paradigm from that of feature based approaches, where the correspondence occurs at the object level, either in 2-D or 3-D, versus individual features. There is a breadth of information, largely from the machine vision community, regarding image rendering correspondence utilizing a number of different techniques, as described in this section. Similarly to the feature based approaches, many of the correspondence approaches rely on distance metrics with somewhat arbitrary thresholds used for matching criteria.

The template matching or image rendering image correspondence approach is based upon using either 2-D or 3-D object models or templates to perform object detection or pose state estimates of that object relative to the imaging sensor. In the case of image rendering, the relative state (position and/or orientation) is determined by perturbing the relative position state and generating rendered or pseudo-images based on a model of the known object and sensor models until a maximum correspondence is achieved. This process is described in Figure 5. Likewise with template matching, 2-D projections or pseudo-images of the object are generated at different poses using known sensor model parameters or captured using the actual sensor and stored in a database. Then object detection and or relative state estimation is performed by finding the best matching database image. This process often relies on some sort of image transformation process, such as edge detection, to convert the image to a more environmentally invariant form

prior to the correspondence process. The navigation estimation process is simply determining the state estimate based on the results of the image correspondence. In the simplest case, the relative navigation state estimate is the rendered pose that yields the best correspondence with the actual sensor image.

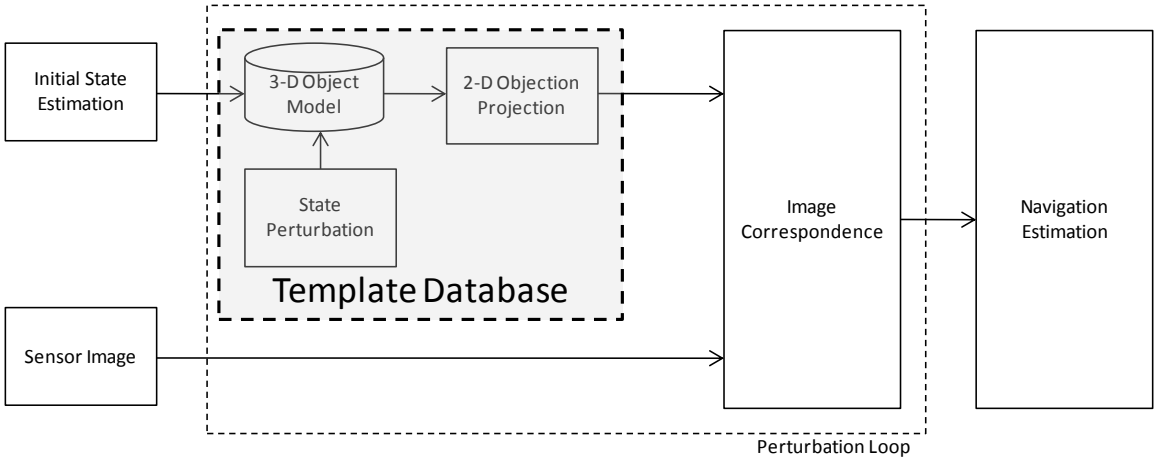

Figure 5: Image Rendering Vision Navigation Processing Overview

As indicated, there are several varieties of image correspondence techniques, each with their unique formation of a conformity or similarity metric. The subsequent section focuses on some of the most commonly used correspondence approaches, particularly for image processing.

2.1.2.1 Sum Squared Difference

Weaver [10] evaluated a sample of techniques for 2-D image correspondence using object rendering image matching. The first technique considered was the Sum-Squared Difference (SSD), in which the normalized grayscale sensor image (I) and the rendered normalized grayscale image using the model (N) were directly compared at the pixel level, across all n pixels, using

$$SSD = \sum_{i=1}^{n} (I_i - N_i)^2 \tag{2.2}$$

The general concept is by perturbing the model by its 6-DOF axes and comparing it to the image, a minimization of Equation 2.2 occurs at the best estimated pose of the object. To improve performance by eliminating aliasing effects, Weaver filtered the images through a Gaussian blur (mean = 0, σ = 1 pixel) to smooth the edges of the object prior to performing the SSD. The results of this approach, evaluated in a simulated environment based on a KC-135 aircraft can be seen in Figure 6 and Figure 7.

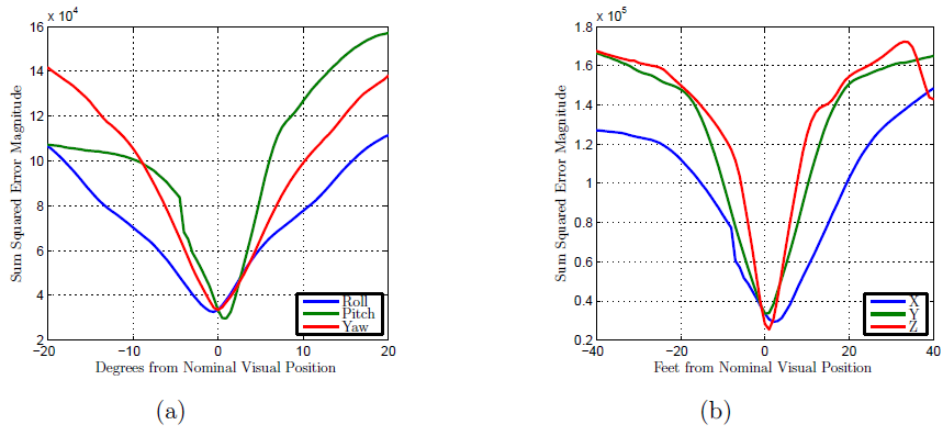

Figure 6: Sum-Squared Difference Example of Independent 6-DOF Perturbations [10]

In Figure 6, SSD image correspondence between a sensor image and a series of rendered images is shown. The rendered images are generated by using the *a priori* model and perturbing the relative states, position and attitude one state at a time, for observability purposes. The concept demonstrates that a minimization is achieved near the nominal or true state. However, as can be seen in Figure 7, the global minimum becomes much less observable when two axes are perturbed simultaneously, in this case

the vertical position and pitch. This problem of observability is guaranteed to only become worse when more degrees-of-freedom are considered. Based on this fact, this research will focus solely on relative position state estimation aspects, assuming that attitude information can be accurately determined through other approaches (such as inertial navigation systems).

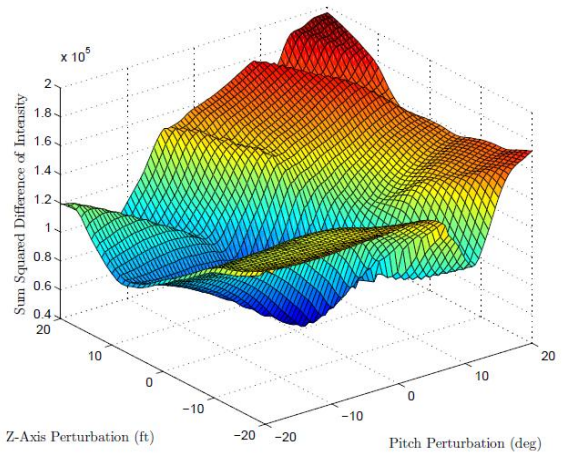

Figure 7: Sum-Squared Difference with Simultaneous Vertical and Pitch Perturbations [10]

2.1.2.2 Magnitude of Gradient

Weaver [10] also evaluated the Magnitude of Gradient technique, in which the image and the model were preprocessed through a Prewitt Filter [40] to determine changes in image intensities between adjacent pixels. The purpose of this approach is to highlight the object edges or corners where large gradients are likely to occur. An example of this approach within a simulated environment is shown in Figure 8 andFigure 9.

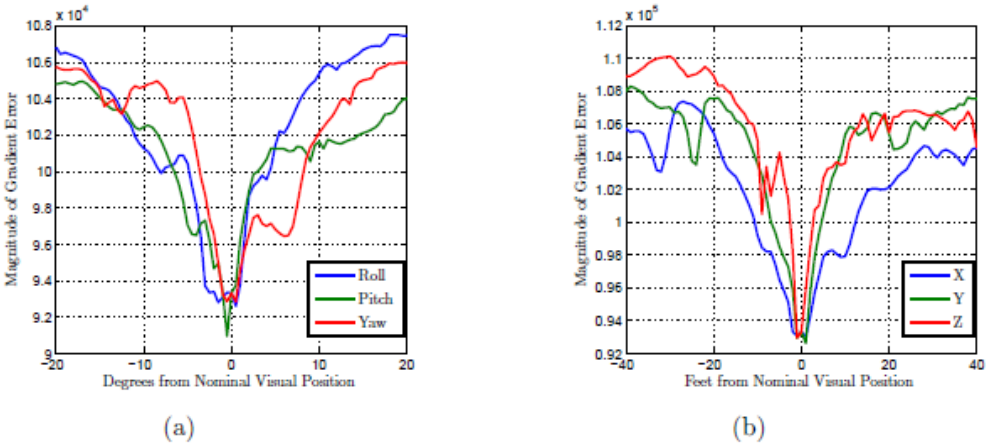

Figure 8: Magnitude of Gradient Example of Independent 6-DOF Perturbations [10]

Like the Sum Squared Difference approach, this technique has merit. In fact, the minimization curves in Figure 8 show a steeper slope about the minimization, a highly desirable attribute when considering aspects of accuracy and uniqueness. However, this approach also shows similar susceptibility to observability problems when multiple degrees-of-freedom are considered, as seen in Figure 9, an issue that might be unavoidable under any vision based concept.

Based on the results documented by Weaver, it is clear that he successfully demonstrated the feasibility of performing pose estimation using these image rendering, 2-D image processing techniques.

Weaver's [10] work utilizing image rendering did not include any relationship between the various minimization techniques to that of accuracy or more importantly for this research, that of integrity. However, we will see in Chapter 2.4 of this dissertation, that his approach has laid the ground work for utilizing these approaches to develop a rigorous integrity concept for image rendering vision navigation.

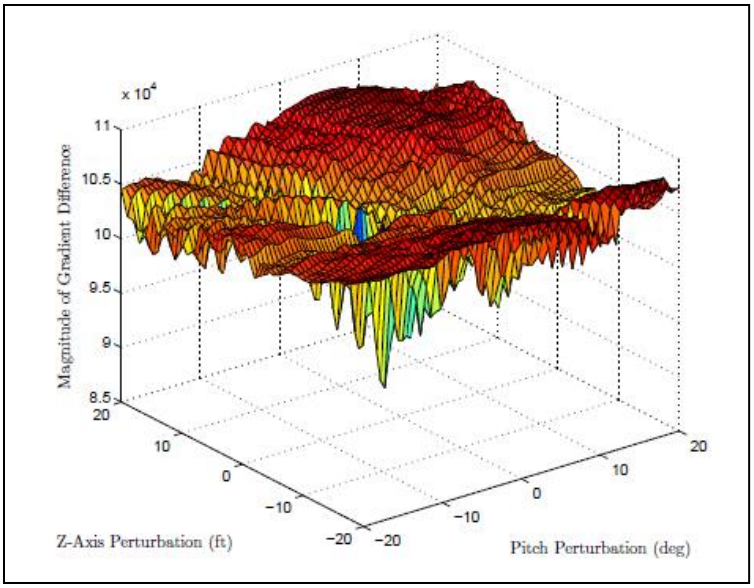

Figure 9: Magnitude of Gradient with Simultaneous Vertical and Pitch Perturbations [10]

2.1.2.3 Minkowski Distance

The Minkowski Distance [41][42] approach, also known as the L_p distance, defined by Equation 2.3, is a set of functions that provides a specific metric for distance between two points or set of points (x_i,y_i), such as an image with N pixels.

$$L_p(x,y) = \left[\sum_{i=0}^{N} |x_i - y_i|^p\right]^{1/p} \tag{2.3}$$

There are a few commonly used versions of the Minkowski Distance. The Manhattan or City-Block distance refers to the L_1 case, when $p = 1$, expressed as:

$$L_1(x,y) = \sum_{i=0}^{N} |x_i - y_i| \tag{2.4}$$

The case when $p = 2$ or L_2 is simply the Euclidean Distance between two points.

$$L_2(x,y) = \left[\sum_{i=0}^{N} |x_i - y_i|^2\right]^{1/2} \tag{2.5}$$

Lastly, the case when $p = \infty$ is called the Chebyshev Distance, and is simply the max distance between the set of two points,

$$L_\infty(x,y) = \max_{0 \le i \le k} |x_i - y_i| \tag{2.6}$$

The Minkowski Distance is a very simplistic matching routine and usually used in conjunction with another metric, such as the Hausdorff Distance (described in Section 2.1.2.6). Like many of the discussed image correspondence approaches, on its own the Minkowski Distance has no metric to describe the proximity of the match or to account for any uncertainty.

2.1.2.4 Eigenspace Matching

As mentioned earlier, the process of image rendering image matching is sometimes approached in different domains, with the image being processed through some type of transformation. In [29] [43][44], the approach is based on transforming the images to eigenspace and comparing it to a set of training images for object detection and pose estimation.

In [29][43], a series of training images are used to generate a database based on the object desired for identification and pose estimation. The training images consist of varying the object over a range of poses and taking snapshot images. Each of these images are then scanned in raster order to form a large matrix array *Z*, where each column

is a training image. The mean of all the training images, c, is then subtracted from each of the images,

$$\bar{Z} = [z_1 - c, z_2 - c, \dots, z_M - c] \tag{2.7}$$

The training image sample covariance matrix is then computed as

$$Q = \bar{Z}\bar{Z}^T \tag{2.8}$$

This covariance matrix provides the series of eigenvalues λ_i and eigenvectors e_i, where

$$\lambda_i e_i = Q e_i \tag{2.9}$$

The input image is then processed in the same manner to generate the eigenvalues and eigenvectors. The pose estimate is calculated by finding the best fit in eigenspace, using a distance minimization processes such as SSD, L1-norm, or Mahalanobis distance, etc. Both approaches use some form of interpolation for estimating the pose between training images.

2.1.2.5 Hough Transformation

Another technique of image matching that involves the transformation of the imaging to a different domain, is that of localization using Hough transformations [44][45]. The Hough transformation was originally applied for line detection and improving edge-detection algorithms in the presence of noise and segment breaks [45]. However, it has since been expanded to a robust technique for detecting other shapes in an image, such circles and ellipses.

The basic concept of the Hough Transformation is to transform data points from the normal (x-y) Cartesian plane to the (ρ,θ) plane of the Hough domain. This transformation is governed by the following relationship,

$$\rho = xcos\theta + ysin\theta \tag{2.10}$$

A simple example of this transformation is illustrated in Figure 10. As expected by Equation 2.10, each point in the x-y plane is subsequently represented as a sinusoidal-like line in the Hough plane, where a line can be identified as the intersection of the curves in the Hough plane. Under normal conditions, where noise is present and can affect this process, the identification of geometric patterns, such as lines, is performed by an accumulation of points, with local peaks representing the best estimate of these patterns.

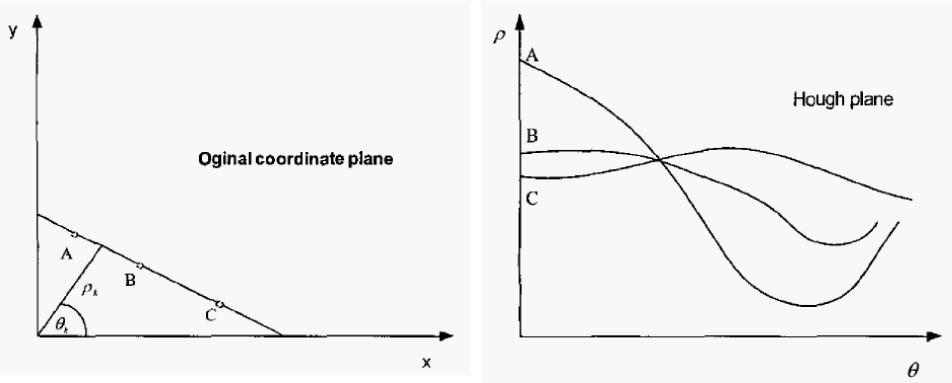

Figure 10: Hough Transformation Example [45]

An example of image correspondence using this approach is presented in [45], in which a mobile robot is being localized by transforming range measurements to the Hough domain and comparing to a predetermined map, also being represented in the Hough domain. The correspondence problem is solved by restricting the search area to a threshold (δ), otherwise expressed as

$$(\theta^M - \theta^S)^2 + (\rho^M - \rho^S)^2 < \delta \tag{2.11}$$

where (θ^M,ρ^M) represent the reference points based on the map and (θ^S,ρ^S) are the data points based on the sensor range measurements. The difference between the map and sensor data points provides information on the motion of the mobile robot. Due to the properties of the Hough transformation, the translational and rotation estimation processes can be decoupled, which is a convenient property. An example of this process is illustrated in Figure 11 for a mobile robot facing a corner. The solid line (a,b) on the left figure represents the map model and the dotted line (a',b') represents the data from the sensor, all in Cartesian space. The same data is shown in the right figure represented in the Hough domain where the local maxima, representing line segments, are shown. Here the map model is represented as circles and the sensor data is represented as crosses. Orientation and position estimates are illustrated by the deltas between the model and sensors, represented as such.

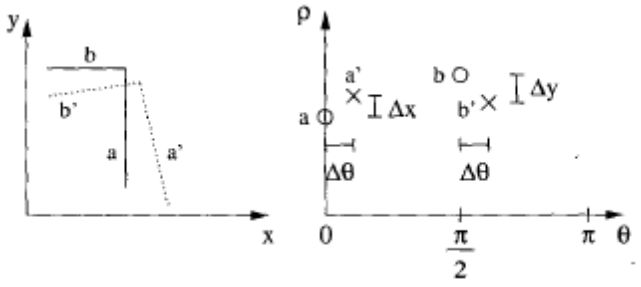

Figure 11: 2-D Hough Localization Example [45]

The approach of transforming image data to the Hough domain begins to mimic the feature based tracking and correspondence approaches described in Section 2.1.1. Where motion estimation and localization problems are solved using discrete point

representations of simple geometric shapes. In addition, this approach appears to use similar non-rigorous similarity metrics as part of the correspondence solution, the δ threshold being an example of this. There is little mention of the derivation and the characteristics this sort of threshold approach has on the statistical uncertainty in the state estimation process, a common issue encountered as part of image correspondence applications.

2.1.2.6 Hausdorff Distance

The Hausdorff Distance approach for image matching [41][46][47][48][49] relies on comparing two entire image sets, potentially of different sizes. To compute the Hausdorff Distance (HD), consider two data sets, *X* and *Y*, that represent two images

$$HD(X,Y) = \max\{\vec{h}(X,Y), \vec{h}(Y,X)\} \tag{2.12}$$

where

$$\vec{h}(X,Y) = max_{x\in X}\, min_{y\in Y}\, d(x,y) \tag{2.12}$$

d(x,y) is any variety of distance measure (e.g. Minkowski distance).

The standard Hausdorff distance scheme calculates the measure or degree of similarity between two images, where one image might be an object template. For the case when detection is used, it effectively is the max-min distance of all the edge points between the template and the image. However, this approach does have some major drawbacks. First, since it requires computing a distance metric between every two points in the image pair, its execution comes at a high computational cost. Secondly, it is

susceptible to a single outlier in even only one of the image pairs. This can become problematic given that most image matching must deal with occlusion and noise. To address these issues, several variants of HD have been developed. This includes Partial Hausdorff Distance (PHD), Hausdorff fraction, and several other variants.

The Partial Hausdorff Distance was introduced in [50],

$$h_K(A,B) = {K^{th}}_{a\in A} min_{b\in B}\, d(x,y) \tag{2.14}$$

where only the *K* best points in *A* that match points in *B* are considered and the distance of points in *B* to *A* are no longer considered. The Hausdorff fraction is performed by determining the number of points in *A* that are within a defined distance (δ), to *B*:

$$h_K(A,B) \leq \delta \tag{2.15}$$

So suppose there are K_δ points in A for which (2.15) is true, hence the ratio of K_δ to the total number of points in A is considered the Hausdorff fraction, or

$$F_\delta(A,B) = \frac{K_\delta}{|A|} \tag{2.16}$$

One of the major drawbacks with the classical Hausdorff distance and other similar distance metrics approaches is that they usually rely on some sort of threshold to determine a match, thus leading to a sharp distinction between matched and unmatched objects. Additionally, there are no probabilistic attributes associated with them, thus preventing the use of prior knowledge and extracting uncertainty information from the process. Again, this lack of uncertainty in this aspect of the correspondence process fails to provide any information for determining navigation integrity.

2.1.2.7 Statistical Approaches

All the previously described techniques for image correspondence have relied on various forms of distance functions to determine matches, none of which have incorporated any sort of statistical probabilities that could be used for integrity determination. The following section describes work that begins to lay the foundation for integrity by incorporating statistical estimation techniques into the image correspondence process.

Wells [51] applies Maximum Likelihood (ML) and Maximum A-Posteriori (MAP) estimation concepts to object detection and pose estimation problem. In order to set the stage for these statistical approaches, some background on the ML and MAP techniques is warranted.

The general solution for the ML approach is given by

$$\hat{x}_{ML} = \max_{x} p(d; x) \tag{2.17}$$

where $p(d; x)$ is the probability density function (pdf) of the observed data (d), given a fixed, but unknown parameter (x). Since the pdf, $p(d; x)$, is viewed as a function of the unknown parameter, it is referred to as a likelihood function. Hence the ML solution is simply the estimate of x that maximizes the likelihood function. One of the key attributes of the maximum likelihood approach is the fact that the unknown parameter, x, is not a random variable, but rather an unknown parameter that influences the underlying pdf of the observed data.

The MAP approach on the other hand is based on the condition when the

unknown parameter (x) to be estimated is random. Hence the solution incorporates additional uncertainty associated with the prior probability of x itself. This is mechanized by utilizing Bayes law part of the maximization

$$\hat{x}_{MAP} = \max_{x} p(\mathrm{x}|\mathrm{d}) = \frac{p(\mathrm{d}|\mathrm{x})p(\mathrm{x})}{p(d)} \tag{2.18}$$

where now knowledge of the *prior* probabilities $p(\mathrm{x})$ and $p(d)$ are required due to the randomness of x. The name comes from the reference to $p(\mathrm{x}|\mathrm{d})$, referred to as the *a posteriori* pdf. This technique is often referred to as Bayesian inference, meaning we are gaining information on the probability distribution of the state (x), using a series of observations (d) and prior knowledge of the state itself. It should also be noted, that much like the ML approach, with MAP the likelihood function plays a central role in the process.

To understand Wells use of these concepts for the object identification and pose challenge [51], consider the following condition described by a set of images (Y) and model (M) features:

$$\begin{gathered} Y = \{Y_1, Y_1, \dots, Y_n\} \\ M = \{M_1, M_1, \dots, M_m\} \end{gathered} \tag{2.19}$$

Wells defines the variable, Γ, that maps the correspondence between these image and model features. Further, a probabilistic model is developed for these correspondences, which we will see, will act as the prior probability in the estimation

process. All image correspondences are assigned one of two levels of uniform probability, associated with either image points that map to model features or background clutter in the image

$$p(\Gamma_i) = \begin{cases} B_i & if\ \Gamma_i \text{is background } (\perp) \\ \dfrac{1 - B_i}{m} & otherwise\ maps\ to\ a\ model\ feature \end{cases} \tag{2.20}$$

Wells makes the assumption that all image correspondences are independent and combines them using the joint probability mass function,

$$p(\Gamma) = \prod_{i=1}^{n} p(\Gamma_i) \tag{2.21}$$

The prior probability for the pose (β) is developed assuming a normal density function,

$$p(\beta) = G_{\psi_\beta}(\beta - \beta_0) \tag{2.22}$$

where $G_{\psi_\beta}(x)$ is defined as

$$G_{\psi_\beta}(x) = (2\pi)^{-\frac{z}{2}} \left|\psi_\beta\right|^{-\frac{1}{2}} \mathrm{e}^{\left(-\frac{1}{2}x^T {\psi_\beta}^{-1} x\right)} \tag{2.23}$$

Here *z* represents the dimensionality of the pose vector and ψ_β is the covariance matrix of the pose prior. Wells notes that if little is known about the pose prior then it can be disregarded and a ML approach can be utilized, something that will be discussed later in this section.

The last building block is the probability models for the image features, conditioned upon the correspondences (Γ) and the object pose (β). Again, Wells takes the

approach of developing two distinct models for image features that map to background clutter and another for features that map to the model, as shown in Equations 2.24-2.26. For the background features another uniform distribution is used, where *W* is dependent upon the image pixel size,

$$p(\mathrm{Y_i}|\,\Gamma,\,\beta\,) = \frac{1}{W_1 \ldots W_v} \;\; if \;\; \Gamma_{\mathrm{i}} = \perp \tag{2.24}$$

For image features that map to the model, similar to the priors, a normal distribution is utilized,

$$p(\mathrm{Y_i}|\,\Gamma,\,\beta\,) = G_{\psi_{ij}}\left(Y_i - \; f\left(M_j, \beta\right)\right) \; if \;\; \Gamma_{\mathrm{i}} = \mathrm{M_j} \tag{2.25}$$

where

$$G_{\psi_{ij}}(x) \equiv (2\pi)^{-\frac{v}{2}} \left|\psi_{ij}\right|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} x^T {\psi_{ij}}^{-1} x\right) \tag{2.26}$$

Once again, independence is assumed and yielding the following expression for the joint probability,

$$p(\mathrm{Y}|\,\Gamma,\,\beta\,) = \prod_i p(\mathrm{Y_i}|\,\Gamma,\,\beta\,) \tag{2.27}$$

With this, all the pieces have been defined to complete the MAP estimation approach. Using Bayesian inference, the probability density function for the correspondences (Γ) and the object pose (β) conditioned upon the realized image features can be expressed as

$$p(\Gamma, \beta|\mathrm{Y}) = \frac{p(\mathrm{Y}|\Gamma, \beta)p(\Gamma, \beta)}{p(Y)} \tag{2.28}$$

Where the MAP estimator is defined as the correspondences (Γ) and the object pose (β) that maximizes this *a posteriori* probability distribution, expressed as

$$\widehat{\Gamma, \beta} = \arg\max_{V} p(\Gamma, \beta|\mathrm{Y}) \tag{2.29}$$

Further, the solution to Equation 2.29 can be narrowed to just the pose estimate (β) by integrating across the range of correspondences to obtain the marginal pdf,

$$p(\beta|\mathrm{Y}) = \sum_{\Gamma} p(\Gamma, \beta|\mathrm{Y}) \tag{2.30}$$

Wells demonstrated these techniques in a laboratory environment for a single object detection and pose estimation problem. However, no clear conclusions are made as to the resulting performance or effectiveness [51].

Following from Wells' work, in [50][52][53] Olson uses a similar approach from the aspect of solving for self-localization, where the images are transformed into binary format or edge maps such that each pixel or feature is then either a 0 or 1. This yields two feature sets, the model features, $M = \{\mu_1, \ldots, \mu_m\}$ and a set of image features, $I = \{\upsilon_1, \ldots, \upsilon_n\}$. Although not explicitly required, it is assumed that exactly one instance of the model appears in the image. The model position within the image is described by the random variable *t*.

Oslon [50][52][53] similarly formulated the likelihood function for *t* as the product of the distance probabilities, $D_1, \ldots, D_m$,

$$L(t) = \prod_{i=1}^{m} p(D_i; t) \tag{2.31}$$

where the distances (D) are defined as the measure between each model pixel to the closest image pixel assuming the model position t, and $p(D_i; t)$ is the probability density function. Equation 2.31, is then put into logarithmic form to maintain model position ordering and improve computational efficiency:

$$\ln L(t) = \sum_{i=1}^{m} \ln p(D_i; t) \tag{2.32}$$

A probability density function (pdf) of the form in Equation 2.33 is used for $p(D_i; t)$ as the matching criterion; however, Olson mentioned that any pdf can be used.

$$\ln p(D_i; t) = \begin{cases} k_1 + k_2, & \text{if } D_i \leq \delta \\ k_1, & \text{otherwise} \end{cases} \tag{2.33}$$

This approach yields that for a particular model position (t), if the model feature lies within δ of an image feature, it will result in a uniform distribution or constant probability. Otherwise a smaller, but still constant probability is assigned. The precise values of these constants are unimportant as long as $k_2 > 0$, with $k_1 = 0$ and $k_2 = 1$ being generally used. The fact that this approach results in a two-valued pdf based on if the distance metric (D_i) is within a threshold (δ), means that it essentially becomes equivalent to the Hausdorff approach of Equation 2.13 [50].

Olson [50][52][53] continued by considering a normal distribution for this approach, thus replacing Equation 2.33 with the following pdf

$$p(D_i; t) = k_1 + \frac{k_2}{\sigma\sqrt{2\pi}} e^{-D_i^2/2\sigma^2} \tag{2.34}$$

where σ is the standard deviation of the feature uncertainty. The constant term k_1 is added to provide a lower bound on the natural log form, thus preventing $\ln p(D_i; t)$ from approaching an arbitrary small number when outliers are encountered. Thus k_1 and k_2 are determined by the frequency of the existence of outliers. Olson made note of the fact that Equation 2.34 is not a true pdf since it does not integrate to unity. However, he also made the statement that this is unavoidable if we are to account for outliers when the distance D_i becomes large and doesn't affect the accuracy of these results in a significant way.

The use of the normal distribution allows for varying levels of uncertainty in the image features or its position. Another important aspect of the Olson's work is that he chose to base his maximum likelihood matching on summing about the model features versus the image features. Olson claimed this could avoid matching on noise or background features that are irrelevant to the object being tracked [52].

Another benefit to this probabilistic approach is the ability to take advantage of prior probabilities [50][52]. In this case, Equation 2.32 can be expanded to the following, where ln *p(t)* is the prior probability of the model position:

$$\ln L(t) = \ln p(t) + \sum_{i=1}^{m} \ln p(D_i; t) \tag{2.35}$$

The position of the model, *t*, that maximizes this expression or has the maximum likelihood, is chosen. Olson utilized a search strategy that divided the pose space into

rectilinear cells, and then recursively searched and pruned these cells to locate the maximum likelihood peak. In order to improve upon the resolution of this approach, Olson presented an approach for achieving sub-pixel localization [54].

As expected, one of the major advantages to Olson's approach is the ability to estimate the probability of failure. This is estimated by considering the ratio of the sum of the area around the peak which maximized Equation 2.35 and the sum of the remaining pose space.

This process was evaluated in a simulated environment with a series of over 100,000 trials [52]. The results, shown in Figure 12, demonstrate that the statistical predictions on the localization errors match extremely well with the predicted performance,

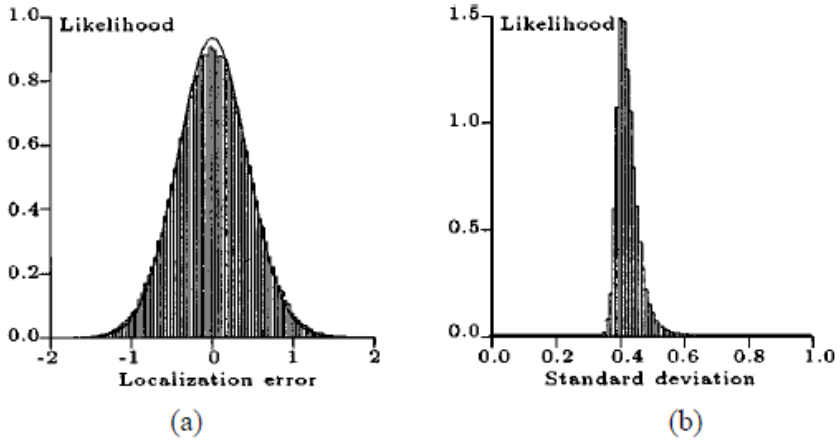

Figure 12: Simulation Results for Maximum Likelihood Localization [50]

In general, Well's and Olson's approach demonstrates the successful incorporation of statistical techniques, such as ML and MAP, to the challenge of object detection and relative pose estimation using a model based image correspondence techniques.

However, their research was focused exclusively on object detection and state estimation, not integrity. For the purposes of this research, limiting the scope to simply the best estimate will not suffice. As will be discussed in Chapter 2.4, this research proposes extending the approach of Wells and Olson, and use Bayesian inference to perform full pdf inference of the state.

2.2 Image Navigation Integrity Monitoring for Feature Correspondence

Larson [14] acknowledged this gap in research in the integrity aspects of feature-based vision navigation by looking to the concept of GPS RAIM [32] and extending it to vision systems. Specifically, Larson looks at the impact of a failure to correctly map a feature to the correct image pixel and how a failure of this kind can affect the navigation state in the position domain. This appears to be the first set of research to address the integrity aspects of vision navigation, specifically related to failure in the feature correspondence process.

To further illustrate this type of failure and its relationship to integrity, consider Figure 13. This figure show the probability that the magnitude of test statics (referred to as the parity vector) exceeds the detection threshold under both fault-free and biased conditions, notated as H_0 and H_1 respectively. The real area of concern in this scenario is a case where an undetected bias exists, but the magnitude of the parity vector is still determined to be below the threshold. The probability of this condition is represented as the red shaded area. The other aspect of this problem is a condition where no bias exist, fault-free, however the magnitude of the parity vector is above the threshold, a condition that would lead to a false-alarm. The probability of this condition is represented as the

blue shaded area. Both scenarios are of critical importance, setting the balance between integrity and continuity of the system.

Similar to the GPS RAIM approach, Larson developed an approach that based on the specific geometry of the system, would determine the largest possible horizontal error given the existence of a single bias in one of the pixel location measurements. Larson used the term "slope" to describe the linear approximation of the relationship between the bias value and the resulting horizontal error.

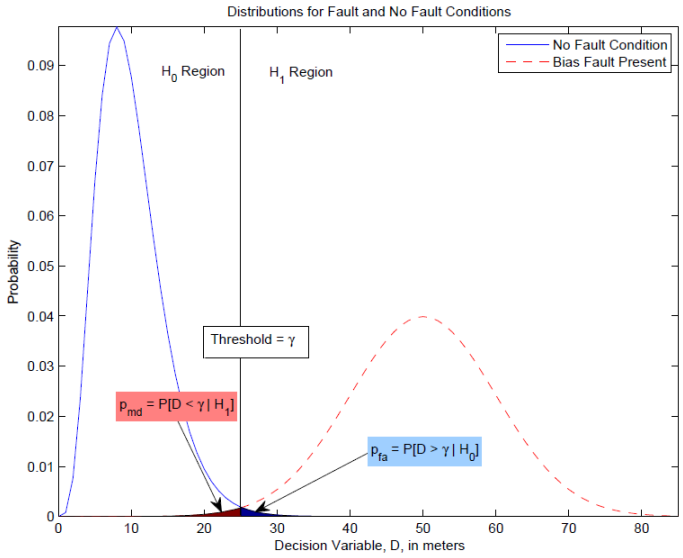

Figure 13: Fault Detection Overview under Fault-Free and Biased Conditions [14]

Figure 14 is an example of this approach based on simulated data. The plot on the left represents the image plane with four features being indentified and tracked. Figure 14 shows the resulting relationships between the parity vector test statistic and the resulting horizontal error.

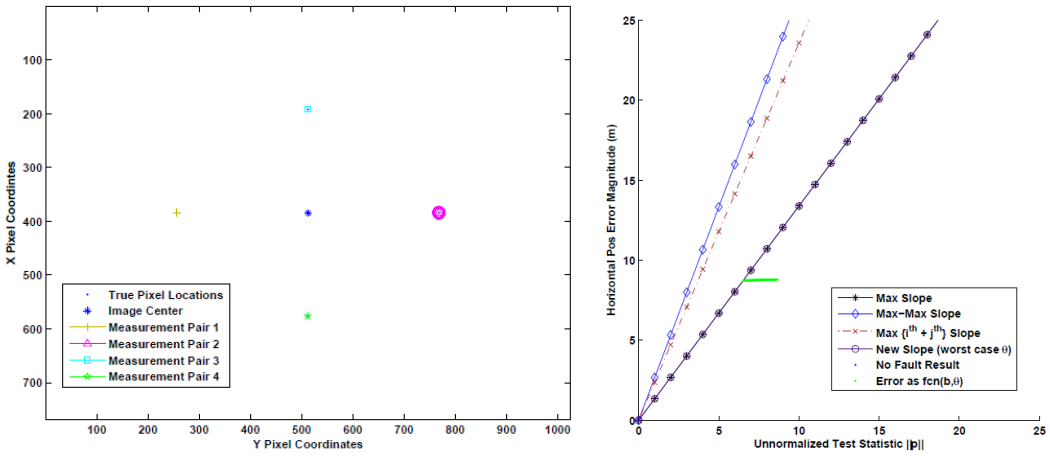

Figure 14: Example of "Slope" Method for Feature Based Integrity Monitoring [14]

The results Larson showed indicated that the extension of the GPS RAIM concepts can be made to that of feature based pixel correspondence errors. However, the specific definition of a faulted-condition, that being a feature pixel error, is only one aspect of the image analysis processing and doesn't address the issue of image or feature correspondence and the resulting pose uncertainty.

2.3 Probability Density Function Estimation

Given that the goal of this research is to develop an integrity methodology for vision relative navigation. This research revolves around the need for rigorous pdf estimation. Therefore, it is appropriate to discuss some existing techniques for pdf estimation. The basic concept of pdf estimation is based upon estimating the governing probability distribution based on a set of realized samples from that distribution. A pdf estimator can be broken into two classical categories — parametric and non-parametric. The defining distinction is that parametric approaches rely on a predefined structural form for the pdf, such as Gaussian or Poisson, turning the challenge into solving for the

associated parameters that define those distributions. In contrast, non-parametric approaches make no presumption on the functional form and allows the data to drive the result.

Given that there is no current basis to assume structural form of the pdf for vision navigation application, this research focuses on non-parametric estimation techniques. If execution of the proposed research yields results that trend towards more classically known probability distributions, this choice can be revisited and parametric estimation techniques could be explored, but this is not expected. The remaining sections are a brief overview of some of the most commonly used non-parametric pdf estimation techniques, along with an analysis of the advantages and disadvantages of such approaches.

2.3.1.1 Histogram Method

The oldest and most commonly used approach for pdf estimation is the histogram method. The histogram method is based on the simple approach of binning observation data, X_i, into either uniformly or varying bin widths and normalizing to ensure unity area. These approaches are defined mathematically for uniform bin widths by

$$\hat{f}(x) = \frac{1}{nh}(no.\,of\ X_i\ in\ same\ bin\ as\ x), \tag{2.36}$$

where h is the bin width and n is the number of observations [55]. For varying bin widths, this takes the form of

$$\hat{f}(x) = \frac{1}{n} \cdot \frac{(no.\,of\ X_i\ in\ same\ bin\ as\ x)}{h_x}. \tag{2.37}$$

where h_x is now the varying bin width containing x.

The histogram method is a straight forward approach for pdf estimation, and very practical for early presentation and exploration efforts. However, the histogram method does have some drawbacks. For example, the histogram can be sensitive to the selection of the origin that defines the binning. As a result of this effect, masking of multi-modal distributions can occur. Additionally, by its nature the histogram has discontinuities between bins, thus causing issues when pdf derivatives and other higher order mathematical processing is required. Figure 15 is an example implementation of the histogram methodology, showing both the true pdf, from which 1000 samples are drawn, and the resulting estimated pdf utilizing 100 bins.

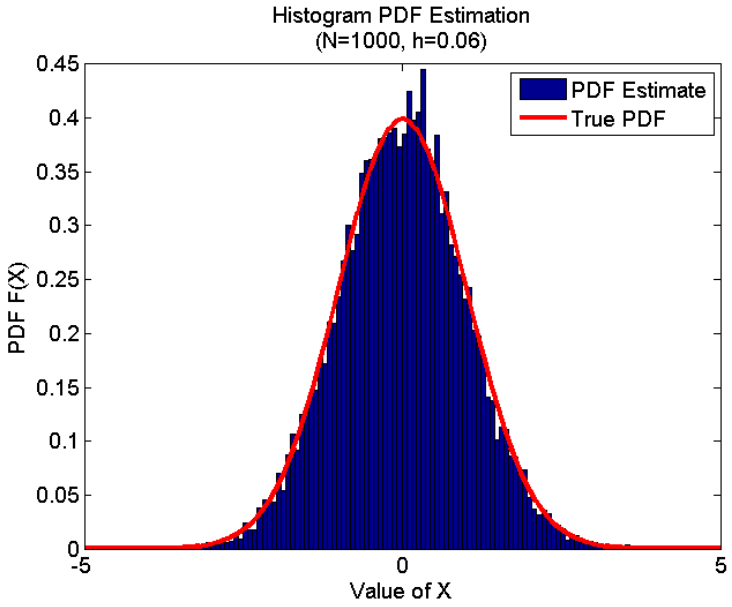

Figure 15: Histogram pdf Estimation Example

The results demonstrate that indeed the estimate pdf mimics the trend of the underlying distribution, although some of the fundamental issues associated with the histogram method are also present, such as discontinuities between bins.

2.3.1.2 Kernel Method

Another approach to pdf estimation is the kernel method, which relies on a predefined kernel function that acts as a weighting function in the construction of the pdf [55]. The Kernel method is really a class of techniques taking on the following functional form

$$\hat{f}(x) = \frac{1}{nh}\sum_{i=1}^{n} K\left(\frac{x - x_i}{h}\right) \tag{2.38}$$

where x_i are the n independent samples drawn from the distribution, h is the window width, also known as the smoothing parameter, and $K(x)$ is a kernel or weighting function satisfying the following condition

$$\int_{-\infty}^{\infty} K(x)dx = 1 \tag{2.39}$$

One common implementation referred to as the naïve method, uses the weighting function defined as

$$K(x) = \begin{cases} 1/2 & if\ |x| < 1 \\ 0 & otherwise \end{cases} \tag{2.40}$$

The normal distribution is also often used, and due to its smoothness it yields very nice differentiability properties.

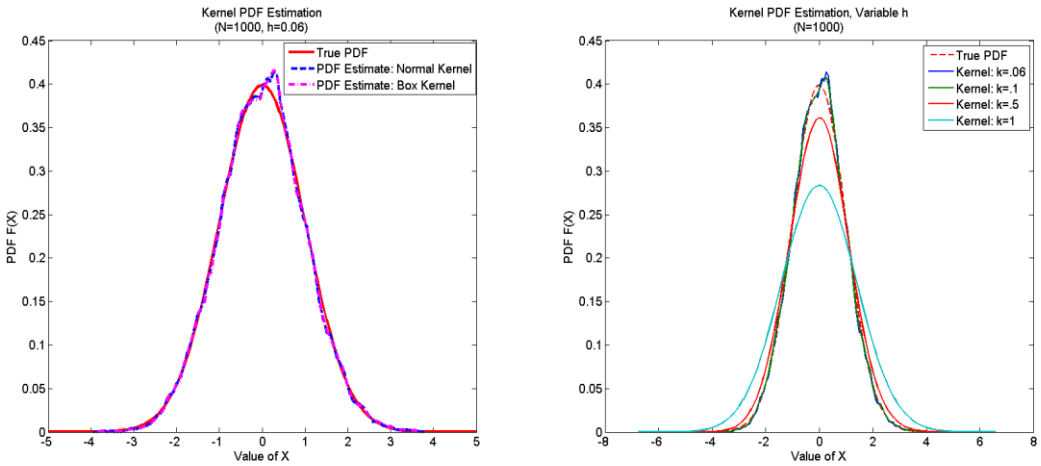

Figure 16: Kernel pdf Estimation Example

Figure 16 is an example implementation of the kernel methodology using two different kernels, a naïve and normal kernel. Additionally, the figure illustrates the smoothing effect that the window parameter, *h*, has on the resulting estimation using a normal kernel.

2.3.1.3 Nearest Neighbor Method

The nearest neighbor or K-nearest neighbor (KNN) method is based on an adaptive smoothing about the "local" density of the data [55]. A simple derivation of this approach can be developed by considering the characteristics of a given a sample size of *n* for a specific pdf *f(x)*. Over an interval of [x - r, x + r], one would expect approximately $2 \cdot r \cdot n \cdot f(x)$ samples to fall within that interval, for all $r > 0$. Now consider the distance, $d_n(x)$, between *t* and all the samples sorted in ascending order,

$$d_1(x) \leq d_1(x) \leq \ldots \leq d_n(x) \tag{2.41}$$

By definition ($k - 1$) observations would fall within the interval of [$t - d_k$(x), $t + d_k$(x)], and using the above approximation yields

$$(k-1) = 2nd_k(t)\hat{f}(x) \tag{2.42}$$

and solving for $\hat{f}(x)$

$$\hat{f}(x) = \frac{(k-1)}{2nd_k(x)} \tag{2.43}$$

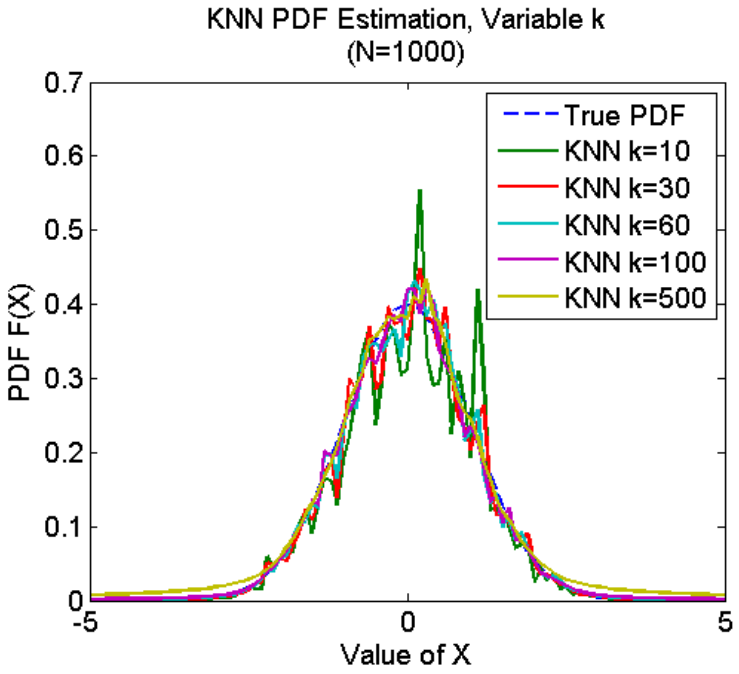

Figure 17: KNN pdf Estimation Example

Figure 17 is an example implementation of the KNN methodology and shows the smoothing effect that *k* has on the resulting estimation.

2.4 Summary

This chapter explored some of the existing research related to vision navigation in order to highlight the current state-of-the-art in this area. It can be observed that the trade-space for image processing is vast, and although many approaches to image and vision navigation have been explored, with the exception of Larson's work, the research findings show the lack of emphasis on rigorous and quantifiable robustness, in terms of integrity. A critical gap that this research is intended, in small part, to address.

Additionally, some background information related to pdf estimation was provided. Being that this is an estimation problem being explored in this research, it should come as no surprise that that these techniques are an important aspect in much of the work presented in this dissertation.

The following chapters captures the core innovations for this research and is broken out into the two main techniques previously mentioned, the development of a discrete integrity monitor, followed by the determination of protection level confidence bounds.

3 Vision-Aided Integrity Monitor

The concept of integrity, particularly for navigation systems, usually assumes unsafe hazards can occur with unacceptably high probability any time an undetected navigation error exceeds some alert limit. Under some applications, such as formation flight or aerial refueling, there are uniquely defined areas that are considered unsafe. In these cases, these unsafe conditions or relative position areas are not direction independent, but rather in the direction of the other aircraft. This fact leads to an interesting conceptual question of whether integrity can be provided by narrowly limiting the probability the relative state is not in an unsafe condition (as opposed to the error exceeding a specified bound). Consider Figure 24, which demonstrates the nominal aerial refueling envelope and contrasts that with an area considered unsafe, being the area beyond the Airframe Safety Boundary. Taking advantage of the fact that the refueling operation can be broken up into two distinct conditions, (1) the relative navigation system estimates the refueling aircraft to be in the refueling envelope and that is in fact the case; or (2) the relative navigation state indicates the refueling aircraft to be in the refueling envelope when in actuality the refueling aircraft has breached the safety boundary. From an integrity perspective, the only condition of concern is (2), the case where the navigation system provides hazardously misleading information (HMI). This proposed research will answers the following question: Can the concept of integrity be defined in terms of these distinct conditions, by developing a discrete detection scheme to determine which condition holds true?

To illustrate this concept, consider Figure 18 which shows a pdf of the image correspondence between the real image and the rendered image, using the navigation estimate, for both safe and unsafe conditions. A notional detection rule, based upon the image correspondence is also illustrated. The shaded areas represent the probability of False Alarm and the probability of Mis-Detection (P_{MD}), two key performance metrics associated with a detection scheme such as this. The integrity risk associated with such an approach is equivalent to P_{MD}, which is the probability that the detection scheme indicates a safe condition, but in fact the true relative state is in violation of the safety boundary.

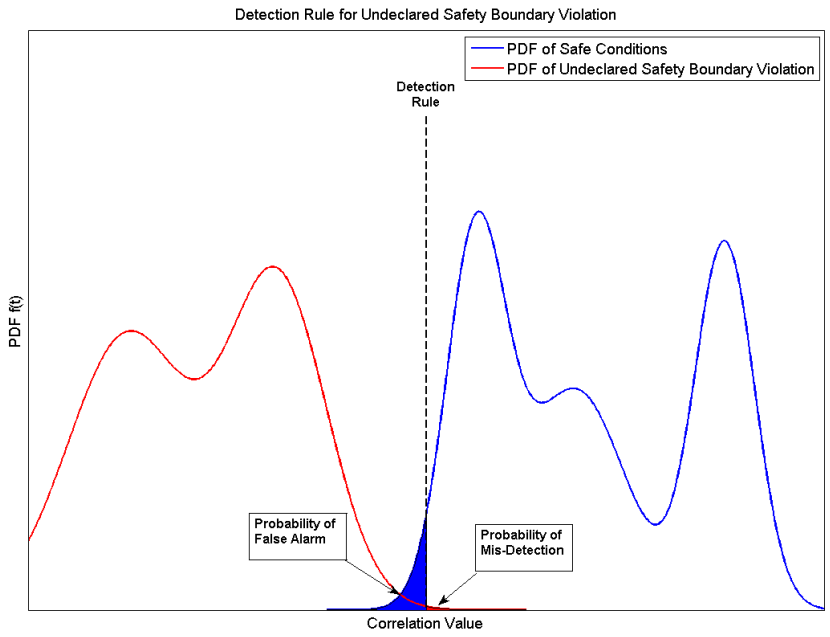

Figure 18: Detection Rule Concept for Undeclared Safety Boundary Violation

The research includes an exhaustive simulation to develop the associated pdf's for each condition. This can be achieved by computing the image correspondence from a

proper sampling of all of the possible relative state conditions within the scope of this aerial refueling scenario. Subsequently, pdf estimation techniques, as described in Section 2.3, will be applied to develop the associated pdf for each condition. Once this is achieved, classical detection schemes can then be applied to determine a detection rule based on the image correspondence.

Several detection schemes are considered for this research, to include Bayesian, Minimax, and Neyman-Pearson [56]. The research in this area is not only include the derivation of the associate detection scheme, but also evaluate predicted performance metrics using methods such as Receiver Operating Curves (ROC). These metrics are key to showing feasibility of this approach to provide an associated relative position integrity using image rendering image techniques. The following sections are descriptions of some of the proposed detection schemes to be explored.

3.2.1 *Bayesian Detection*

To understand the Bayesian detection scheme let H_0 represents the safe condition and H_1 the undetected safety boundary violation. $C_{i,j}$ are the associated cost parameters, which are essentially scalar weights for the cost of guessing H_i when H_j is true. Consider the expression for risk (R), associated with these cost parameters, defined as [56][57]

$$R = E[C] = \sum_i \sum_j C_{ij} P[guess\ H_i; H_j\ is\ true] \tag{3.1}$$

which can be further expressed mathematically as

$$R = E[C] = \sum_i \sum_j C_{ij} P[H_i | H_j] P[H_j] \tag{3.2}$$

It should be noted that $P[H_i|H_j]$ is the only term affected by the decision rule. Let Γ be defined as the entire state space of $\bar{x}$. Further, let Γ_0 be the set of relative states, $\bar{x}$, that leads to a selection of the H_0 condition, and likewise a set Γ_1 that leads to a selection of the H_1 condition, such that the entire state space Γ is categorized as Γ_1 or Γ_0. The risk expression can then be expanded to

$$\begin{aligned} R = {} & C_{00} P_0 \int_{\Gamma_0} p_0(\bar{x})\, d\bar{x} + C_{10} P_0 \int_{\Gamma_1} p_0(\bar{x})\, d\bar{x} \\ & + C_{11} P_1 \int_{\Gamma_1} p_1(\bar{x})\, d\bar{x} + C_{01} P_1 \int_{\Gamma_0} p_1(\bar{x})\, d\bar{x} \end{aligned} \tag{3.3}$$

Since the union of Γ_0 and Γ_1 is equivalent to the entire state space, and using the property that the integral of any pdf across the state space equals to 1, results in the following equality for any pdf

$$1 = \int_{\Gamma_0} pdf + \int_{\Gamma_1} pdf \tag{3.4}$$

Equation 3.4 can be rewritten as

$$R = C_{00}P_0 \int_{\Gamma 0} p_0(\bar{x})\, d\bar{x} + C_{10}P_0 - C_{10}P_0 \int_{\Gamma 0} p_0(\bar{x})\, d\bar{x} + C_{11}P_1 \\ - C_{11}P_1 \int_{\Gamma 0} p_1(\bar{x})\, d\bar{x} + C_{01}P_1 \int_{\Gamma 0} p_1(\bar{x})\, d\bar{x} \tag{3.5}$$

Grouping terms and defining all integral terms as $I(\bar{x})$ then yields

$$R = C_{10}P_0 + C_{11}P_1 + \int_{\Gamma 0} I(\bar{x})\, d\bar{x} \tag{3.6}$$

To minimize *R* averaged over all $\bar{x}$, we need only to minimize $I(\bar{x})$ for each $\bar{x}$. However, the only choice is whether or not each $\bar{x}$ is in Γ_0. Therefore, we can define a set rule such that for a given $\bar{x} \in \Gamma_0$ if and only if $I(\bar{x}) < 0$. This yields a specific set definition for Γ_0 that minimizes the overall risk, *R*, defined by

$$\Gamma_0 \equiv \{\bar{x}: (C_{00} - C_{10})P_0 p_0(\bar{x}) + (C_{01} - C_{11})P_1 p_1(\bar{x}) < 0\} \tag{3.7}$$

This expression for the definition of the H_0 condition can now be simplified into the inequality expression

$$(C_{01} - C_{11})P_1 p_1(\bar{x}) < -(C_{00} - C_{10})P_0 p_0(\bar{x}) \tag{3.8}$$

Finally, solving this expression in terms of the standard Bayesian Likelihood Ratio Test (LRT), $\Lambda(x)$, is described in Equation 3.9 [56][57]

$$\Lambda(x) = \left(\frac{p_1(x)}{p_0(x)}\right) \underset{H0}{\overset{H1}{\gtrless}} \left(\frac{C_{10} - C_{00}}{C_{01} - C_{11}} \cdot \frac{P_0}{P_1}\right) = \eta \tag{3.9}$$

where $p_0(x)$ and $p_1(x)$ are the probability density functions for the corresponding H_0 and H_1 hypothesis. In this case H_0 represents the safe condition and H_1 the undetected safety boundary violation. $C_{i,j}$ are the associated cost parameters, which are essentially scalar weights for the cost of guessing H_i when H_j is true. The prior probabilities (P_0, P_1) represent the *a priori* probabilities, if known, associated with each condition. The testing parameter in this case, *x*, is chosen to be the image correspondence value between the model and the actual image.

3.2.2 *Minimax Detection*

A minimax detection rule follows the same basic scheme as Bayes, the difference being that it is based on no knowledge of the prior probability (P_0) for the H_0 condition. The minimax rule minimizes, over all decision rules (δ), the maximum of the conditional risks, $R_0(\delta)$ and $R_1(\delta)$, Equation 3.10 [56][57]. In effect, this minimizes the Risk (R) assuming the worst case P_0.

$$\min \max \{R_0(\delta), R_1(\delta)\} \tag{3.10}$$

This expression can be further understood by expressing the risk as a function of P_0

$$R(P_0, \delta) = P_0 R_0(\delta) + (1 - P_0) R_1(\delta) \tag{3.11}$$

Equation 3.11 is simply an affine function of P_0. Consider for each prior P_0 letting δ_{P_0} be a Bayes decision rule for that prior, and defining

$$V(\mathrm{P}_0) = R\left(P_0, \delta_{P_0}\right) \tag{3.12}$$

Then the minimax rule is simply the Bayes rule for the prior P_0 that maximizes $V(P_0)$. This concept is illustrated in Figure 19, where δ_{P_L} is the minimax decision rule.

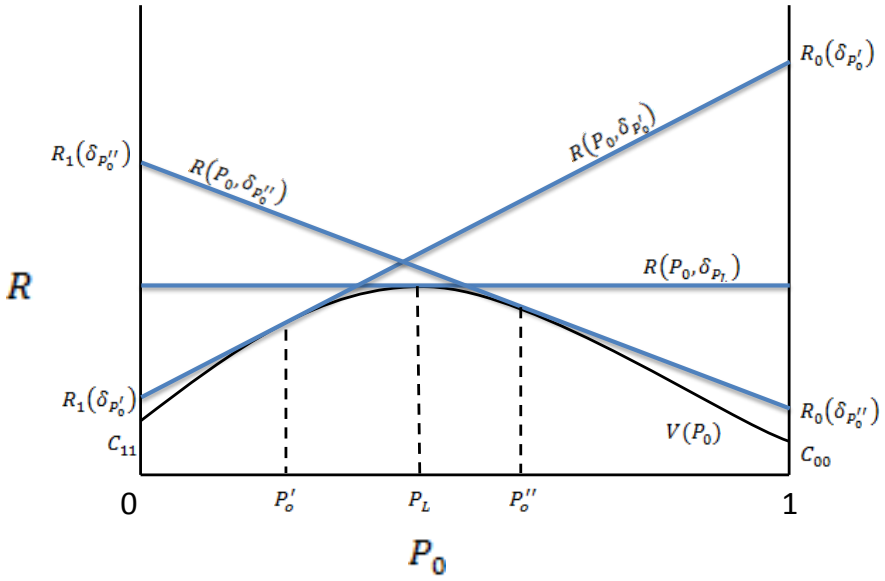

Figure 19: Minimax Detection Rule Illustration

3.2.3 *Neyman-Pearson Detection*

Another detection approach often used is the Neyman-Pearson (NP) hypothesis test [56][57]. The NP approach is based upon a constrained optimization problem, specifically designed to develop a decision rule, δ_{NP}, that optimizes the detection rate (P_D) while constrained by a maximum rate of false alarm (α):

$$\delta_{NP}(\bar{x}, \alpha) = \arg\max_{\delta} P_D(\delta) \; such \; that \; P_F(\delta) \leq \alpha \tag{3.13}$$

It is also worth noting and easily observed in Equation 3.13, that the NP approach does not rely upon a cost structure that the Bayes and minimax utilize.

The NP decision rule can be developed both analytically and graphically. One solution method for the analytic method uses a Lagrange multiplier to solve the

constrained optimization problem. However, instead of applying the Lagranian optimization constraint directly to Equation 3.13, which is a maximization problem, the expression is transformed into a minimization problem. This is achieved using the known relationship between detection and Missed-Detection (P_{MD}), $P_D = 1\text{-}P_{MD}$. Using this and applying the Lagrangian (λ), Equation 3.13 becomes

$$\delta_{NP}(\bar{x}, \alpha) = \arg\min_{\delta} (P_{MD}(\delta) + \lambda[\, P_F(\delta) - \alpha]) \tag{3.14}$$

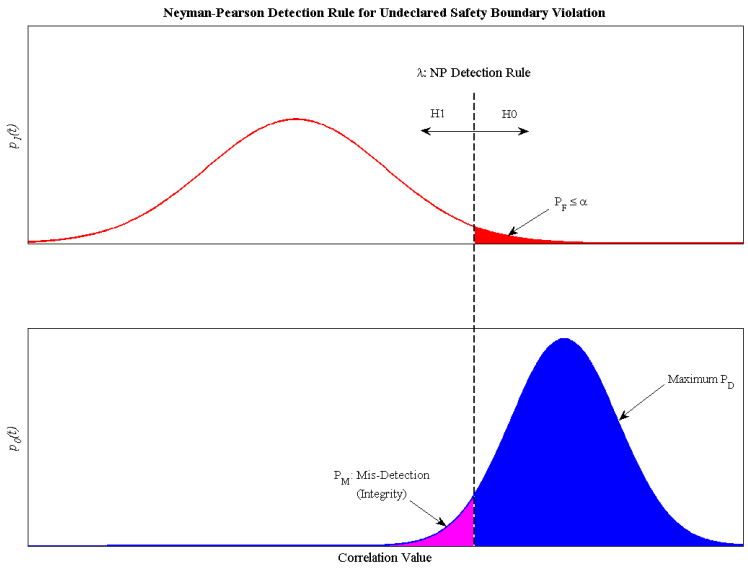

Figure 20: Graphical Neyman-Pearson Detection Rule Example

As indicated in the preceding sections, this research is based on defining a vision-aided integrity monitor in terms of detecting when the system navigation state (x) is within a defined operating region, defined as

$$X_{OR} = \{x | x \in Operating\ Region\}, \tag{3.15}$$

where X_{OR} defines the set of states within the operating region, versus being within the alert region state space (X_{AR}), which will defined as

$$X_{AR} = \{x | x \in Alert\ Region\} \tag{3.16}$$

There are four distinct conditions the integrity monitor can result in, rejection (P_R), missed-detection (P_{MD}), detection (P_D), and false-alarm P_{FA}). Figure 21 illustrates these four conditions based on the navigation estimate and the actual state of the system or truth. The performance of this type of binary (H0/H1) detection scheme can be characterized using just two of these metrics, detection and false-alarm rates, which will be the two primary performance metrics for this research. P_D is the primary metric measuring navigation integrity, describing the probability that the monitor successfully detects the condition when $x \in X_{AR}$.

Truth / Hypothesis	$x \in X_{OR}$	$x \in X_{AR}$
$H_0 : x \in X_{OR}$	P_R	P_{MD}
$H_1 : x \in X_{AR}$	P_{FA}	P_D

Figure 21: Integrity monitor detection trade space

Bayesian, Minimax, and Neyman-Pearson (NP) are a few of the detection schemes available to solve this type of binary detection problem[56]. These detection schemes rely on the knowledge of the underlying statistics of the H0 and H1 condition, often characterized in terms of the probability density functions (pdf). The main

difference between these approaches is the resulting detection rule value (δ). Once δ has been established, the resulting theoretical performance of the detector are computed by integrating the underlying pdf's of the H0 and H1 conditions, (P_{H0}) and (P_{H1}) respectively. The probability of detection (P_D) is computed as

$$P_D = \int_{\delta}^{\infty} p_{H1}(x)dx \tag{3.17}$$

The integrity performance of the monitor can also be described in terms of integrity risk or probability of missed detection (P_{MD}), which is computed as

$$P_{MD} = 1 - P_D = \int_{-\infty}^{\delta} p_{H1}(x)dx \tag{3.18}$$

Similarly, the probability of false-alarm (P_{FA}) is computed as

$$P_{FA} = \int_{\delta}^{\infty} p_{H0}(x)dx \tag{3.19}$$

This is represented graphically in Figure 22.

The pdf's represent the statistical distributions of image correspondence value for the respective H0/H1 condition. The general detection rule premise is such that for a given sensor image, the underlying pdf for the "best" image correspondence with the rendered reference set is sufficiently distinct when the sensor image is in an H0 condition versus H1.

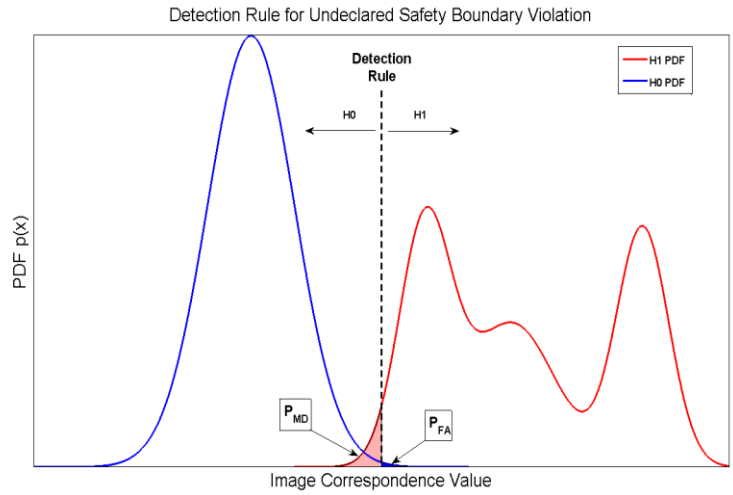

Figure 22: Graphical illustration of detection performance

The characteristics of the H0/H1 pdf's that dictate the monitor performance are dependent on many factors, including the fidelity and accuracy of the world model, the general observability of the image rendering process, and the image correspondence approach for the specific application. For this research two image correspondence techniques were utilized to evaluate the overall integrity monitor approach.

The first image correspondence technique evaluated is a simple binary silhouette (SIL). In this approach, both the sensor image $I_S(x)$ and reference image set $I_R(\hat{x})$ are converted to a silhouette using pre-defined thresholds to first convert the RGB images to grayscale then subsequently to a binary image. The image correspondence function computes the percentage of overlap between the silhouettes using the following ratio expression:

$$img\_corr(I_s, I_R) = \frac{\left|I_{S\_SIL} \cap I_{R\_SIL}\right|}{\left|I_{S\_SIL} \cup I_{R\_SIL}\right|} \tag{3.20}$$

where I_{S_SIL} and I_{R_SIL} represent the set of silhouette pixels from the sensor and rendered images respectively. The resulting image correspondence is based on the ratio of the cardinality of these sets. The navigation state estimate (x) that yields the maximum image correspondence value from the set of rendered reference images or template database is considered the most likely for that particular image sensor (I_S),

$$X_{IC\_SIL} = \max_{x}[img\_corr(I_{S\_SIL}, I_{R_{SIL}}(x))] \tag{3.21}$$

The second image correspondence utilizes edge features for the image correspondence process, which was previously demonstrated to provide a means for navigation estimation[10] [15]. Under this approach, magnitude of gradient (GRD) processing is used, in which the sensor image and the rendered reference images are preprocessed through a Prewitt filter [40] to determine changes in image intensities between adjacent pixels. This process computes the components of the gradient, *Gx* and *Gy*. The gradient magnitude is computed by RSS of the x-y components and normalized, resulting in an edge detection. A Gaussian blur filter [58][59], is then applied to the output of the edge detection. The normalized Gaussian blur filter array, h(x,y), is constructed as described by

$$h(x_{pix}, y_{pix}) = e^{-\frac{x_{pix}^2 + y_{pix}^2}{2\sigma^2}} \tag{3.22}$$

where x_{pix} and y_{pix} are the corresponding pixel coordinates, and σ is the associated standard deviation of the Gaussian distribution [60]. The application of the Gaussian blurring compensates for the spatial discrepancies between the discrete reference set or

template database and the sensor image. Finally, the resulting feature images, including both the reference image (I_{R_GRD}) and the sensor image (I_{S_GRD}), are processed through a sum squared difference (SSD) image correspondence, described by

$$img\_corr(I_S, I_R) = \sum_{i=1}^{n} (I_{S\_GRD_i} - I_{R\_GRD_i})^2 \tag{3.23}$$

where *n* is the number of image pixels. The resulting pdf's are based on the best image correspondence with the RE reference set, which is the minimum for the GRD processing.

$$X_{IC\_GRD} = \min_{x}[img\_corr(I_{S\_GRD}, I_{R\_GRD}(x))] \tag{3.24}$$

These image correspondences build the basis of the detection metric, utilizing both the sensor image (I_S) and the rendered reference set (I_R), which is spatially distributed across the operating region, illustrated by Figure 23. This illustrated example shows instances of both a H0 and H1 sensor image (blue and red, respectively). The underlying H0/H1 pdf's for establishing the detection threshold are determined by sampling sensor images from X_{OR} and X_{AR} and computing the image correspondence against I_R. This can be done through a combination of high fidelity simulation and/or test data. The overall performance of the integrity monitor will be dictated by these underlying distributions. The following sections show the results of this integrity monitor approach for an aerial refueling application.

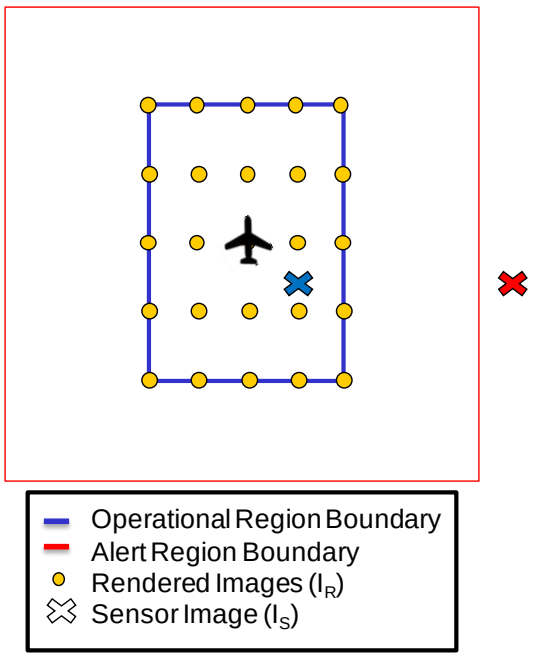

Figure 23: Simplified example of rendered reference set (IR) illustrating image correspondence process for integrity monitoring

3.3 Simulation Evaluation

To explore the performance of the proposed integrity monitor approach, an aerial refueling (AR) application was modeled within a simulation environment. The AR operation lends itself well to the construct of the proposed integrity monitor and is developed to show that the system (refueling aircraft) is in the Refueling Envelope (RE), and has not violated the alert limit, which in the AR case is the Safety Boundary (SB). In this operational case, H0 is defined as the condition when the integrity monitor determines the refueling aircraft is in the refueling envelope (RE), and H1 as the case when the integrity monitor determines the refueling aircraft to be within the SB. A validity region is also defined in order to bound the problem. in which it is assumed that the refueling aircraft is always within, under both H0 and H1 conditions, as shown in Figure 24.

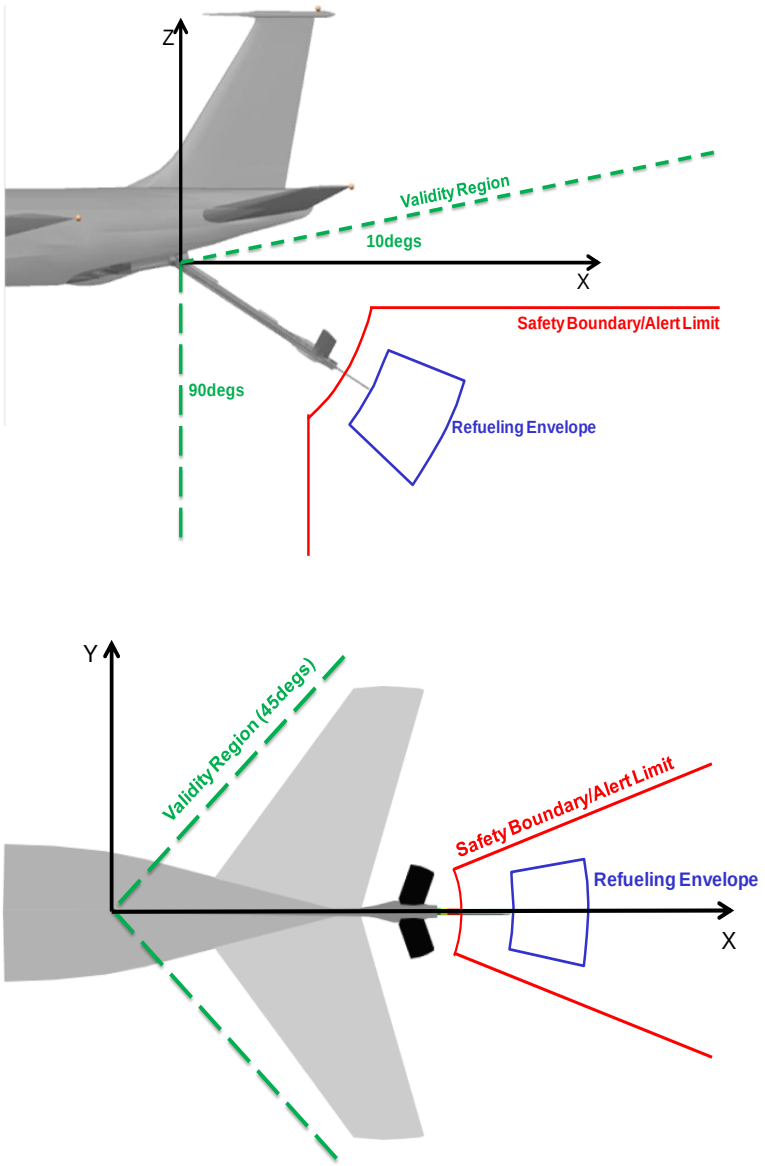

Figure 24: Integrity regions of interests for Aerial Refueling (AR) application

To determine the underlying H0/H1 distributions, a set of reference images uniformly sampled from the RE was rendered using the associated tanker and camera models. This rendered image set is used as the common basis for performing the image correspondence with the actual sensor image.

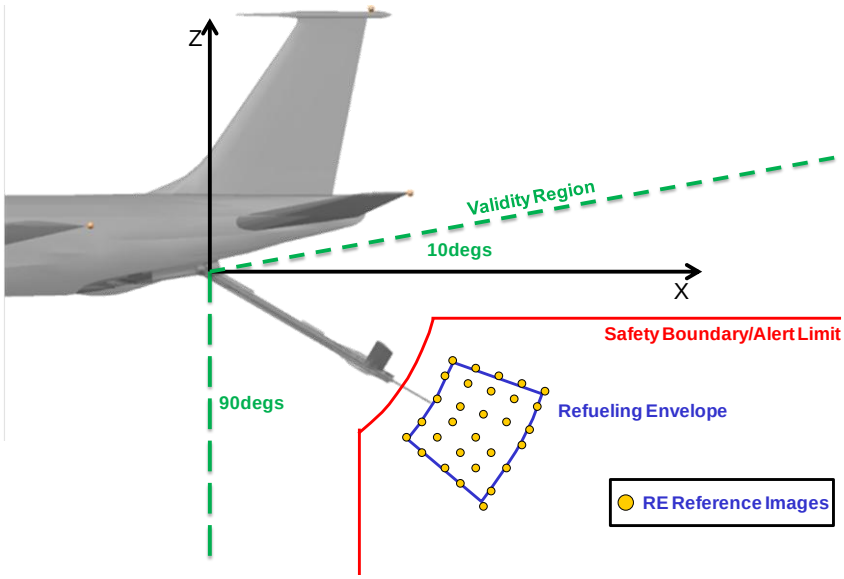

Figure 25: Illustrated example of a rendered H0 image set for the Refueling Envelope (RE) used as the correspondence basis for the integrity detection metric

The baseline RE reference set utilized for this research was developed using 504 rendered images distributed in a spherically uniform manner across the entire RE volume. Then two random sets of simulated sensor images were generated and drawn from both RE and SB regions. It is assumed that the refueling aircraft and corresponding sensor images are within the validity region in order to bound the simulation. This bounding assumption is an acceptable constraint, given that the system most likely had to pass several operational checks to ensure the refueling aircraft is in the general region of the RE as defined by the validity region. In order to get detailed statistical representation of the pdf's, particularly at the tails of the distribution, both RE and SB image sets included over 100,000 simulated sensor images, representing true states of the refueling aircraft. The simulation environment for this analysis utilizes the same refueling tanker model for the sensor images and the RE reference set, which eliminates the effects of modeling

errors. Additionally, variations in the attitude are currently not considered. The resulting pdf's for H0 (blue) and H1 (red) conditions are shown in Figure 26 and shows generally good distinction between the H0 and H1 hypothesis, which is a necessary condition to achieve good detection performance. Several techniques were evaluated for determining the pdf including histogram, nearest neighbor, and kernel with a Gaussian weighting function [55]. These underlying H0 and H1 distributions will be used as the basis for designing the detection thresholds, based on the image correspondence of the sensor image with the RE reference set. These results assume uniform prior distributions across the RE and SB regions; however, it would be relatively straightforward to incorporate non-uniform prior information, based on a particular application, as available.

Detection schemes are often characterized using Receiver Operating Characteristics or ROC curves [56][57], which illustrate the detection monitor trade-off between probability of detection and probability of false-alarm. The predicted detection performance for this AR application is a function of these underlying H0/H1 pdf's, and this performance is captured in the ROC curves shown in Figure 27. The ROC curves in Figure 27 demonstrates that 10^{-3} level integrity monitor detection performance (P_D) is realizable for both SIL and GRD image correspondence approaches, while still maintaining a reasonable probability of false-alarm (P_{FA}) of less than 0.05 (5%). The SIL approach demonstrates slightly better performance than GRD under the chosen image resolution and RE reference set density.

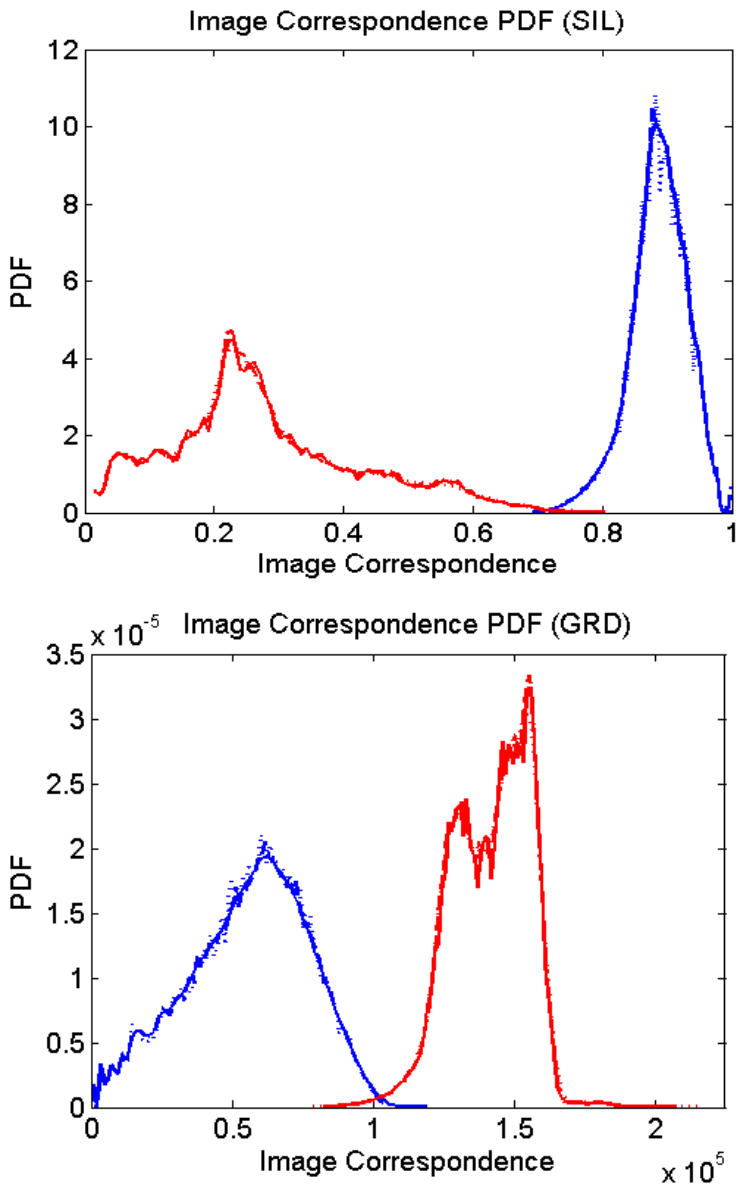

Figure 26: Underlying image correspondence distribution for H0 (Blue) and H1 (Red) conditions

Normally, theoretical ROC curves would extend through the whole range of values [0,1] for both P_d and P_{FA}, however this assumes unbounded pdf's. Doing so would require an infinite amount of simulation cases and is obviously not practical for a simulation evaluation to gain statistics necessary to extend the pdf's near the entire theoretical ranges. Overbounding of the pdf tails [19] could be performed to extrapolate and extend

the tails of H0/H1 pdf's to determine the integrity detection performance beyond the current ranges, but this was not performed as part of this research.

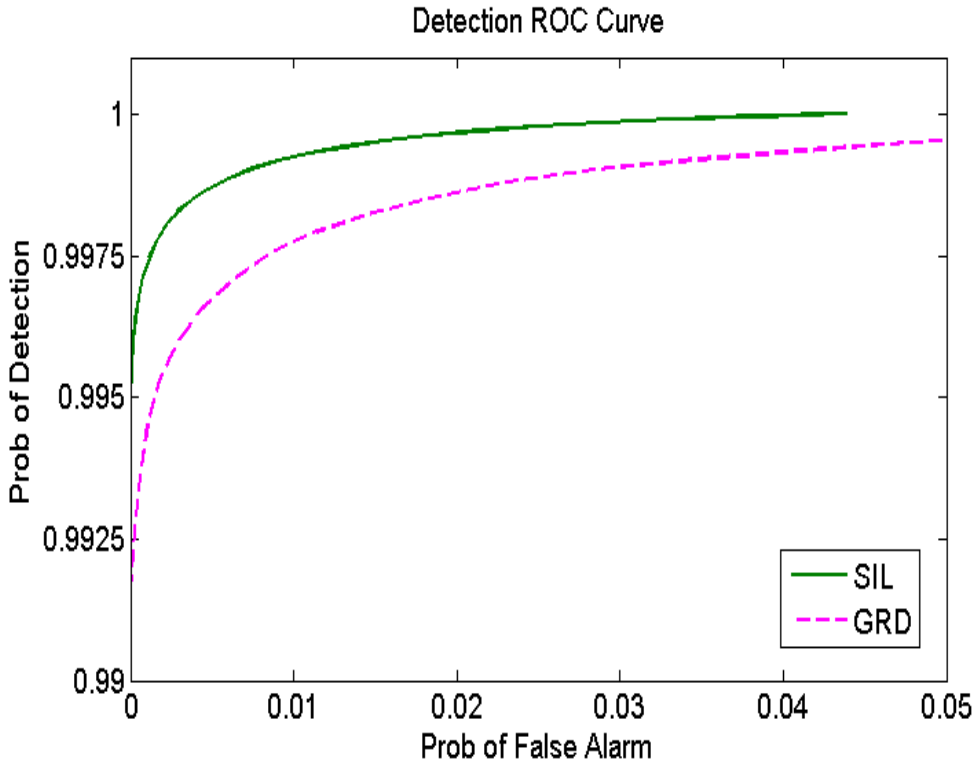

Figure 27: Predicted integrity detection performance for both SIL and GRD image correspondence techniques

In most applications, conditions exist that are outside of the nominally defined operational envelope, but yet are not significant enough deviations to be considered safety risks that require alerts and action. Such a case exists for the refueling operation under consideration in this research, where there exists a region outside of the RE, but not in the SB, which we will refer to as the operational limit volume (OLV). The current definitions of H0 and H1 for the image-aided integrity monitor approaches developed above only consider conditions within the RE or the SB volume, and not within the OLV

volume. OLV conditions were omitted, since they technically aren't considered a safety or integrity risk. However, it is possible under certain implementations and operational considerations that integrity monitoring coverage is desired under these OLV conditions.

Using the same analysis process as the original evaluation, an updated simulation was performed, this time considering all points within the validity region, including the OLV points. To construct a detection scheme under this new paradigm, the OLV conditions must be either mapped to existing H0 or H1 hypothesis, or a new hypothesis must be defined, possibly creating a M-ary hypothesis scenario. The approach taken for this research was to consider OLV conditions as a safety risk, which is a conservative approach, rather than defining a new hypotheses. The resulting image correspondence distributions is shown in Figure 28 - Figure 31. Figure 28 and Figure 30 show the difference the OLV points have on the underlying pdf distributions. As expected, when the OLV points are excluded, the pdf's track the original distributions quite well. The impact of including sensor locations from the OLV is clear from these figures, yielding a much bigger overlap between the H0/H1 conditions.

Much like the pdf's, the ROC curves align with the previous results quite well when the OLV conditions are omitted, but take a order of magnitude integrity performance hit when OLV are captured under the existing H0/H1 definition and detection thresholds. Even under this conservative assumption, the overall monitor performance still yields a 0.96 (96%) detection rate at a 0.05 (5%) false-alarm rate, as illustrated by the ROC curves shown in Figure 30 and Figure 31. It is likely that these

results can be significantly improved by redefining the terms of the H0 and H1 conditions or defining a H2 condition specifically for the OLV region.

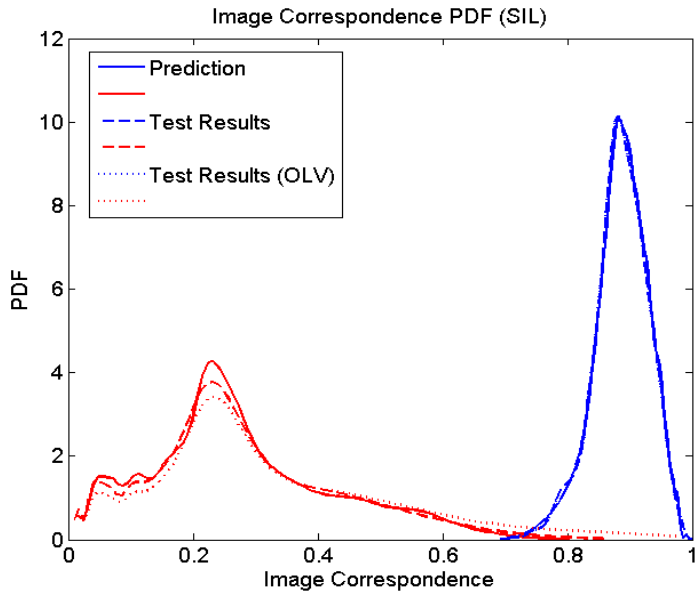

Figure 28: Simulation Testing Results Account for OLV States (pdf - SIL)

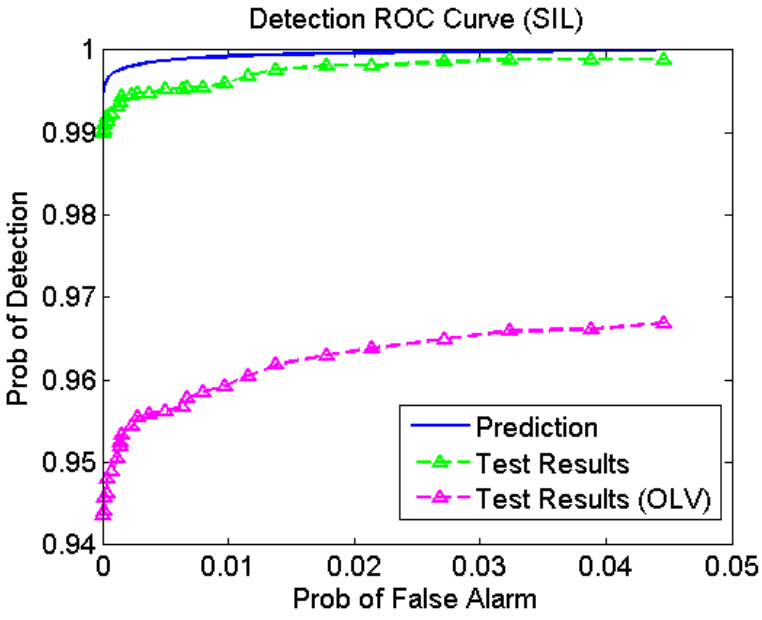

Figure 29: Simulation Testing Results Account for OLV States (ROC - SIL)

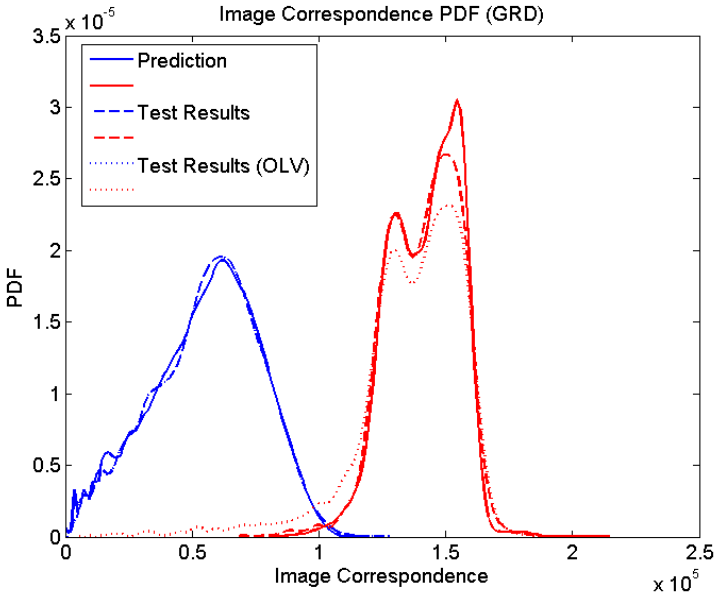

Figure 30: Simulation Testing Results Account for OLV States (pdf - GRD)

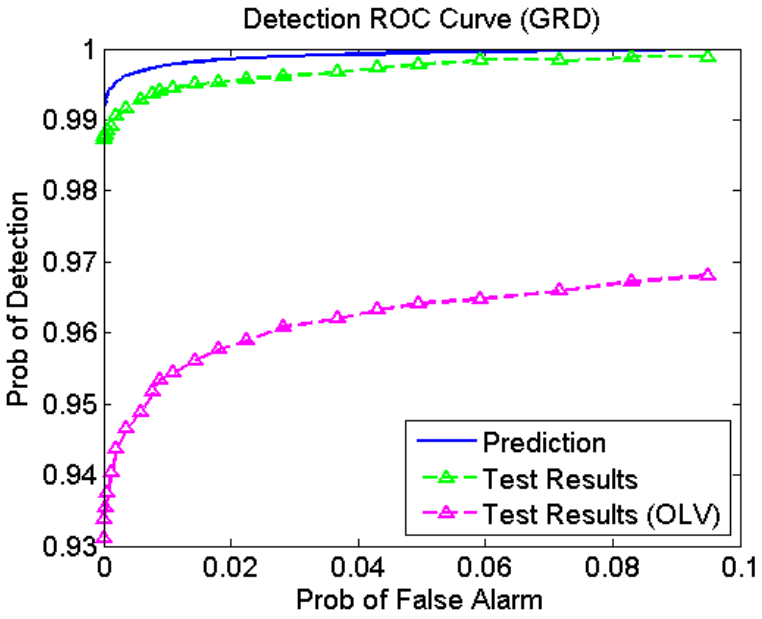

Figure 31: Simulation Testing Results Account for OLV States (ROC - GRD)

3.4 Sensitivity Analysis

In addition to the baseline integrity monitor results, various sensitivity studies were performed to evaluate the integrity monitor performance impacts of environmental and hardware considerations. These sensitivity evaluations focused on common vision-based considerations such as sensor distortions and lighting conditions, and monitor design choices such as pixel resolution and reference image density. The sensitivity aspects that were evaluated under this research including (1) number of reference images, (2) the effects of image distortion, and (3) pixel resolution.

3.4.1 *Reference Set Density*

Figure 33 shows how the density of the rendered reference set affects the underlying H0/H1 distribution. A reference set of 504 RE images was used as the baseline for this research, but given no hard integrity requirement to aim for, this value was somewhat of an arbitrary choice between performance and processing speed. To explore this trade space, a 50% increase and decrease in the number of reference images was evaluated (288 and 729 images, respectively). As expected, if the number of reference images used is increased, hence higher density reference image set, the distinction between H0 and H1 becomes improved, which would lead to improved integrity detection performance. Subsequently, when the number of reference images is reduced, a small but noticeable increase in the overlap is realized. The results of this analysis also showed that the result is mainly on the H0 pdf and has somewhat minimal effect on the H1 pdf. This tradeoff between performance and number of reference images, and associated processing cost, is potentially beneficial, allowing a designer to

trade-off processing power for performance as necessary for that particular application and the associated integrity monitor performance requirements.

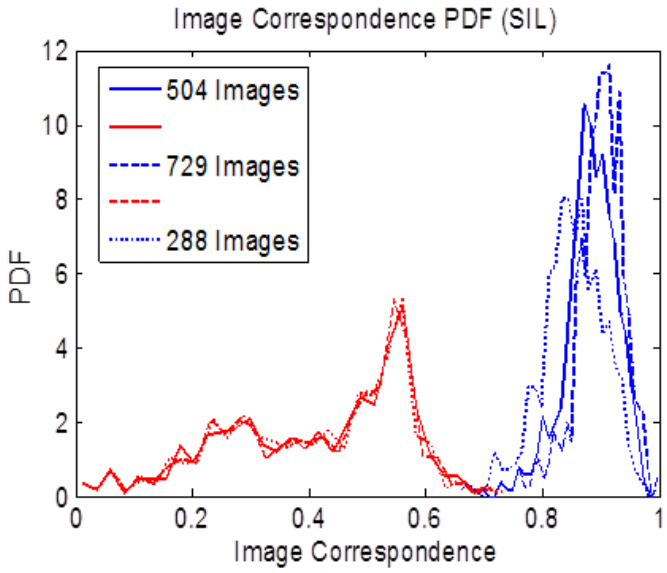

Figure 32: Image Density Sensitivity Analysis

3.4.2 *Image Distortion*

Since the initial evaluation assumed perfect sensor data, the sensitivity analysis evaluated the effect of lens distortion on the H0/H1 pdf's. The applied distortions were based on radial and tangential error models [61][62]. The radial distortion, as the name implies, distorts image pixel points (u_i,v_i) radially in the image plane and is expressed as

$$\begin{bmatrix} \delta u_i^{(r)} \\ \delta v_i^{(r)} \end{bmatrix} = \begin{bmatrix} u_i(k_1 r_i^2 + k_2 r_i^4) \\ v_i(k_1 r_i^2 + k_2 r_i^4) \end{bmatrix} \tag{3.25}$$

where k1, k2,… are the radial error model coefficients dictating the extent of the distortion, and

$$r_i = \sqrt{u_i^2 + v_i^2} \tag{3.26}$$

For most applications only the first one or two coefficients are utilized for the radial distortion model. The tangential error model accounts for situations where the centers of curvature of the lens are not strictly collinear, resulting in a distortion that has both radial and tangential components. The tangential error model is expressed as

$$\begin{bmatrix} \delta u_i^{(t)} \\ \delta v_i^{(t)} \end{bmatrix} = \begin{bmatrix} 2p_1 u_i v_i + p_2 (r_i^2 + 2u_i^2) \\ p_1 (r_i^2 + 2v_i^2) + 2p_2 u_i v_i \end{bmatrix} \tag{3.27}$$

where p1 and p2 are the tangential error model coefficients. Combining these two error effects, we arrive at the distortion model applied for this analysis.

$$\begin{bmatrix} u_i \\ v_i \end{bmatrix} = \begin{bmatrix} u_i + \delta u_i^{(r)} + \delta u_i^{(t)} \\ v_i + \delta v_i^{(r)} + \delta v_i^{(t)} \end{bmatrix} \tag{3.30}$$

Many tools and techniques exist for calibrating out many of these distortions effects; however, residual errors still occur. Therefore, this analysis evaluated the impact of these residual effects. Recent flight tests evaluating vision-aided systems for aerial refueling were conducted [15], where the distortion parameters were estimated, including an estimate in the calibration uncertainty and possible residual distortion errors. These calibration uncertainty parameters were used to apply distortion effects to the simulated sensor images (I_S) and applied such that they represented a 95% certainty of the residual error to represent an outer envelope case for this type of sensor. As seen in Figure 33, the

impact on the H0/H1 pdf's are very minimal, and the results demonstrate a potential robustness to this common type of sensor effect.

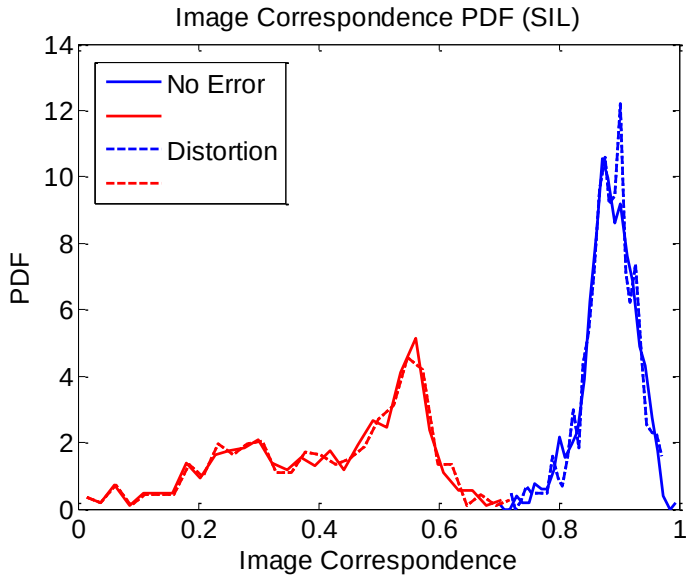

Figure 33: Image Distortion Sensitivity Analysis

3.4.3 *Pixel Resolution*

Since this research is mainly a feasibility analysis where post-processing is acceptable, processing cost was an inexpensive commodity. However, one of the big trade-offs that is often encountered when using vision systems, particularly when considering real-time applications, is image resolution and often the associated system performance versus processing costs. To consider this aspect, a sensitivity analysis was performed to evaluate how image resolution, which had a strong relationship to the processing time for this analysis, would affect the performance of the proposed integrity monitor. To this end, 8 different pixel resolutions were evaluated, and the resulting H0/H1 pdfs are shown in Figure 34. The results show a surprising robustness to pixel resolution, indicating only marginal performance impacts down to extremely limited

pixel densities. An interesting phenomena is that reducing pixel density actually shifts the distributions to higher image correspondence values. This phenomena is a result of the decrease in pixel resolution causing a spatial lowpass filter type effect on the images, allowing slightly varying images under higher resolution. This robustness, particularly the results shown in Figure 34, indicates a large trade space for reducing the processing burden for a real-time implementation of this integrity monitor concept.

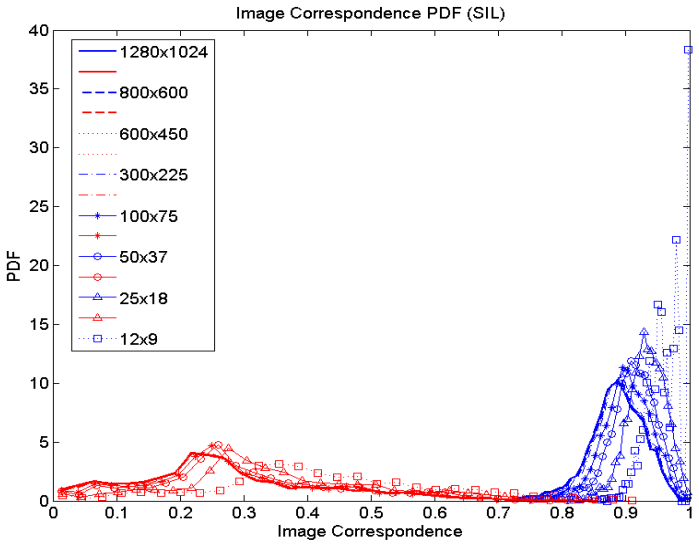

Figure 34: Image Pixel Resolution Sensitivity Analysis

3.5 Ratio Test Integrity Test

The initial integrity monitor results discussed thus far only utilized reference images from the operational region, RE. However, it is also possible to utilize a reference image set created with rendered images from the alert region, SB, by including an additional image correspondence process between the sensor image and rendered SB reference set. This is done to create a ratio test statistic as the detection metric. In this case Equations (16) and (17) express the ratio of the highest image correspondence

between the RE and SB reference sets. This approach is very analogous to the use of ratio tests for integer fixing [63].

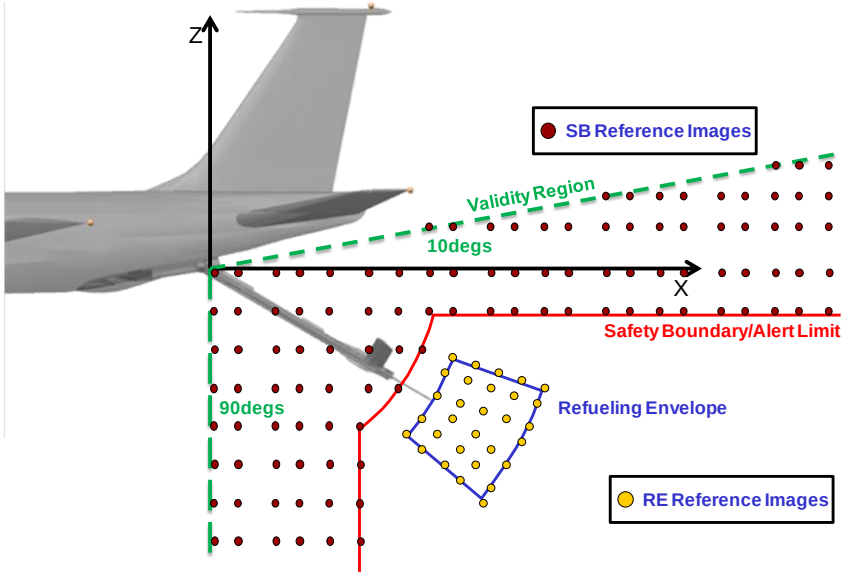

Figure 35: Illustrated example of the rendered H1 image set of the Safety Boundary (SB) region used as the correspondence basis for the ratio test integrity detection metric

The expression for this type of ratio test, utilizing the silhouette image correspondence techniques, is expressed as

$$X_{IC\_SIL} = \frac{\max\limits_{R\_RE}[img\_corr(I_s, I_{R\_RE})]}{\max\limits_{R\_SB}[img\_corr(I_s, I_{R\_SB})]} \tag{3.29}$$

and the corresponding gradient threshold ratio expression is

$$X_{IC\_GRD} = \frac{\min\limits_{R\_RE}[img\_corr(I_s, I_{R\_RE})]}{\min\limits_{R\_SB}[img\_corr(I_s, I_{R\_SB})]} \tag{3.30}$$

Figure 36 show the results of this approach. The top figures in these plots are identical to the distributions using solely the RE reference set, shown in Figure 26. The SB pdf (middle figure) shows the results from utilizing the new SB reference set of over 1100 data points and corresponding rendered images, distributed uniformly about the SB area. The SB set distribution has the opposite trend, in this case when the sensor is in the RE it has lower correspondence with the SB reference set, as compared to the RE set. It also easily observed that when using the SB reference set, there is less distinction between the underlying H0/H1 pdf's. This is mainly attributed to the density discrepancy of the reference sets, which is a result of the much larger SB volume. The SB reference set was limited for processing time considerations. However, Figure 37 was shown that when the ratio detection threshold is utilized, an improved detection performance is realized.

The resulting ROC detection performance of the ratio threshold approach is shown in Figure 38. As with the single RE reference set, the SIL image correspondence approach yields the best H1 detection performance, resulting in the best integrity protection. In fact, the ratio threshold yields an order of magnitude detection performance improvement over the performance using only an RE reference set. The GRD ratio detection performance also yields improved performance, and is comparable to the SIL image correspondence approach solely with RE reference set. The SIL ratio test yielded on the order of 10^{-4} to 10^{-5} level integrity detection performance. In fact, the SIL ratio approach was able to successfully detect all ~100,000 H1 conditions evaluated while maintaining a probability of false alarm of just over $5x10^{-3}$.

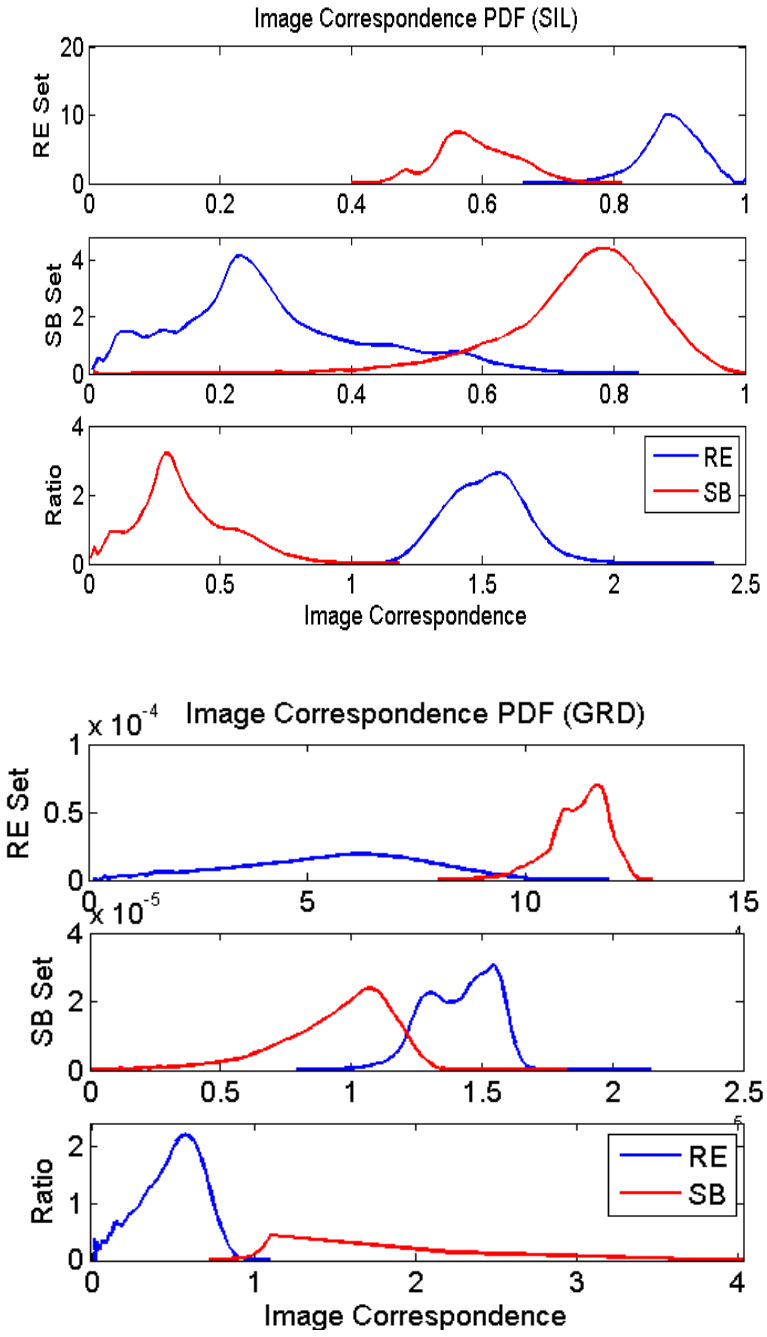

Figure 36: Underlying image correspondence distribution for (Top) H0 (Blue) and H1 (Red) using RE reference set, (Mid) using SB reference set, (Bottom) using ratio test

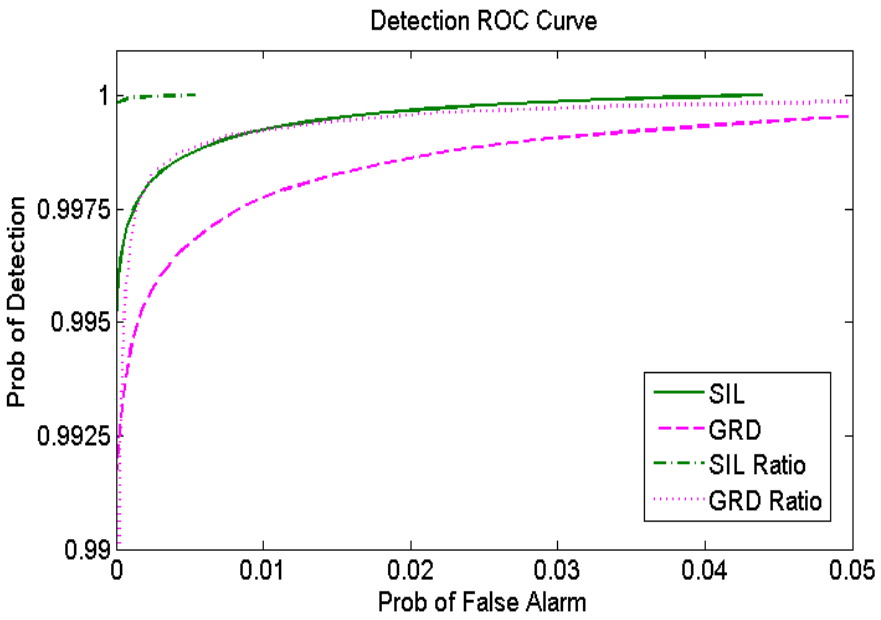

Figure 37: Predicted integrity detection performance for both SIL and GRD image correspondence techniques using a ratio of RE and SB reference sets

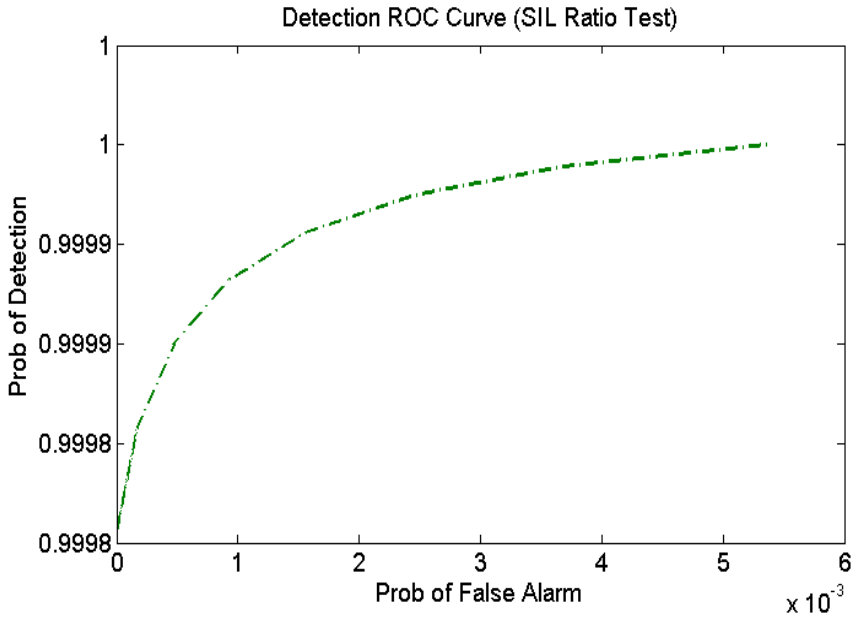

Figure 38: Predicted integrity detection performance for SIL image correspondence techniques using a ratio of RE and SB reference sets

3.6 Summary

This chapter investigated the feasibility of a vision-aided integrity monitor for precision relative navigation systems. The research posed the relative navigation integrity problem within the context of an aerial refueling application, and using image rendering, where an imaging sensor and high fidelity 3-D model is utilized. The approach was inspired by Weaver [10], in which image rendering was shown as a feasible approach for navigation estimation. This research expanded upon some of Weavers techniques, which were subsequently demonstrated in flight test [15], to show that image rendering processing can be utilized for integrity purposes. This chapter showed in fact 10^{-3} to 10^{-5} level of integrity monitoring is attainable for aerial refueling and formation flight applications, using rendering image correspondence. Having this level of independent monitoring could provide tremendous relief to a GPS based precision relative navigation system from a system safety and certification perspective. The research demonstrated the proposed integrity monitor was robust against several degrading imaging effects, including lens distortions, and reductions in pixel resolutions. The next chapter will build off of this approach and extend it to develop techniques to gain further insight into the uncertainty of the vision navigation process, by determining run-time protection levels or confidence bounds that dynamically reflect the uncertainty of the navigation solution.

4 Vision Navigation Integrity using Bayesian Inference

The next concept developed in this research for determining the relative navigation integrity, using vision techniques, follows a similar approach described by Well's [51] and discussed in Section 2.1.2.7, using Bayesian inference. This approach for integrity ultimately relies on the ability to determine a rigorous estimate of the pdf for the relative state (x) by using a known set of models and conditioned upon a set of observed measurements. This is not unlike classical navigation and estimation filter approaches, for example Kalman filtering, where the state estimate is based upon an optimal combination of modeled and measured states and further, in some cases, the uncertainty estimate is solely based on modeling assumptions. Using Bayesian inference, a full pdf determination of the navigation state is the goal. This is a progression from Well's work in which only a maximum likelihood was utilized to specifically solve for the estimate, rather than the full pdf.

The core of the Bayesian inference process is based upon use of the likelihood function, $p(d|x)$, as constructed from observed image correspondences (d) between the real image and the rendered images using the object model, and conditioned upon the relative state (x). Using this likelihood information, in conjunction with prior probability of the state $p(x)$, allows through Bayes law for the determination of the real objective, the pdf of the relative state conditioned upon the realized image correspondences,

$$p(x|d) = \frac{p(d|x)p(x)}{\int p(d|x)p(x)dx} \tag{4.1}$$

The general processing for this approach falls within the same core images processing steps discussed in Section 1.2, that of image transformation, image correspondence, and state estimation, as previously illustrated in Figure 3. Even further, this approach follows a very similar process flow to navigation state estimation, illustrated in Figure 5, the main difference being that now the main estimation objective is to determine the entire pdf of the relative position state, not just the best estimate. It is shown that by expanding the estimation process to the entire pdf allows for a determination of integrity almost identical to current navigation integrity approaches, as described in Section 1.1, where real-time protection levels are computed that bound the estimation error to an associated probability.

Figure 39 illustrates this concept and is drawn to depict the distinct processing steps involved, specifically (1) Sensor Data Generation, (2) Image Transformation, (3) Image Correspondence, and (4*)* pdf Estimation. This research focuses primarily on the last two steps, Image Correspondence and pdf estimation; however, aspects of this research involved investigating how the choice of image transformation techniques can influence the effectiveness of integrity determination.

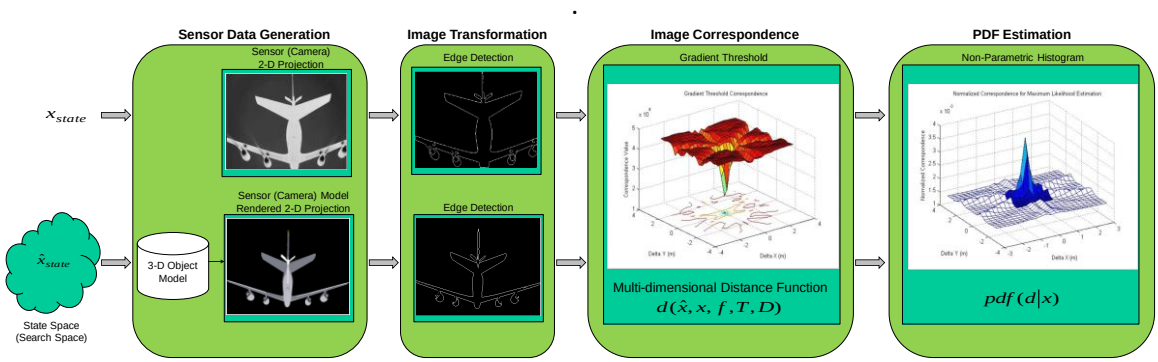

Figure 39: Conceptual Overview for Determining Integrity for Image Rendering Image Navigation

Figure 39 shows a conceptual process flow of the Bayesian inference integrity approach, with specific techniques that could be used for each step, for example the choice of edge detection as the image transformation method. Ultimately, what we are trying to solve for is $p(x|d)$, the pdf of the relative position state *x*, conditioned upon the observed correspondence distance, *d*. The mapping between the correspondence metric and the pdf estimate is a critical aspect of this research. It has been shown there are a number of different correspondence metrics available to compare images, and in all cases considered some form of distance metric is determined, *d.* The distance metric is a multi-dimensional function dependent on several considerations, including the true relative state (*x)*, estimated relative state ($\hat{x}$), the sensor imaging parameters (*f)*, the image transformation domain (*T)*, and even any potential modeling errors between the actual object and the *a priori* model (*D*).

These correspondence distance values, d, can be treated as measurements to create a likelihood function, $p(d|x)$ or *L(x)*, based upon a fixed, but unknown true relative state (*x*). Another way to look at the likelihood function is that it represents the probability of the realized distance values, given the fixed true state, *x*. Having this likelihood function,

Bayes inference can be performed to compute $p(x|d)$, which is the stated goal. Using Baye's Law, this can be performed using the expression

$$p(x|d) = \frac{p(d|x)p(x)}{p(d)} \tag{4.2}$$

where $p(d|x)$ is the conditional probability of the correspondence distance conditioned upon the relative position state, and $p(x)$ and $p(d)$ represent the prior probabilities of the relative position state and correspondence distance function, respectively. The expression in Equation 4.2 can be further reduced by replacing $p(d)$ with an equivalent expression in terms of numerator parameters,

$$p(x|d) = \frac{p(d|x)p(x)}{\int p(d|x)p(x)dx} \tag{4.3}$$

yielding an expression which is a function of only the likelihood function and the prior for the state.

Once solving for $p(x|d)$, the process for obtaining integrity is straight-forward using a protection level computed as,

$$Protection\ Level \equiv \arg\min_{V} \int_{V} p(x|d)dx \geq 1 - I_{req} \tag{4.4}$$

where I_{req} is the integrity risk requirement.

The result of Equation 4.4 is a minimum containment volume (V), such that the probability that the relative navigation state estimation error is outside of that

containment volume is less than the integrity risk requirement. As noted in Equation 1.1, this is typically broken out into vertical and horizontal/lateral directions. A further illustration of this concept for a single directional axis is shown in Figure 40.

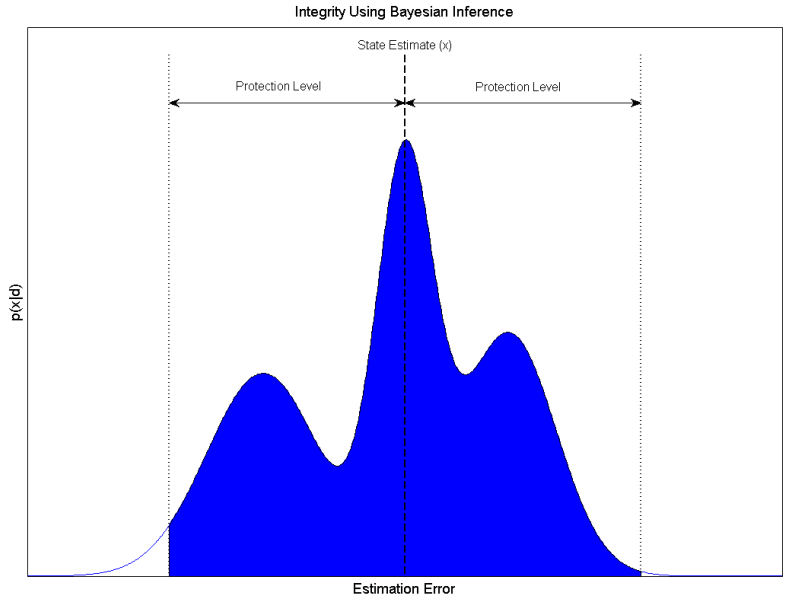

Figure 40: Protection Level Determination using Bayesian Inference

The construction of the likelihood function, $p(d|x)$, will be a major focus of the research. To achieve this, we will leverage Well's and Olson's work described in Section 2.1.2.7, in which the $p(d|x)$ is constructed based on the difference in edge intensity level, on a per pixel basis, between the real image and the rendered image. Utilizing a pixel-based approach for the image correspondence allows for multiple independent measurements to improve the likelihood function. Whereas a likelihood function relying on a consolidated image correspondence, such as SSD that yields a single scalar value per image pair, only provides a single measurement. As a result, the likelihood function with a wide distribution will occur. With multiple independent measurements, such as the case

with a per pixel image correspondence approach, the independent likelihood functions can be combined into a joint likelihood function, thus improving the distribution.

To illustrate this beneficial effect that multiple measurements have, consider the following example illustrated in Figure 41 that shows how the uncertainty is reduced by incorporation of additional measurements. In this example we are trying to estimate the parameter x, given a sensor measurement corrupted by Gaussian random noise.

$$d_i = x + w_i \tag{4.5}$$

where w is $N(0,\sigma)$. The results show how the *a posteriori* distribution, $p(x|d)$ Equation 4.3, is influenced by multiple independent measurements.

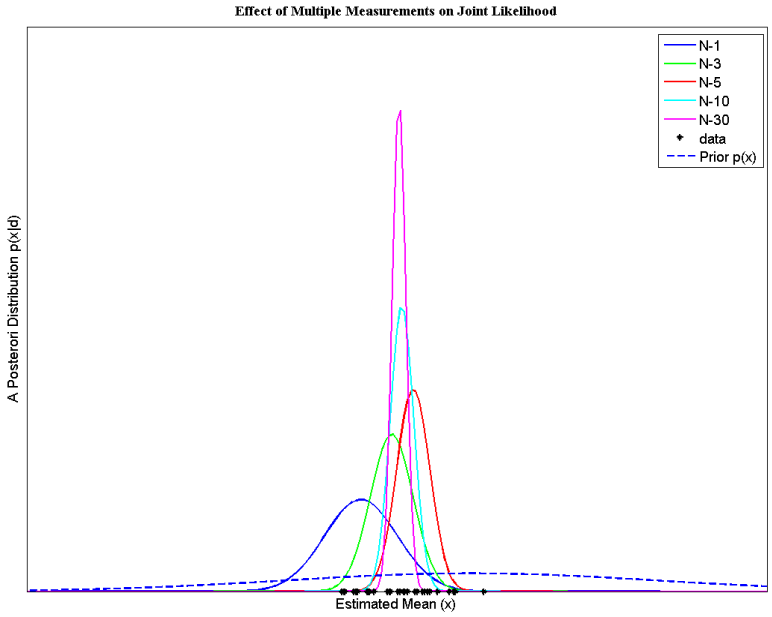

Figure 41: Effect of Multiple Measurements on Joint Likelihood Function

The key in understanding the results shown in Figure 41 lies within the interpretation of the likelihood function. When actual data is available, the pdf $\boldsymbol{p}(\mathbf{d}|\mathbf{x})$ is treated as a likelihood, and becomes a function of the unknown parameters, x. For multiple measurements, a joint likelihood is formed as

$$p(\mathrm{d}|\mathrm{x}) = \prod_{i=1} p(\mathrm{d_i}|\mathrm{x}) \tag{4.6}$$

In the above example, this joint likelihood function would be expressed as

$$p(d|x) = \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^N e^{-\frac{1}{2\sigma^2}\|d_i - x)\|^2} \tag{4.7}$$

were N is the number of independent measurements.

A key assumption in use of Equation 4.7 for constructing the joint likelihood function is that all measurements are independent and drawn from the same underlying distribution. Therefore, close attention will need to be paid to these assumptions in executing the proposed research, ensuring that independence is properly justified. In determining the underlying distribution, this will be developed using non-parametric pdf estimation techniques described in Section 2.3. The data for the pdf estimation will be generated through high-fidelity simulation with some lab or flight test confirmation possible.

An important fact of this research is that the only sensor measurements in this process are the light intensities from the imaging sensor, which are used to generate features using an image transformation technique, such as edge detection. One critical

aspect of this research will be how to best make use of these measurements within the context of the Bayesian inference process described herein. One aspect of this is to give justification for multiple independent measurements and more importantly ensure the proper utilization of the image correspondence distance. The benefit of multiple independent measurements was made clear above.

Lastly, the prior probability, $p(x)$, will require some additional assumptions in order to have a meaningful representation. One approach that will be considered is the use of an additional navigation source, such as an inertial unit. The benefit that the inertial system provides is good observability into the dynamic states, where upon standard propagation techniques can be applied to determine the prior probability. Therefore, some analysis will be required as part of this research to develop the appropriate system models to utilize such an approach. Although this approach will rely on knowledge of the prior probability, if multiple independent measurements are realized, the weight of the prior becomes quickly diminished. This effect can be seen in Figure 41.

4.1 Measurement Feature Extraction

In order to utilize the imagery from optical sensors in an effective way, the imagery must be transformed into feature space. To understand this process, let's first consider how object features $F(x)$, based on the true object state (x) with respect to the sensor coordinate frame, are transformed into the image frame described by intensities in the pixel frame (x_{pix}, y_{pix}), which can be described mathematically as

$$\begin{bmatrix} x_{pix} \\ y_{pix} \end{bmatrix} = \frac{1}{r_{\perp}} T_c^{pix} F(x) \tag{4.8}$$

where $r_{\perp}$ is the true orthogonal distance from the object feature to the optical center of the sensor, and T_c^{pix} is the homogenous transformation matrix from sensor frame to pixel frame [5]. Once the scene is transformed into the pixel frame of the sensor, digital imaging techniques can be performed to extract features within that scene.

Edge detection is an extremely useful and commonly used technique to detect discontinuities in intensity values, usually used to define the shape of an object[64]. A common method for edge detection is through the use of the gradient, defined by Equation 4.9

$$\nabla \mathbf{f} = \begin{bmatrix} G_x \\ G_y \end{bmatrix} = \begin{bmatrix} \frac{\partial f}{\partial x} \\ \frac{\partial f}{\partial y} \end{bmatrix} \tag{4.9}$$

where f(x,y) is a 2-D function, in this case representing an image. A desirable property of the gradient vector **f**, is that it points in the direction of the maximum change of **f** at coordinates (x,y).

To illustrate the process of gradient calculation of an image, consider a simple example of 3x3 image pixel plane, as shown in Figure 42, where z_5 is the center pixel.

z_1	z_2	z_3
z_4	z_5	z_6
z_7	z_7	z_8

Figure 42: Image Neighborhood Mapping for Edge Detection

A simple approximation of the partial derivative in the vertical direction at the center point could be computed by the difference

$$G_x = (z_7 + 2z_8 + z_9) - (z_1 + 2z_2 + z_3) \tag{4.10}$$

The horizontal gradient is computed in a similar fashion, and the magnitude of the gradient vector is then computed as shown in Equation 4.11.

$$g = \left[G_x^2 + G_y^2\right]^{1/2} \tag{4.11}$$

The gradient is then normalized according to

$$g_n = \frac{g - \bar{g}}{g_\sigma} \tag{4.12}$$

where $\bar{g}$ and $g_\sigma$are the mean and standard deviation of the g.

An edge is declared when this magnitude, g_n, is above a defined threshold, T.

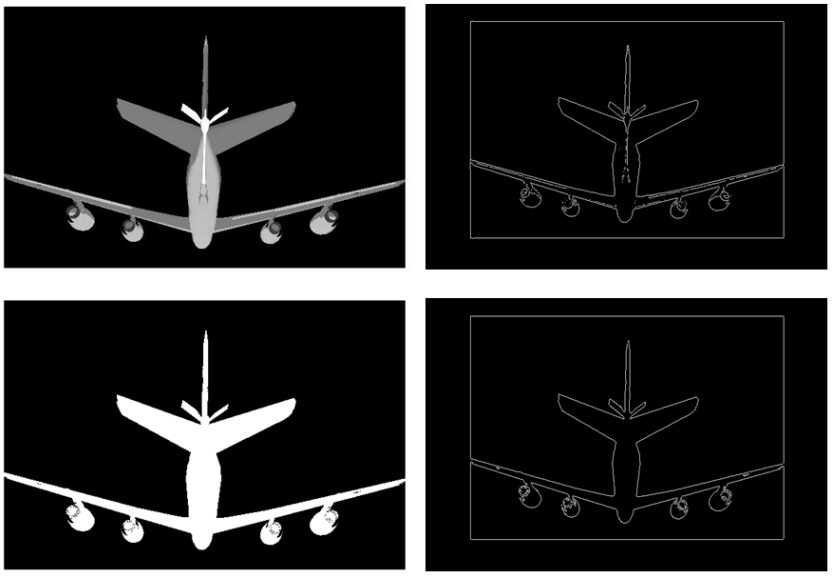

Figure 43: Edge Detection Gray-Scale versus Black-White Example

An example output of this process is show in Figure 43. The figures on the left side show gray-scale and black-white images in which the sobel edge detection function in MATLAB was applied, resulting in the figures at the right. It can easily be seen that in a general sense, the technique works and performs reasonably well in extracting the edge shapes of various objects in the image. However, some initial differences applying the edge detection can be observed, depending on the whether the gray-scale image or a black-white silhouette is used. In particular, the noise level from extraneous objects in the image is noticeably higher in the gray-scale image versus the black-white silhouette. Some of this can be equalized utilizing the threshold values in either the conversion from gray-scale to black-white or in the threshold value of the edge detection itself. The following subsections are dedicated to describing the different edge detection approaches evaluated and showing some preliminary results. In general it is difficult to discern

performance differences from the following results; rather the intent was to gather an understanding and familiarity with the process.

4.1.1 *Prewitt*

The Prewitt edge detection method, Figure 44, is very similar to that of the Sobel algorithm and can yield slightly noisier results in favor of a small reduction in computational complexity.

-1	-1	-1
0	0	0
1	1	1

$$G_x = (z_7 + z_8 + z_9) - (z_1 + z_2 + z_3)$$

-1	0	1
-1	0	1
-1	0	1

$$G_y = (z_3 + z_6 + z_9) - (z_1 + z_4 + z_7)$$

Figure 44: Prewitt Edge Detection [64]

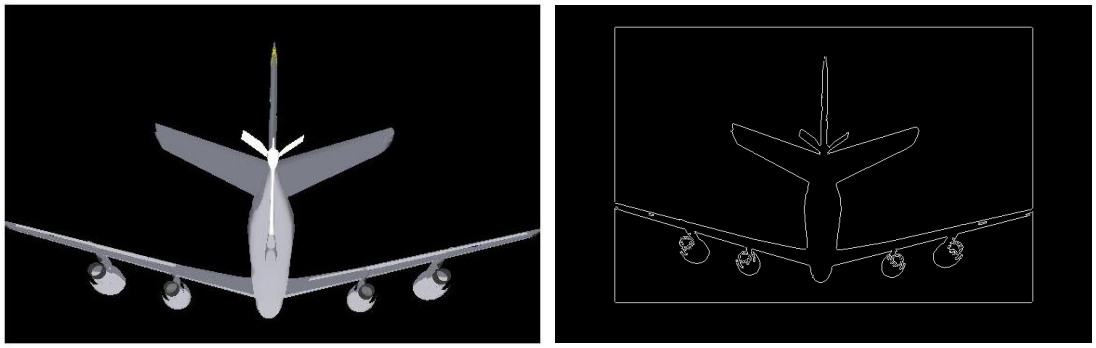

Figure 45: Edge Detection Example - Prewitt

4.1.2 *Gaussian Blurring*

One approach for dealing with modeling errors is through the application of a Gaussian Blur filter [58][59], applied to the output of the edge detection step described in the previous section. The normalized Gaussian Blur filter array, *h(x,y)*, is constructed as described by [60]

$$h(x_{pix}, y_{pix}) = e^{-\frac{x_{pix}^2 + y_{pix}^2}{2\sigma^2}} \tag{4.13}$$

where x and y are the corresponding pixel coordinates, and σ is the associated standard deviation of the Gaussian distribution. Figure 46 illustrates several examples of the resulting Gaussian Blur filter array as a function of sigma (σ).

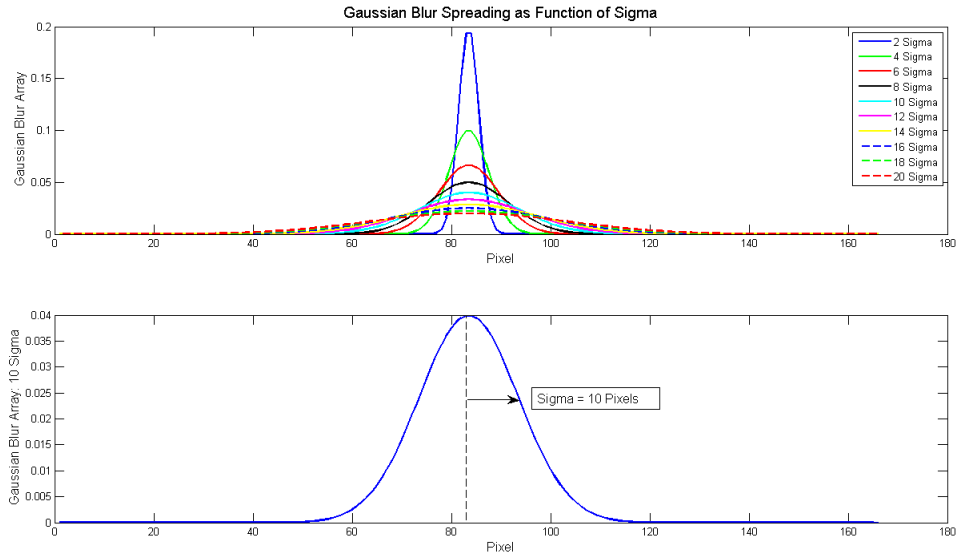

Figure 46: Gaussian Blur as Function of Sigma

The resulting image, I_{GB}, is a result of convolving the original image, I, with the Gaussian Blur filter array in the pixel domain

$$I_{GB}(x_{pix}, y_{pix}) = h(x_{pix}, y_{pix}) * I(x_{pix}, y_{pix}) \tag{4.14}$$

The Gaussian Blur filter is inherently a low-pass filter, thus the primary effect on image edges is to smooth or spread them out over the corresponding Gaussian distribution. The nature of edges, as represented in image processing, is that of an impulse. Thus, energy exists across all spatial frequencies. The Gaussian Blur filter attenuates the higher-frequencies, resulting in the desired spreading effect.

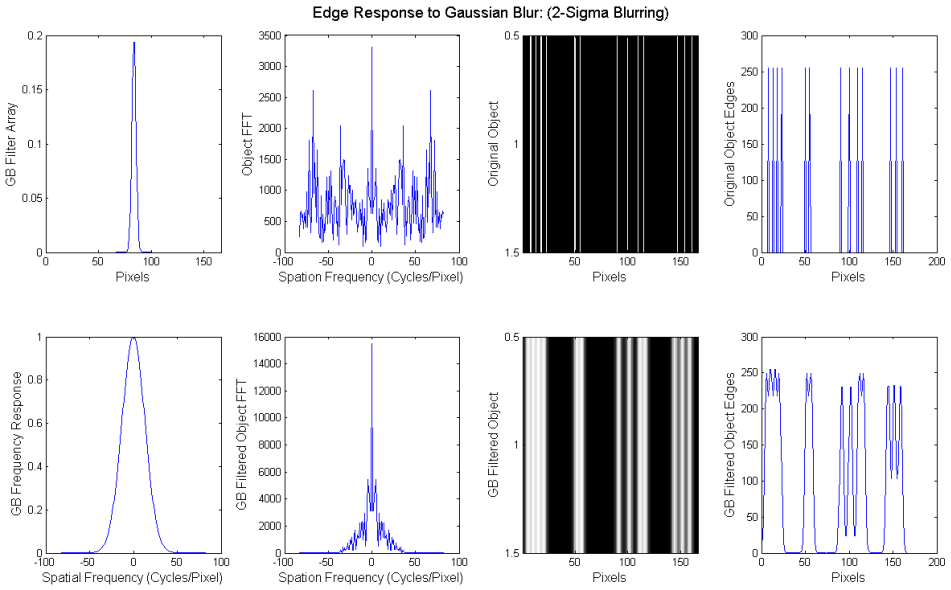

Figure 47: Gaussian Blur Filter Response. (a)-(h) From the top, left to right.

Consider Figure 47(b)-(d), which shows an example edge array and the associated spatial frequency. As expected, there is energy across the entire observable frequency range. Figure 47(f)-(h) shows the resulting image and associated spatial frequency when a 2-sigma Gaussian Blur filter is applied. As expected, the results are lower frequency edges spread across the designed Gaussian distribution. Another important aspect

demonstrated by this example is how distinct edges can become merged or grouped into larger regional edges. This grouping effect is a natural outcome of having too much model uncertainty relative to the spatial separation of edge features.

4.2 Likelihood Determination

The first step in the Bayesian inference approach is to determine the underlying likelihood function of our vision navigation approach. As indicated in Section 4.1, all of the imagery is transformed into edge features, including both the rendered and sensor image, with a Gaussian blur filter applied. The blurred edge feature intensities of the sensor image become the primary measurement for estimation process, in which the difference between the sensor image edge intensities and the rendered images become the basis for the likelihood distribution, as shown in Equation 4.15

$$g_b(x_{pix}, y_{pix}) = I_{GB_S}(x_{pix}, y_{pix}) - I_{GB_R}(x_{pix}, y_{pix}) \tag{4.15}$$

Where I_{GB_S} and I_{GB_R} is the blurred edge intensities of the sensor image, which is a function of the true state (x), and rendered image, which is a function of the estimate state ($\hat{x}$), respectively, as described by Equation 4.8.

The distribution of $g_b(x_{pix}, y_{pix})$, is a function of many factors, including image feature choice, spatial resolution of the rendered state space search, and geometric characteristics of the operational application. Hence, it is critical that each implementation for a particular application be specifically characterized. The resulting underlying likelihood distribution for the aerial refueling application, based on sensor edge features that meet a predefined threshold, is illustrated in Figure 48. The actual pdf

shows the distribution of $g_b\left(x_{pix}^n, x_{pix}^n\right)$ for all blurred edge features, that meet a predefined threshold, where $\left(x_{pix}^n, x_{pix}^n\right)$ represent the pixel locations for those features. The results shown in Figure 48 are the ensemble statistics of g_b taken across the state space within the operational envelope of the application. In the case of the aerial refueling application, it is determined using rendered image, I_{GB_R} specifically selected to minimize the state error between the sensor image and the rendered image state space.

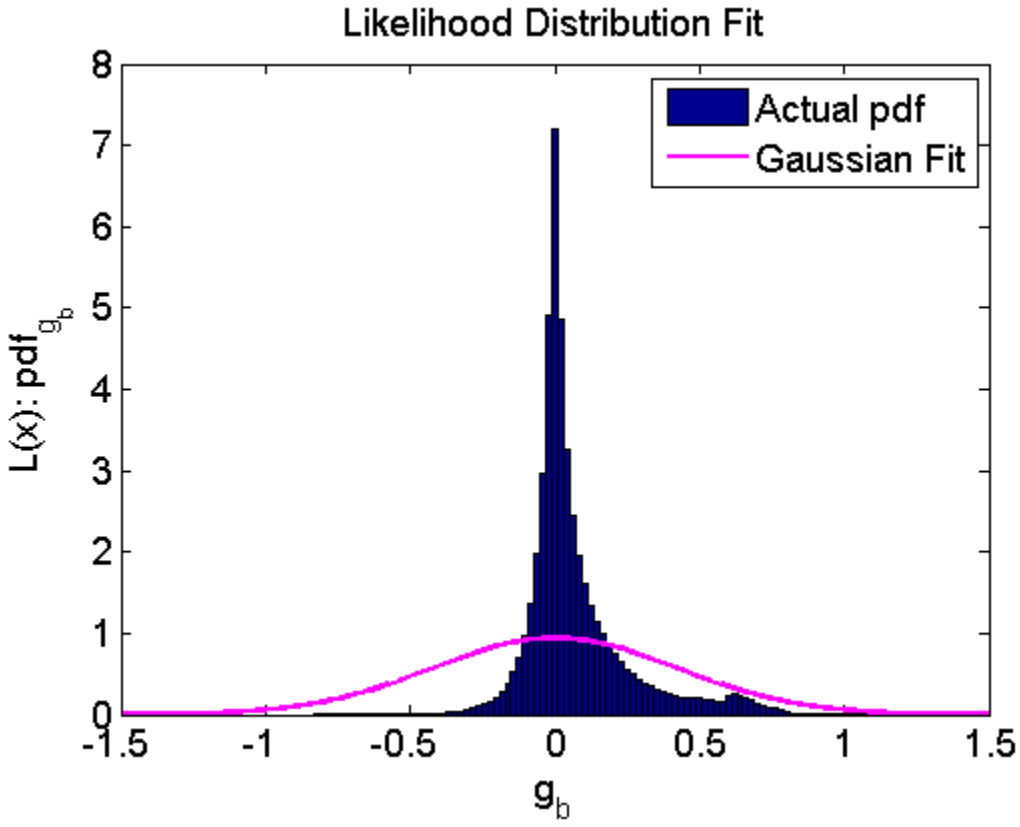

Figure 48: Aerial Refueling Likelihood Distribution

The reason only the rendered reference images that minimize the state error are used, is because that is the purpose of the estimation process, to minimize the estimation error. One of the assumptions of this process is that the true state lies within the rendered search space, hence a minimum estimation error exists within that search space, and the goal of

the estimator is to find that state. Therefore we want to characterize the measurement distribution when this is state error minimization is satisfied.

The distribution of $g_b\left(x_{pix}^n, x_{pix}^n\right)$ closely fits that of a Laplacian distribution, with the exception of a small local maximum near the positive tail. The tail artifact is a result of instances where there are sensor images, but no strong rendered edge features at this location. This occurs for two main reasons. First, the discrete nature of the rendered state space search, in this case 0.5m, results in some outlier cases where the Gaussian blurring is insufficient to make up spatial discrepancies, resulting in slight mismatches between the sensor image and rendered images, even under perfect modeling conditions. Secondly, this research chose to use the presence of sensor images as the constraint in performing any image correspondence. Therefore, this characteristic is only present on the positive tail. Ideally, under no image processing throughput constraints, the state space resolution can be fine enough to eliminate this characteristics, an example of which will be shown in subsequent sections. To account for this tail effect a Gaussian distribution was chosen to over-bound the inherent pdf, as described in Equation 4.16,

$$L_n(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{\frac{-\left(g_b(x_{pix}^n, y_{pix}^n) - \mu\right)^2}{2\sigma^2}} \tag{4.16}$$

where the likelihood, $L_n(x)$, is determined for each of the *n* edge feature pixels that meet the predefined threshold (i.e. edge threshold). The mean (μ) and standard deviation (σ), again are based on the specific state-space search choices discussed above. At this step in the process, the measurement has been realized and the likelihood is now a function of the state space.

4.3 PDF Determination using Bayesian Inference

Once the likelihood is determined, Baye's law can be used to infer information on the state and ultimately the uncertainty associated with that state, as given by

$$p(x|d) = \frac{L(x)p(x)}{\int L(x)p(x)dx} \tag{4.17}$$

The prior, $p(x)$, is assumed non-informative, hence a uniform distribution across the state search space is utilized. However, part of the sensitivity analysis, which is covered in subsequent sections evaluates a more informative prior.

To incorporate the multiple measurements, the likelihoods from each individual pixel measurement is incorporated into the overall likelihood function,

$$L(x) = \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^N e^{-\frac{1}{2\sigma^2}\left\| g_b(x^n_{pix}, x^n_{pix}) - \mu) \right\|^2} = \prod_N L_n(x) \tag{4.18}$$

and subsequently the resulting pdf is determined as

$$p(x|d) = \frac{\left(\frac{1}{\sqrt{2\pi}\sigma}\right)^N e^{-\frac{1}{2\sigma^2}\left\| g_b(x^n_{pix}, x^n_{pix}) - \mu) \right\|^2} p(x)}{\int \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^N e^{-\frac{1}{2\sigma^2}\left\| g_b(x^n_{pix}, x^n_{pix}) - \mu) \right\|^2} p(x)dx} \tag{4.19}$$

4.3.1 *Log-Likelihood*

It has been previously discussed in this research that one of the benefits of going to a pixel edge feature intensity based approach is the abundance of potential independent

measurements available for the inference. Figure 49 demonstrates a possible utilization of the multiple measurements available on a single image for a refueling application. This particular example represents 100 Gaussian-blurred edge features points selected from a simulated EO sensor. However, one of the unexpected outcomes of this multiple measurement approach was dealing with the resulting machine precision issues as the number of measurements used in the Bayesian inference process increased. The machine precision issues arise out of the fact of system noise, the presence of modeling error, and quantization of the rendered reference set state space. All of these contributors result in the same effect, errors in the measurements between the sensor and the rendered image features. The net result is the spread in the measurement likelihood function, $L(x)$, that was observed. For example, if the conditions existed for no sensor noise, perfect model knowledge and negligible state space search quantization, a Dirac function would result for $L(x)$. Since this is not the case, and the fact that likelihood values are less than one, means that as the number of measurements grows, the product of those likelihoods functions, Equation 4.18, will start converging to zero, thus causing machine precision issues when computing the resulting pdf using Bayesian inference, as shown Equation 4.19.

To demonstrate this effect, consider the example scenario in Figure 50, in which a varying number of measurements are used to compute $p(x|d)$. When normal computation approaches are taken the resulting pdf will eventually hit machine precision limits and result in zeros, in this case somewhere between 1000 and 2000 points. However, the log-likelihood approach is successful in computing the posteriori pdf up to

the 5000 data points evaluated. The point at which machine precision begins to effect the overall Bayesian inference process is highly dependent on the application and the nature of the underlying likelihood function.

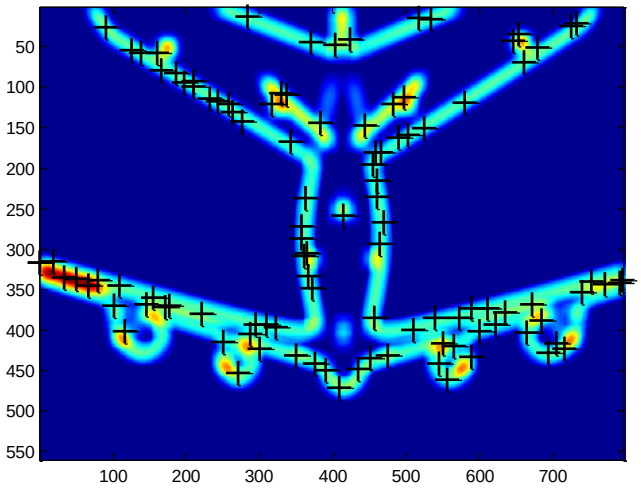

Figure 49: Multiple Measurements Illustration for Aerial Refueling

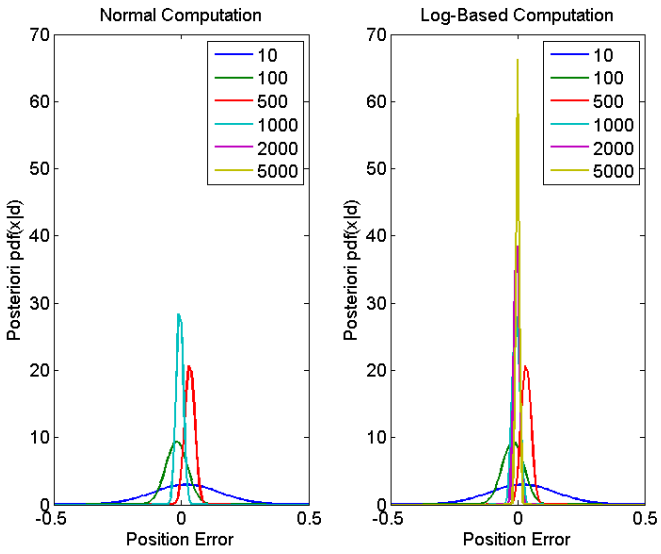

Figure 50: Likelihood Machine Precision Issues

The use of the log-likelihood means that Equation 4.18 and subsequently the computation of the *a posteriori*, $p(d|x)$ turns into a summation, rather than a product, effectively minimizing the effect of machine precision issues.

$$\log\left(L(x)\right) = \sum_{N} \log\left(L_n(x)\right) \tag{4.20}$$

4.4 PL calculation

Once the pdf, $p(d|x)$, is determined all of the information that was able to be inferred from the imaging measurement and image processing is known, and no further information is brought to bear on the state and state uncertainty estimation process. What remains at this point, is using the resulting pdf to determine the appropriate uncertainty bounds.

Equation 4.4 described the governing equation for computing the protection level, once the posteriori pdf is determined, $p(d|x)$. As stated previously, the idea is to compute the minimum state space confidence interval or protection level, centered upon the state estimate, that satisfies the integrity risk requirement. Since the state space is multi-dimensional, so too must the determination of the protection levels. To illustrate this process consider, Figure 51, which shows the 3-dimensional posteriori pdf at two discrete vertical positions within the rendered search space. The top two plots show the result at the vertical state closest to the state estimate, $\hat{x}$, in both a 3-D view (right) and a 2-D contour view (left). The figure also shows the true state and resulting horizontal protection level corresponding to a integrity risk of 10^{-6}. A second view, at a slightly higher vertical rendered state is also shown, and as expected there is less area under the

curve or confidence in these states, as captured by $p(d|x)$. The fact that there is significant pdf is spread across the state space is a direct result of uncertainty in the inference process and this is reflected in larger protection levels to capture this uncertainty.

Although Equation 4.4 describes the process of protection level determination as a function across a continuous function, in practice, this is how it is computed in practice. Due to natural limitations in computer processing throughput, a discrete state search space is utilized. For this research, a somewhat simple approach was taken to the multi-dimensional optimization for the protection level determination. The implementation for this research incrementally increases all axes simultaneously for the protection level until the integrity risk requirement is met. Subsequently, the protection level (PL) computation is discretized,

$$PL(N) \equiv \sum_i \sum_j \sum_k p(x_i, y_j, z_k|d)\, dv \tag{4.21}$$

where $p(x_i, y_j, z_k|d)$ are the values of the *a posteriori* pdf at each of discrete (*i*,*j*,*k*) state space across the search volume, and *dv* is the discretized search volume increment. The state space indices are determined by all discrete states within the defined integral range, as shown for the x-axis and similarly for the other axes,

$$\{i|x_i \in [\max(\hat{x} - Ndx, x_{min}), \min(\hat{x} + Ndx, x_{max})] \tag{4.22}$$

N is a positive natural number and dx is the discritize state resolution, in this case, 0.5m. The final protection level is determined by finding the minimum integral space, or min (N) such that PL(N) is greater than or equal to the integrity requirement.

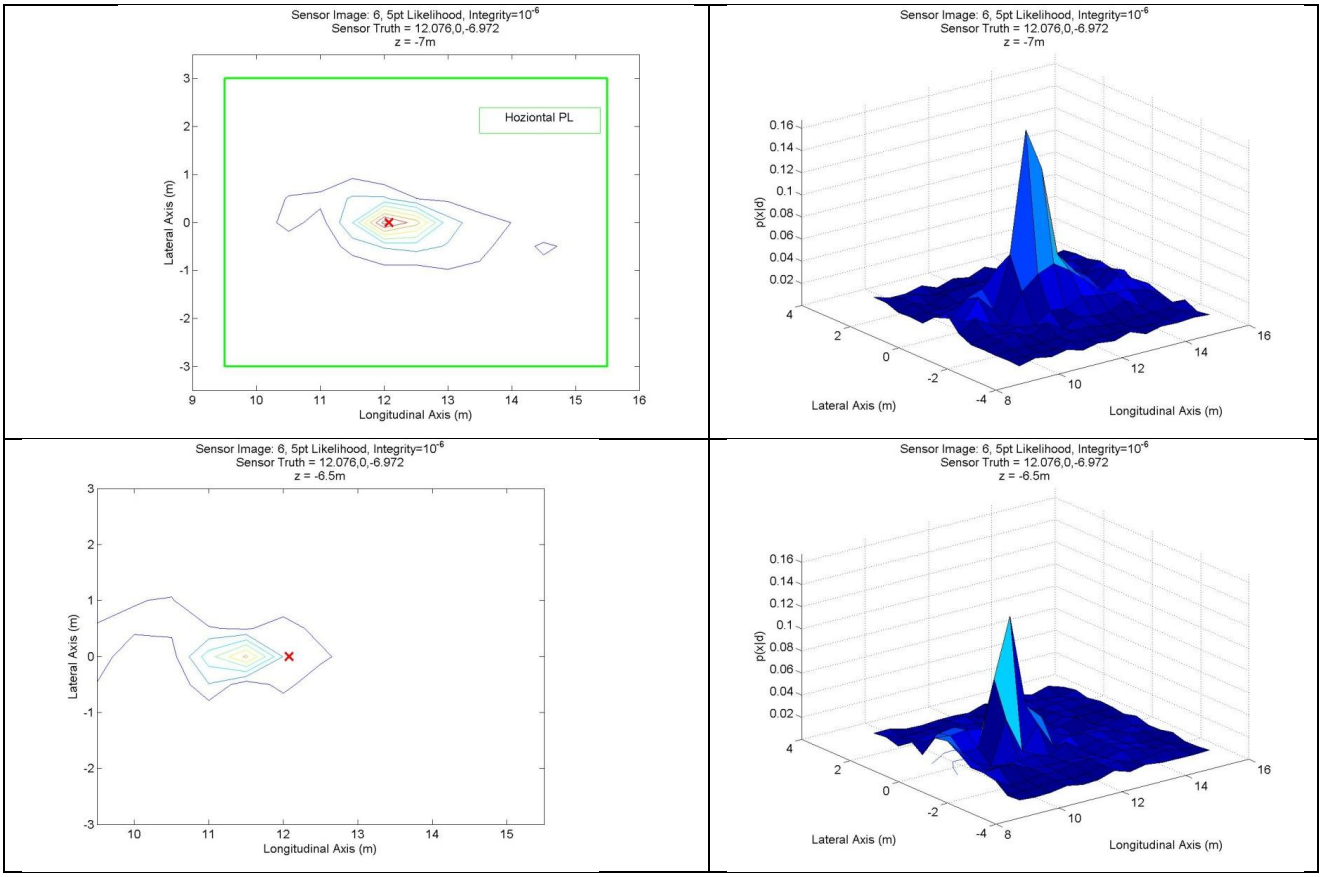

Figure 51: Example Posteriori *PDF show at two (2) distinct vertical states*

Consider Figure 52 in which two levels of integrity risk are evaluated for the same example, 0.5 and 10^{-6}. The underlying pdf, $p(d|x)$, is the same in both cases, and the protection levels are centered about the mode of the $p(d|x)$, which is used as the state estimate. As expected a larger protection level is required for the high integrity requirement of 10^{-6}. The spatial size differences of the protection levels are not proportional to the change in integrity, because most of the probability of $p(d|x)$ is

contained around the mode and diminishes rapidly as the state space diverges from the mode. It can also be observed that the protection level for the 10^{-6} case is skewed in the x-axis. This is a result of the fact that the search space is limited and in this case, hit the lower x-limit in that axis.

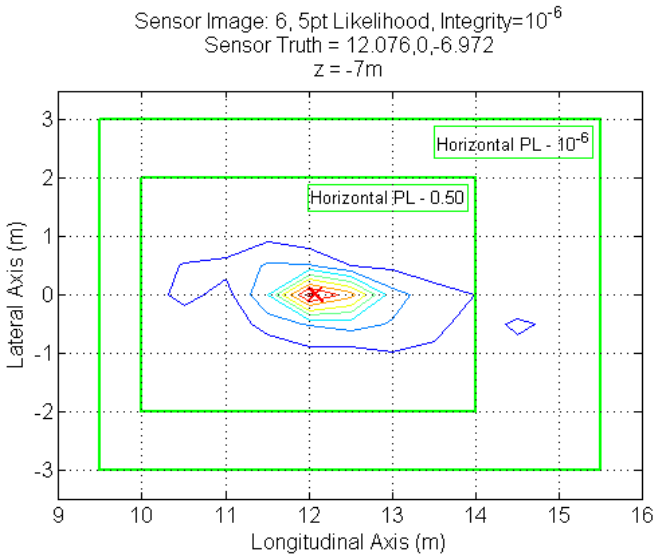

Figure 52: Protection Level Computation Illustration for Different Two Risk Requirements

4.5 Simulation Evaluation

To explore the performance of the Bayesian inference integrity approach, the research turns to an aerial refueling (AR) application within a simulation environment. The simulation environment is nearly identical to that described in Section 3.3, used to demonstrate the integrity monitor approach. However, in this case, there is less interest in demonstrating the detection of discrete instances when the state violates a safety boundary. Rather the simulation for this aspect of the research, and the associated evaluation trajectory, as shown in Figure 53, is intended to validate the approach across the operating range of the application and ensure the protection level's bound the error to

the appropriate probability. The baseline simulated evaluation trajectory included 181 positions across the aerial refueling operating region, some of which included safety violations.

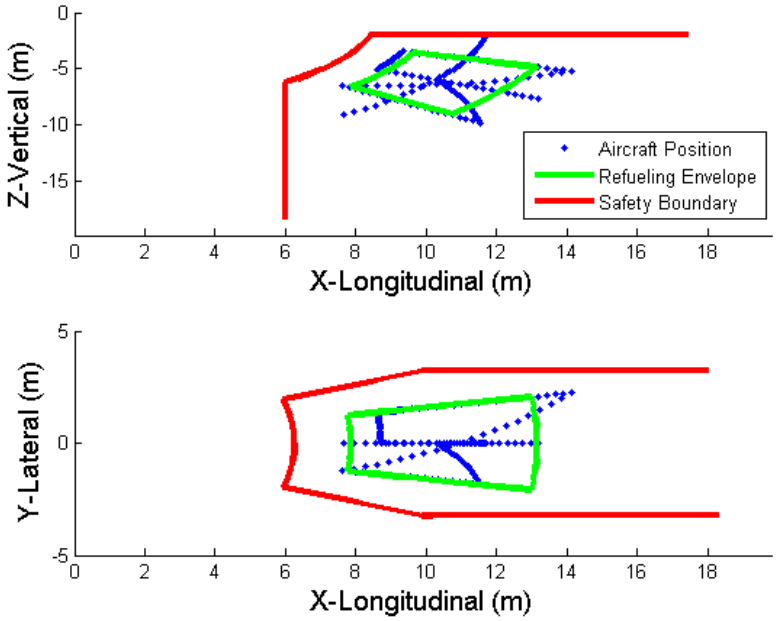

Figure 53: Bayesian Inference Simulation Trajectory

The range of state values that the evaluation trajectory represents, with respect to the refueling tanker, is captured in Table 4-1. Additionally, this table captures the span of the state search space of the rendered image database. Based on this state space range and spatial resolution, the result is a full rendered reference image database of 53,366 images for possible use in the image correspondence process, depending on the initial estimate of the relative position state. However, in order to keep the image processing and computer simulation to a tractable level, only a subset of the image reference database is utilized. Offline analysis provided guidance on the size of the subset rendered reference database

and it was concluded that having rendered images a minimum of +/-6 meters in any axis of the true state of the refueling aircraft / vision sensor, or 2197 images, would be sufficient to capture the statistics up to the desired 10^{-6} integrity risk. However, if additional integrity performance beyond this is desired, additional rendered state space would potentially be required.

Table 4-1: Rendered Image State Space

Relative Position State	x-axis (m)	y-axis (m)	z-axis (m)	Resolution (m)
True State Space	[7.5, 14.5]	[-2,2.5]	[-10, -2]	N/A
Search State Space	[1.5, 20.5]	[-8,8.5]	[-16, 4]	0.5

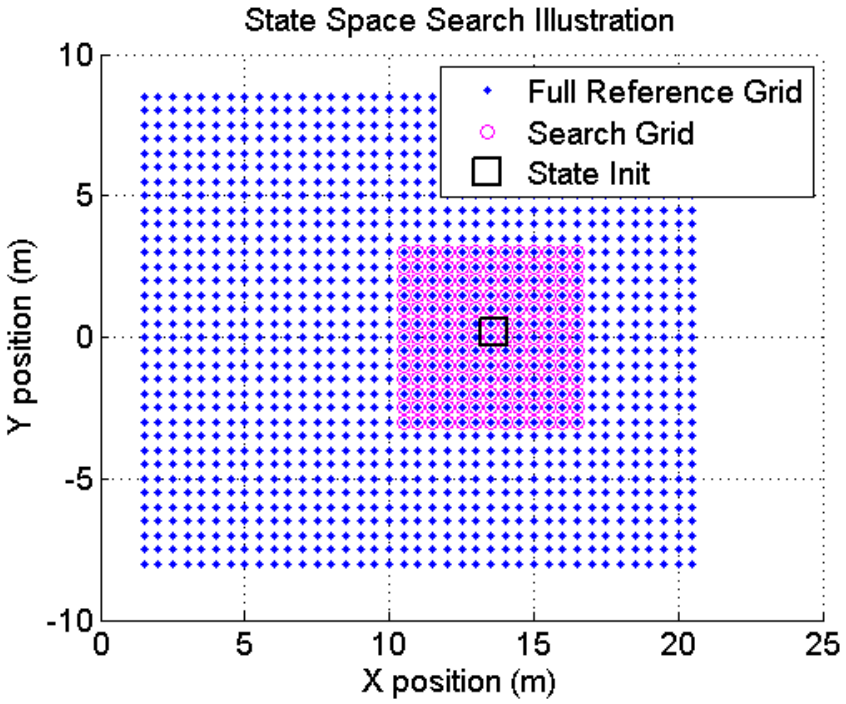

Figure 54: Bayesian Inference Dynamic Reference Database Illustration

This initial state is based on an erroneous true state to represent a projected estimate from the previous state estimate,

$$\hat{x}_{INIT} = N(x_{true}, \sigma_{proj}) \tag{4.23}$$

where σ_{proj} is set to 0.5m for the simulation. Figure 54 illustrates this rendered reference subset selection process for the x-y axis; However as indicated, this subset selection is applied each time epoch to all axes in the simulation.

For the baseline analysis, a integrity risk threshold of 10^{-6} was selected and the Bayesian inference was based on using 100 edge feature measurements for each sensor image. Figure 55 shows the results of this process for each time epoch across the entire simulated trajectory. The results are broken out in terms of horizontal and vertical error components, a choice commonly made when evaluating navigation systems. The figure shows both the navigation error and corresponding protection level / confidence bound. Since the integrity requirement is set so high, relative to the number of time epochs, no violations of the protection levels are expected, which is in fact the case. The estimated relative position state, $\hat{x}$, is chosen by selecting the mode of the pdf, $p(d|x)$, and $\hat{x}$ also serves as the center point of the protection level computation, as indicated previously. The quantization in the protection levels, a result of the rendered image spatial quantization, is easily observed in these results. One important characteristic of these results is somewhat weak correlation between navigation error and protection level. This is not entirely unexpected, since the intent of the protection level is to capture the uncertainty in the navigation estimation process, which is not always directly related to accuracy.

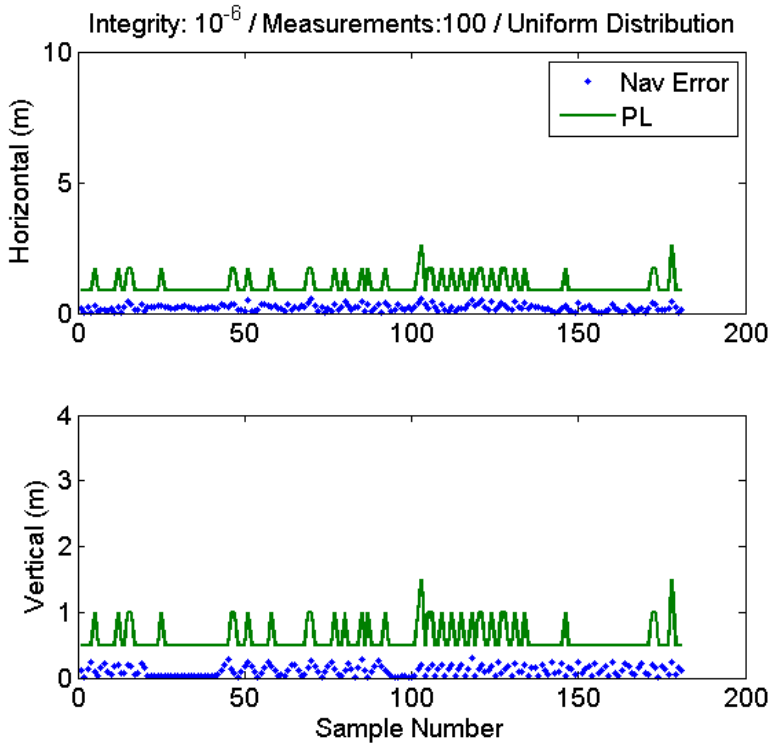

Figure 55: Protection Level Error Bounding (10^{-6} Integrity Risk, 100 Likelihood/Measurement Pts)

To explore this relationship between accuracy and integrity, a single dimensional (x-axis) test was developed. In this test a single edge feature point was strategically selected to highlight the case of ambiguity or uncertainty in the vision navigation estimation process. Figure 56 illustrates this case, in which the measurement point on the sensor image is selected as one of the trailing edge points of the wing. The hypothesis being that moving along the x-axis would highlight the ambiguity of the position state when either the leading edge or trailing edge of the rendered reference set corresponds to this measurement point.

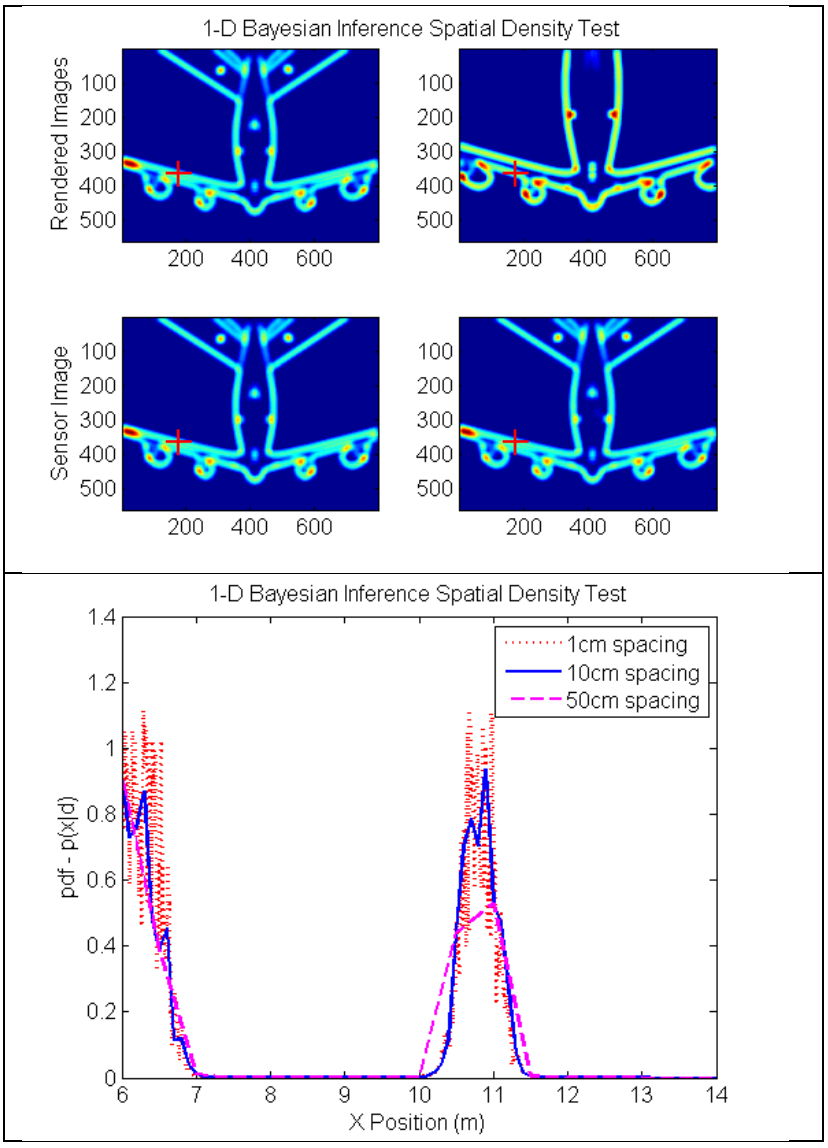

Figure 56: 1-D Bayesian Inference Spatial Density Test

The lower image plots in Figure 56 show the sensor image and the corresponding measurement location. The upper plots show the rendered images that correspond to the peak correspondences for the leading and trailing edge, right and left figures, respectively. The bottom plot shows the resulting pdf, $p(d|x)$, across the x-axis state space. The first peak represents the leading edge and the second peak represents the trailing edge. It is obvious, by looking at the similarity in the images, that the true state is

represented most accurately when the sensor image feature measurement (red cross) aligns with the trailing edge of the rendered image. This is also evident in the peaks of $p(d|x)$ for which the trailing edge peak has the highest pdf value or likelihood and is selected as the navigation solution. By choosing this peak as the state estimate, the error is minimized, and more important to this research $p(d|x)$ captures the uncertainty in the process. Hence, when protection levels are to be computed, a very large state space would be required to capture the necessary uncertainty in $p(d|x)$. This example highlights many of the concepts at work with the Bayesian inference process and is a perfect example of the showing the distinction between accuracy, in this case the estimation error is very low, and integrity or uncertainty, which is very high under these conditions. While this was a simple example, the use of multiple measurements across the sensor image is quite effective at removing this ambiguity, but the ability to address this ambiguity is something that is very scene dependent.

In addition to ambiguity and uncertainty aspects, the lower subplot of Figure 56 also shows the results of an analysis performed to evaluate the impact on the spatial resolution and quantization effects of the rendered reference set state space. To explore this aspect of the process, three different spatial resolutions along the 1-D x-axis were evaluated; 1cm, 10cm and 50cm. All three show similar trends and capture the expected uncertainty; however, it does appear that some information is lost at the 50cm spacing that is captured by the higher resolution reference sets.

This could cause potential issues in the navigation error; however, further analysis showed that impact to the integrity protection level bounds resulted in almost negligible performance difference. In contrast, the 1cm case demonstrates a fair amount of high

frequency noise components in the areas of interest. This was traced to artifacts in the simulation rendering process, a limitation causing small pixel noise in the rendered imagery and nothing systematic in the overall process.

4.6 Sensitivity Analysis

In addition to the baseline evaluation, presented above, additional sensitivity studies were performed to evaluate the various trade-offs related to the vision-based Bayesian inference process. These sensitivity evaluations focused on some of the previously discussed algorithmic considerations such as number of likelihood measurements utilized, integrity risk threshold, and informative versus non-informative prior information. It will be demonstrated that these algorithmic consideration come down to a performance trade-off between integrity and continuity, as previously defined in Section 1.4. This is the classic performance balance that almost all navigation and other "service-like" systems face, where you can gain safety by improving your integrity risk, but this almost always comes at a cost of continuity in the form of alerts to the user that the system should no longer be used. Although continuity is not explicitly evaluated in this research, the general trend can be inferred by the size of the protection level, since larger protection levels indicate higher uncertainty, and likely a higher probability the system is deemed unusable.

4.6.1 *Likelihood Measurement Number*

One of the most beneficial attributes of imaging sensors is the density of measurements available for each image taken at a single instant in time. Modern imagery sensors that contain tens of millions of pixel measurements are common today, even in

our daily consumer electronics, such as phones, and watches, and are almost standard on unmanned vehicles of all classes. Although not every pixel can be treated as independent, due to digital imaging design and processing of the sensor itself, the access to large measurement sets are ripe for exploitation in this process. Figure 57 illustrates the potential large trade-space available for the use of multiple measurement features for the aerial refueling application, even for a low resolution image of 800x600. Example of 20 and 100 measurement sets are show in Figure 57. This sensitivity analysis is intended to explore the impact of the number of measurements on the Bayesian inference integrity determination results. Section 4.3.1 touched upon some of the effects of multiple measurements, and it can be easily observed in Figure 50 that use of simultaneous multiple measurements in the Bayesian inference process can improve the uncertainty in the *a posteriori* pdf.

To explore this sensitivity, the Bayesian inference process was performed across a wide range of number of measurement, 5 through 100. The number of measurements has two direct impacts on the performance, impacting both the navigation error and the spatial size of the protection level, as shown in Figure 58.

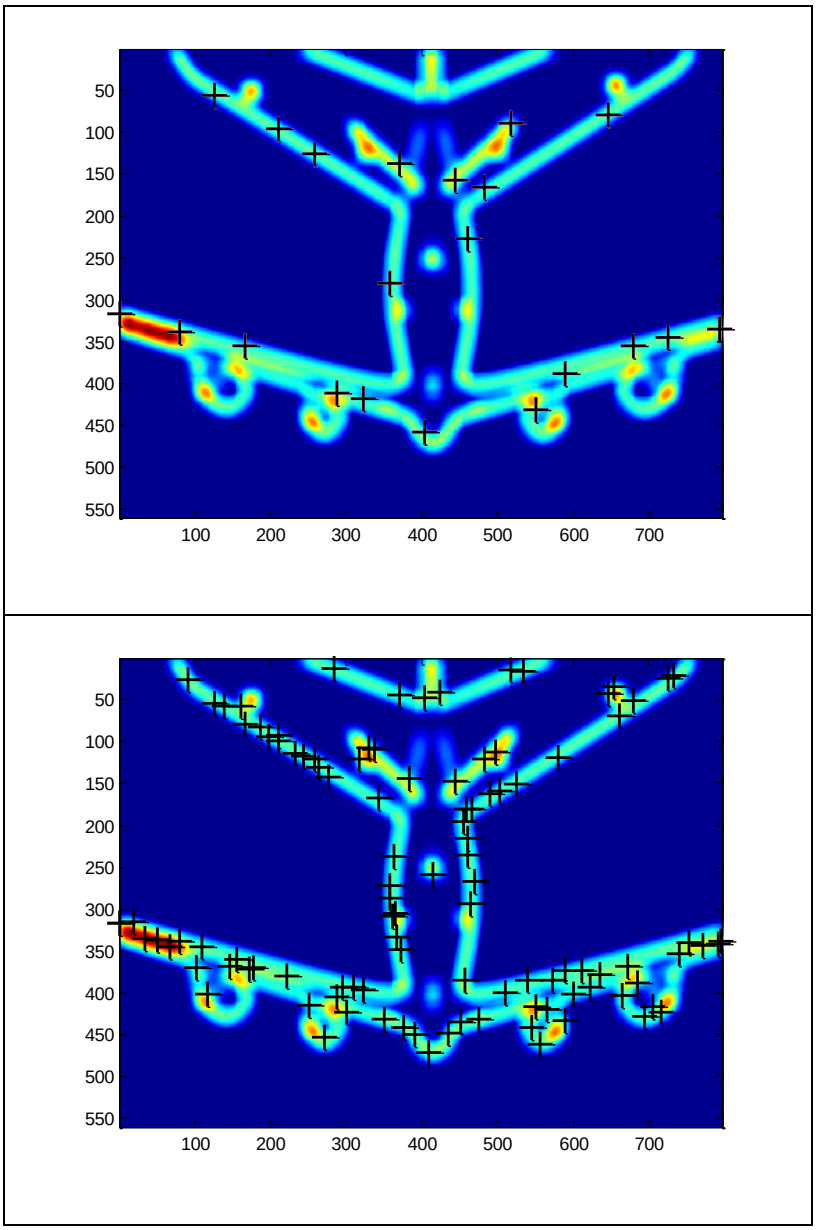

Figure 57: Example Image Showing Multiple Likelihood Measurements: (Top) 20 Measurements, (Bottom) 100 Measurements

The impact to the spatial size of the protection level is most striking, being reduced by almost a factor of 10, in both the horizontal and vertical axes, when going from 20 measurements to 100 measurements utilized in the likelihood function. It is clear from Figure 58 that the lower limit on the protection level is being constrained by the 0.5m

spatial resolution of the rendered reference set, so it is not evident how tight the convergence between the protection level bound and the navigation error can be made if there were no such constraint. Additionally, the navigation error is improved marginally as additional measurements are considered, and both cases exhibit possible performance limitations as a result of the reference set quantization. These search space quantization limitations are illustrated in by flat-line errors in the vertical dimension. Figure 59 shows the relationship across the entire sensitivity range, plotting the computational difference between the protection level and the navigation error. What clear is that there is a striking break point for the aerial refueling application between the 20 – 40 measurement point, at which point it settles out to what looks like the spatial resolution limit as the measurements increase. This trend is very likely to be application specific, since it is being driven by the geometric ambiguity of the world model.

It is clear from Figure 58 and Figure 59 that including more measurements in the Bayesian inference process has a much more impactful effect on the integrity, in terms of the size of protection levels, than on accuracy of the navigation solution, which is negligibly effected.

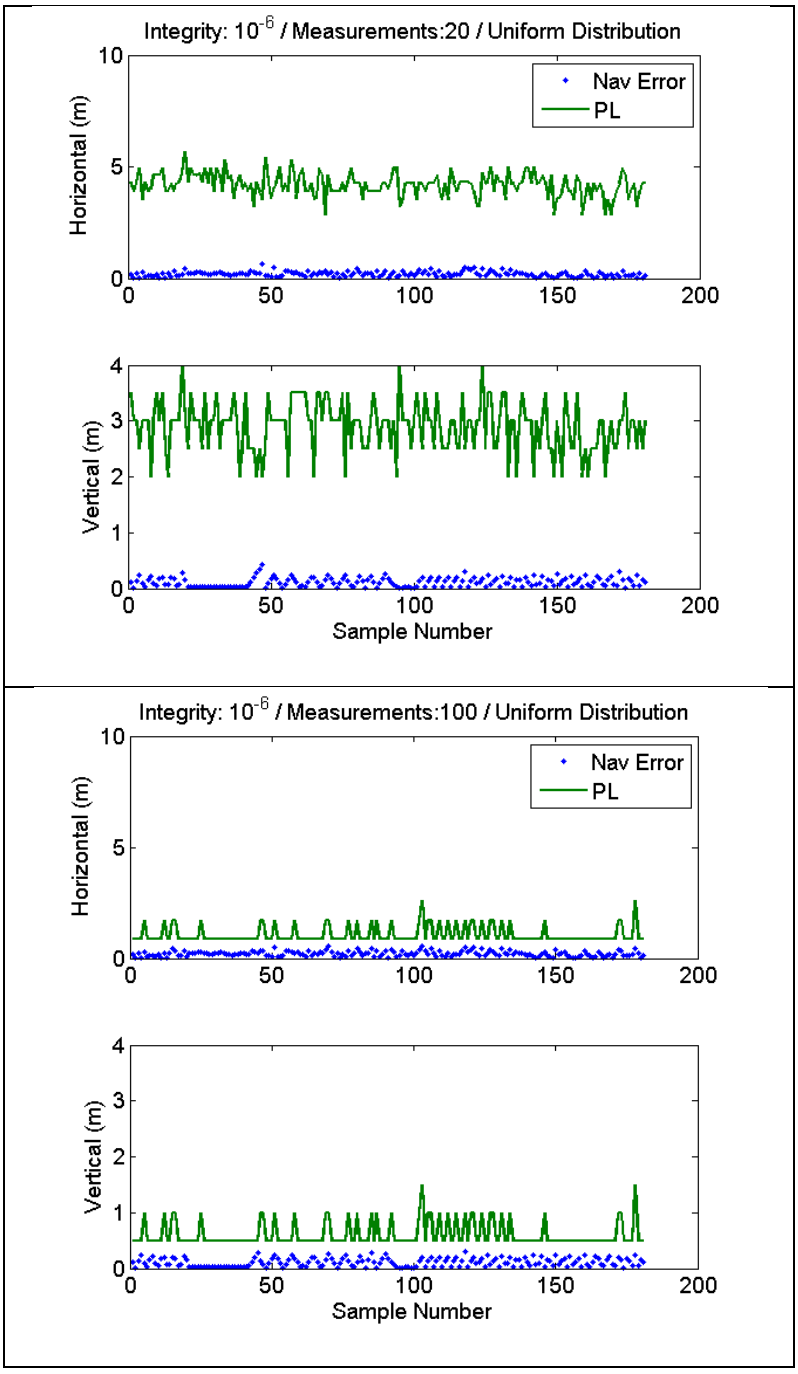

Figure 58: Comparison of Bayesian Inference between 20 vs 100 Measurements

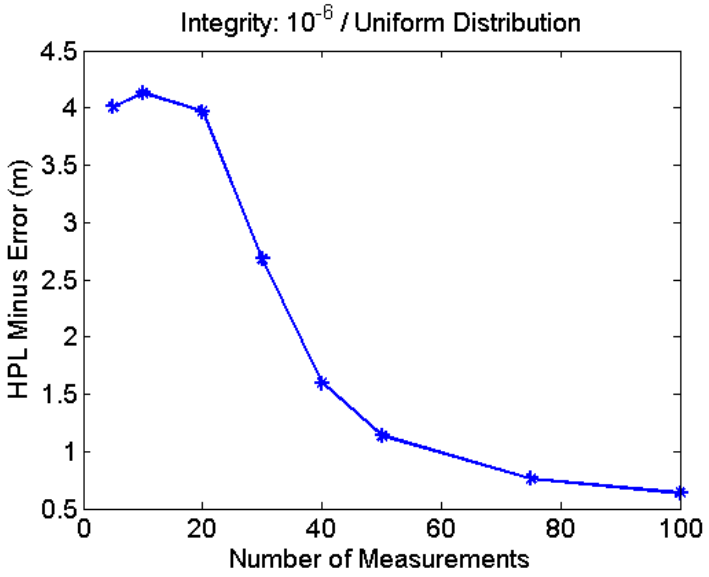

Figure 59: Protection Level Bound minus Navigation Error as a Function of Measurement Number

4.6.2 *Impact of Integrity Requirement*

Each operational application requiring high-precision navigation have potentially unique performance requirements, leading to operational specific trade-offs between integrity and continuity. Not all applications require extremely high levels of integrity requirements and can perform with less confidence in solution. This analysis examines the impact of changing the integrity risk requirements for the Bayesian inference approach. The protection level requirement is really an after the fact operation, once the Bayesian inference process has determined the *a posteriori* pdf. Therefore, the pdf and the corresponding navigation accuracy is not impacted by the integrity risk requirement – only the size of protection level's. This can be easily inferred by considering the governing equations of the protection level, Equation 4.21, and illustrated by Figure 60, which shows the results for vastly different integrity requirements, that of 10^{-6} and 0.05.

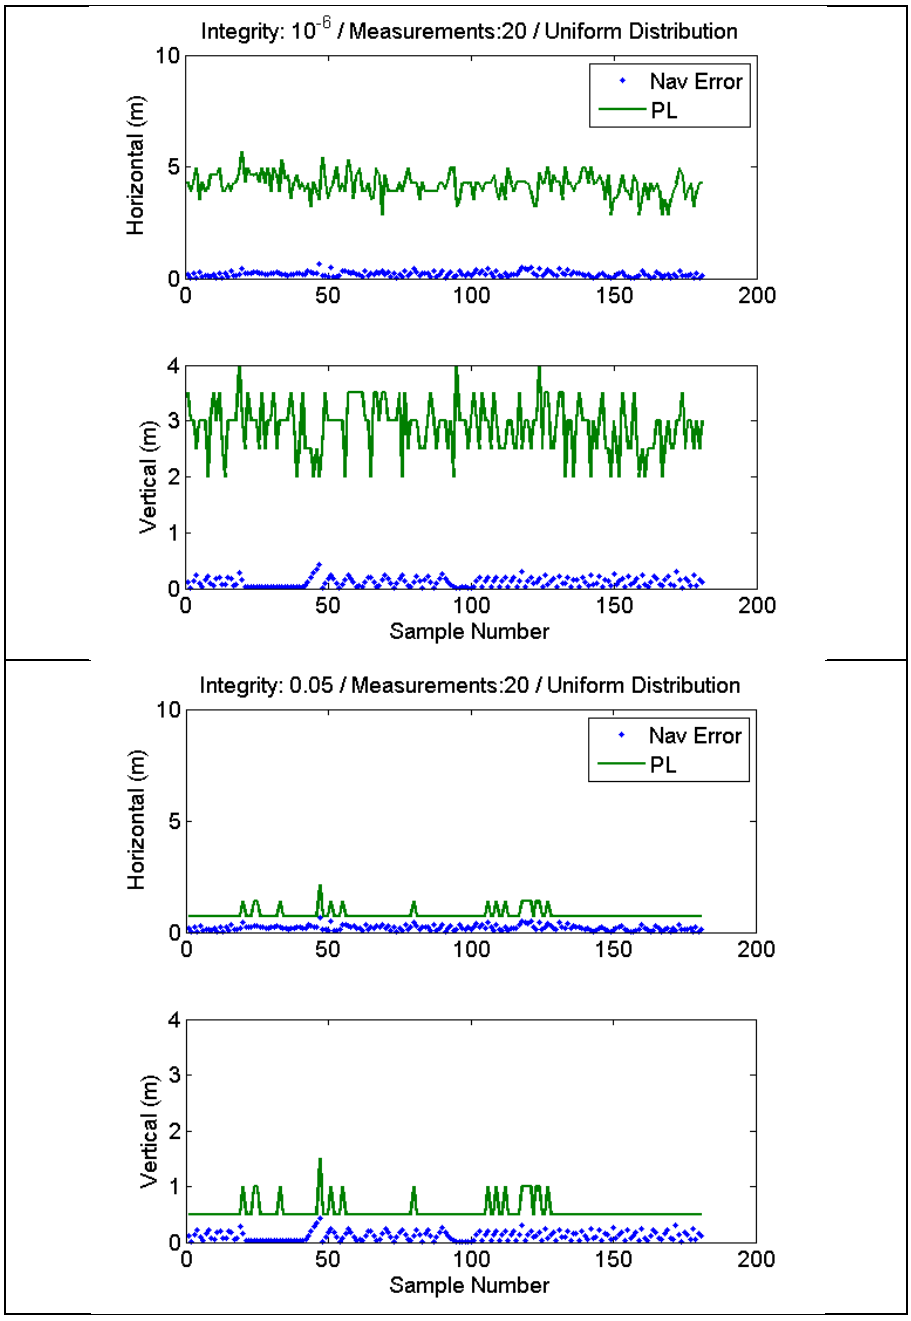

Figure 60: Impact on Integrity Risk Requirement on Protection Level

As expected, Figure 60 demonstrates that the integrity risk requirement has a direct impact on the protection level bound, where the larger integrity requirement, or lower integrity risk requirement, the larger the protection level bound. The size of the protection

level can subsequently have a significant impact on the overall continuity and usability of the system as a whole. This example also demonstrates that the integrity risk requirement has no impact on the navigation error, which remains unchanged between the 10^{-6} and 0.05 scenarios.

The last major observation of this analysis is the lack of dynamics in the protection level for the 0.05 (95%) integrity case. This indicates that most of the probability of the pdf (more than 95%) is contained within the rendered reference set quantization distance of 0.5m.

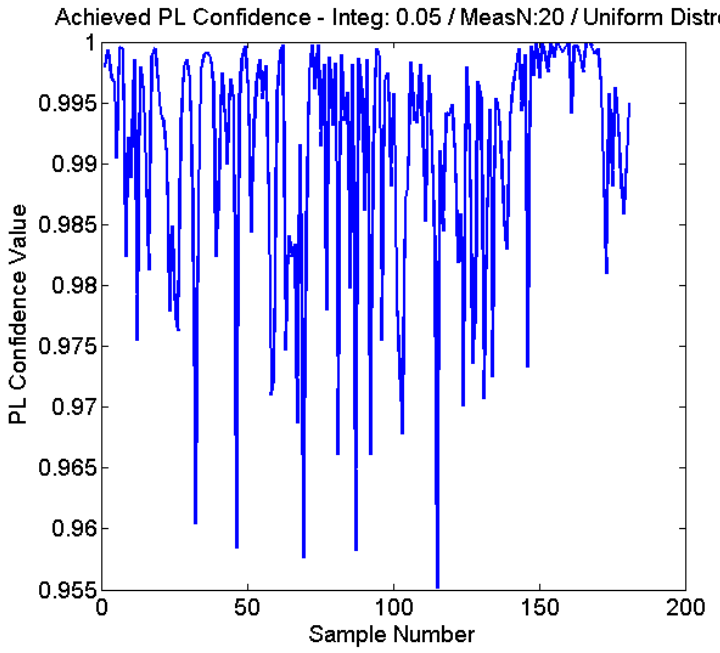

Figure 61: Achieved PL Confidence Probability for 0.05 Integrity Risk Requirement

To further illustrate this point, consider Figure 61 which shows the realized probability bounded by the protection level. Even thought the integrity risk requirement was set to 0.05, or a containment of 0.95, the figure shows that in a vast majority of the cases, a

much larger probability is contained by the protection level, with many cases being greater than 0.99. Again, this characteristics is a result of the 0.5m spatial resolution in the rendered reference set, hence causing similar quantization in the resulting pdf, $p(d|x)$. Therefore, when integrating $p(d|x)$ for the protection level, getting the precise integrity risk requirement becomes problematic. There are several ways this characteristic can be addressed, if it poses a performance issue. The most direct method is to increase the spatial resolution of the reference set, or potentially interpolating $p(d|x)$ after the fact. However, when performing interpolation, it is important to ensure that $p(d|x)$ maintains the properties of a pdf (i.e., integrates to unity) and even more importantly, that the interpolation regions are representative of the true uncertainty.

4.6.3 *Informative versus non-informative Prior*

The initial analysis of the Bayesian inference approach relied on a non-informative prior, $p(x)$, which was results in a constant value for $p(x)$across the search space, and computed as

$$p(x) = \frac{1}{V}, \forall x \tag{4.24}$$

where V is the volume of the state search space. However, a case can be made that some informative prior can be determined based on propagating previous navigation estimates in a manner similar to Equation 4.23, which was limited only to determining the search space. To test the case where a more informative prior can be assumed, an evaluation was performed with a Gaussian distribution of the form,

$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2\sigma^2}\|x - x_{INT})\|^2} \tag{4.25}$$

with a mean of x_{INT}, which is the center of the search space, and with the same σ. The use of this informative prior will effectively increase the values of the *a posteriori* pdf, $p(d|x)$, closest to x_{INT}. The resulting pdf, $p(d|x)$ now becomes a blend between the prior and the likelihood, $L(x)$, according to Equation 4.17 and driven by the uncertainty in each.

The purpose of this analysis is to evaluate the impact the use of informative versus non-informative priors can have on the overall system performance. Figure 62 shows the results from the simulation for each of these cases and shows in general, the use of informative priors does not significantly impact the protection level or navigation error, however there are some interesting findings. The hypothesis going into this analysis was that with additional informative knowledge of the state (x), that the overall performance would be better. However, the results actually show the contrary, that the overall navigation performance is slightly better with non-informative $p(x)$, where as the Gaussian prior case shows small sporadic error divergences, as compared to the uniform case.

To understand the cause of this characteristic, consider Figure 63, which shows the navigation error and the error associated with the mean of the prior, x_{INT}. The figure shows a correlation between the cases where the Gaussian navigation solution diverges and where the prior distribution has a excessively large mean, particularly in the vertical axis. This indicates, that statistically, the prior can draw off the mode p(d|x) sufficiently to change the estimated position within the search space.

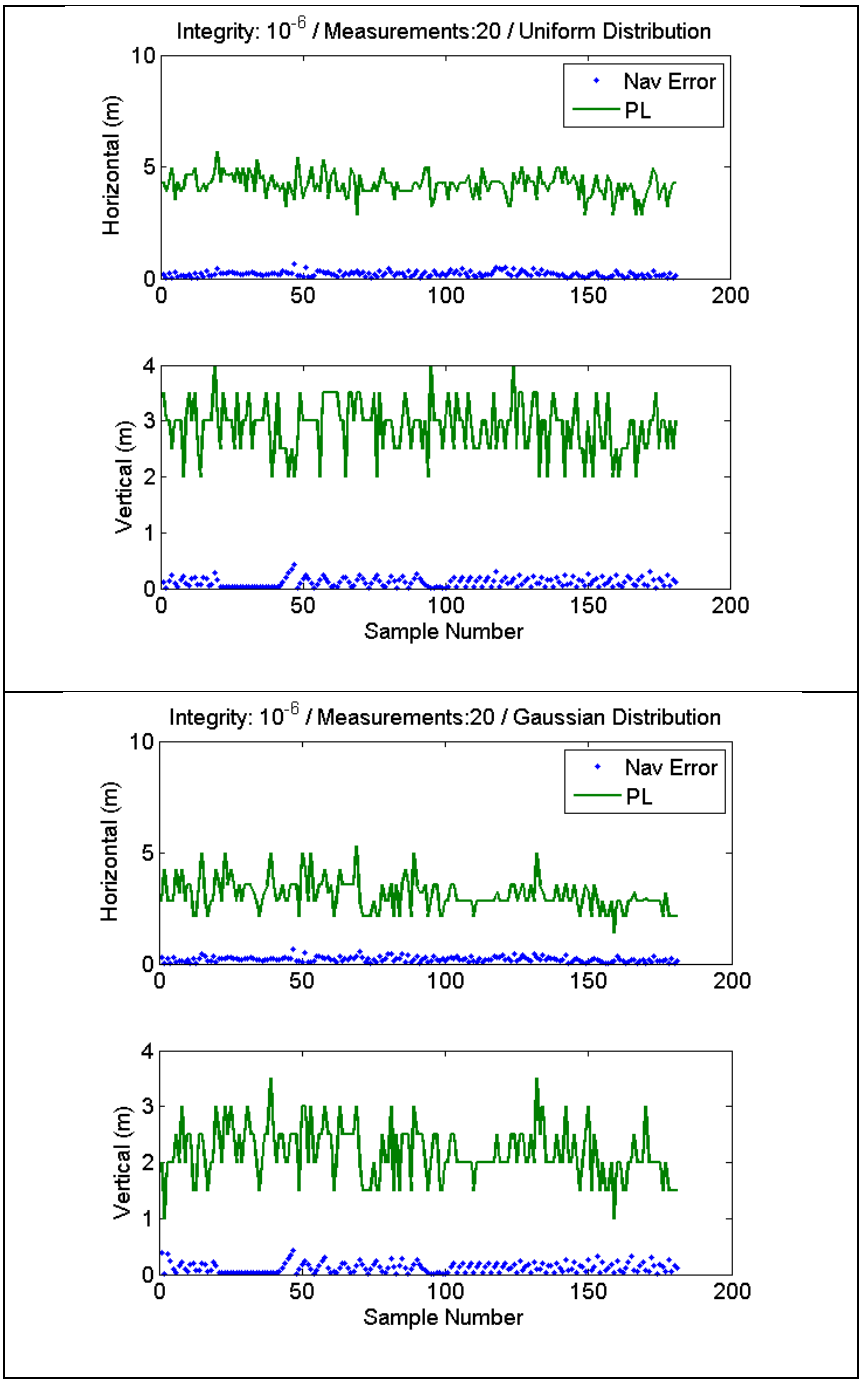

Figure 62: Uniform versus Gaussian Prior Bayesian Inference Performance

Again, the net effect is small, but non-negligible, in terms of navigation accuracy. The effect on the protection level is not as observable in Figure 62, and the protection level values are almost identical between the Gaussian and uniform distribution.

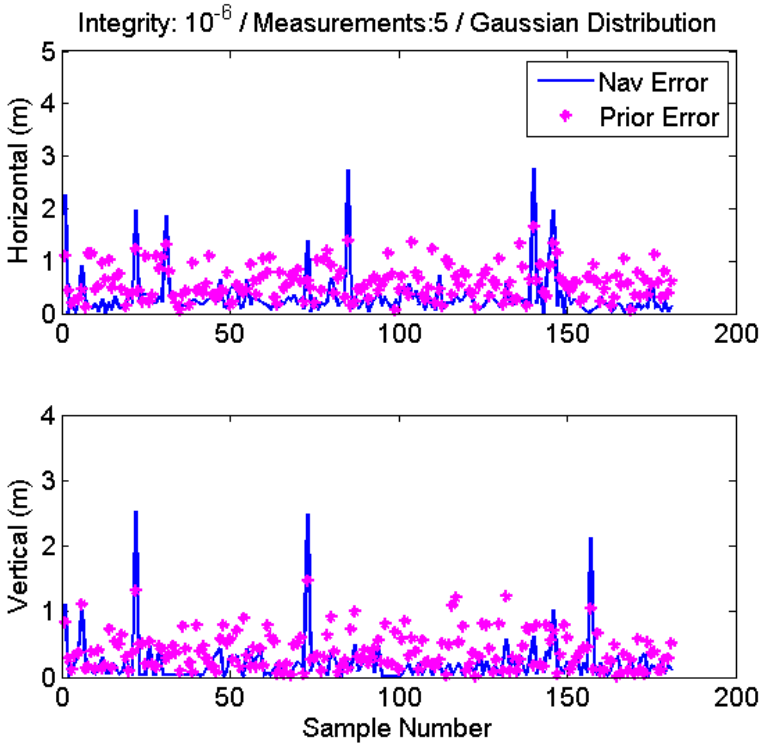

Figure 63: Horizontal Error Performance vs Gaussian Prior Error

4.7 Summary

This chapter investigated a Bayesian inference approach for determining integrity for a vision-based precision relative navigation system. Like the previous Chapter, the research posed the relative navigation integrity problem within the context of an aerial refueling application and using image rendering, where a imaging sensor and high fidelity 3-D model is utilized. The chapter provided a detailed description of the Bayesian inference approach by expanding on work by Well's [51], which used a Bayesian

approach for a maximum likelihood estimation (MLE) for a localization application. This research showed how this can be generalized for a full pdf inference at the pixel level for a vision navigation application, including the computation of the protection levels that can be used for integrity purposes.

The research demonstrated the proposed Bayesian inference integrity approach within a simulation environment. It was shown how the use of multiple image feature measurements can improving the uncertainty in the navigation solution, and subsequent size of the protection levels. In addition, the impact that integrity risk requirements and informative versus non-informative prior information was investigated. The aggregate results show that the Bayesian inference approach is a robust and flexible method of providing fault-free integrity for a vision based navigation system.

5 Conclusions

This dissertation addresses the lack of quantitative integrity approaches for vision navigation by developing two distinct integrity determination techniques using image rendering correspondence. This chapter describes the findings from this research, in particular the major conclusions and touches upon possible future work. Finally, a brief overall summary of this dissertation is presented.

5.1 Conclusions

This research developed two frameworks, or approaches, for determining integrity for vision navigation in a compatible format with existing navigation systems, in the form of a discrete monitor or in the form of a real-time computed protection level bound on the relative state estimate. System integrity includes many aspects, this research the focused on the fault-free, algorithmic aspects of the vision navigation process.

The first contribution of this dissertation was demonstrating the feasibility of a discrete vision-aided integrity monitor. This vision-based navigation integrity monitor generalized the concept of integrity in terms of operating and alerting regions. The integrity monitor approach was based on the concept that most navigation operations have well-defined regions that require a certain navigation performance, whether the operation is based around a glide-slope, a formation flight position, or even a flight path clearance. The concept of this vision navigation integrity monitor was to detect large divergences from these operating regions, that can become safety risks.

As part of this research an aerial refueling simulation was developed and using image rendering correspondence, where a imaging sensor and high fidelity 3-D model is utilized. The approach was an expansion of work by Weaver [10], in which image rendering was shown as a feasible approach for navigation estimation. This research expanded some of Weaver's techniques, such as the use of edge features with SSD as a image correspondence metric, and which were subsequently demonstrated in flight test [15]. This research showed that image rendering processing can be utilized for integrity purposes in addition to navigation estimation. Two different image feature approaches were evaluated, that of edge features and image silhouettes, for which the research showed that 10^{-3} to 10^{-5} level of integrity monitoring is attainable. Additionally, this research showed, in the form of ROC curves, the trade-space between detection probability (integrity) and false-alarms (continuity). The research demonstrated that the proposed integrity monitor was robust against several effects, including lens distortions, spatial rendering quantization and reductions in pixel resolutions. Finally, the use of a ratio-test approach for the integrity monitor, which indicated that another order of magnitude of integrity monitoring performance was realizable.

Another major contribution this research presented was the development a second methodology for determining integrity for a vision navigation system, based on a Bayesian inference approach. Instead of a discrete detection monitor, this approach determined real-time protection levels that bound the navigation error to a user configurable integrity risk requirement. The use of protection levels, mimics that of high precision GNSS navigation systems, such as LAAS and terrestrial navigation systems.

The Bayesian inference approach was inspired by work by Well's [51], which used a Bayesian approach for a maximum likelihood estimation (MLE) for a localization application, in which only the mode of the likelihood function was desired. This research showed how the Bayesian framework can be generalized for a full pdf inference operating at the pixel edge feature level. These innovations allowed for successfully demonstrating the error bounding of a simulated aerial refueling operation up to a 10^{-6} integrity risk threshold. A critical component of this approach is availability and utilization of multiple measurements for a single sensor image. The research showed how the use of multiple measurements was a power factor in reducing the uncertainty in the vision navigation estimation process and providing greater levels of integrity without sacrificing continuity performance. Along that line, the research investigated the algorithmic trade-space between safety (i.e., integrity risk) and operational alerts (continuity risk). Various factors presented within this trade-space include integrity risk threshold, number of measurements, and informative versus non-informative prior information. The key finding is that there are choices the user can make to improve the integrity of the system, such as utilization of more measurements or increasing the integrity risk threshold directly. These options will provide a useful flexibility for future utilization of these integrity determination techniques.

5.2 Future Work

This research provided a good foundation for developing means for determining safety critical integrity for vision navigation system, and although this work has

demonstrated the success of these approaches in a simulated environment, additional work is required to continue maturing these techniques.

The primary means for evaluating the vision navigation integrity algorithms presented in this dissertation was through the use of simulation. Although these simulations proved very valuable to develop these methods, further analysis is required to understand the impact of real-world considerations. As with all sensors and navigation systems, vision-based navigation solutions have a real-world challenges; some unique to imaging sensors, such as lighting conditions; and others are common to many navigation and estimation approaches, such as model errors and real-time processing considerations. Not surprisingly, lighting effects can be one of the biggest challenges when working with vision sensors as a primary means for navigation. Further research is needed to evaluate these vision integrity techniques under different lighting conditions. There has been much work to develop image processing approaches that are somewhat resistant to lighting effects by transforming images to feature space (e.g. SIFT, edges, etc...), however this research did not explore these aspects in any sort of rigorous way.

This research relied on *a priori* models to perform image correspondence with the sensor imagery. The use of models is by no means unique to these research approaches, and in fact many estimation and navigation algorithms rely heavily on the use of models in this manner, Kalman filter being a obvious example; however, as with all these algorithms, this makes them susceptible to model mismatching. In [15], some limited evaluations showed somewhat minimal effects from small model mismatching on the overall vision navigation process. Additionally, the use of the Gaussian blur filtering can

be used to mitigate some of these effects, however a more thorough investigation is likely required to truly understand the impact and limits of this effect.

The possible combined effects of lighting and model mismatching conditions can occur where image features can be present in the sensor image that are not present in the rendered image, or vice versa. These mismatches can cause potential issues in the likelihood determinations, critical to both integrity approaches presented. A possible impact to this to the Bayesian inference approach would be a extremely low likelihood for a measurement or several measurements for a rendered image that is closely aligned to the true state. These small likelihoods can lower the overall likelihood value for that estimated state and possibly effect overall performance. A possible way to mitigate this would be to set a minimum floor in the underlying likelihood functions, so that under these conditions the overall impact would be constrained. This type of scenario is quite likely in operational use of vision navigation systems and further research is warranted to explore this.

Lastly, additional investigation is warranted to evaluate how these techniques can apply to other operations. The evaluation for this research was based on a refueling operations , hover the underlying theories should apply to the broader class of navigation estimation challenges. To adapt the research to other applications, it is likely that one of the most critical aspects would be the derivation of the underlying image correspondence distribution. Both the integrity monitor and Bayesian inference approaches relied heavily on the underlying statistics of the correspondence between the rendered and sensor images. Therefore to apply these integrity approaches to new applications a similar rigorous analysis would be required to determine those statistical distributions, possibly

unique to that application and feature selection. A possible impact to broadening the operational envelope is dealing with non-stationary statistical distributions for the image correspondence. Hence a method for dealing with possible multiple distributions, depending on the operational state and variations in the model confidence are all areas for future consideration.

5.3 Summary

The development of trust in any system is vital to its utilization and operational use, particularly for safety critical applications, such as navigation. This dissertation described the rationale for developing vision-based precision navigation systems and the motivation for addressing the lack of quantitative integrity in this area. This research is a step forward in addressing this integrity gap and bringing the realization of vision navigation to the maturity and robustness necessary for making it a viable navigation solution.

Bibliography

[1] R. Andreas, L. Hostetler and R. Beckmann, "Continuous Kalman Updating of an Inertial Navigation System using Terrain Measurements," in *Proceedings of the IEEE National Aerospace Electronics Conference*, 1978.

[2] C. Baird and M. Abramson, "A Comparison of Several Digital Map-Aided Navigation Techniques," in *Proceedings of IEEE Position, Location and Navigation Symposium*, 1984.

[3] J. Rodriquez and J. Aggarwal, "Matching Aerial Images to 3-D Terrain Maps," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. 12, no. 12, 1990.

[4] T. Hoshizaki, D. Andrisani, A. Braun, A. Mulyana and J. Bethel, "Performance of Integrated Electro-Optical Navigation Systems," in *Proceedings of the Institute of Navigation GNSS*, Portland, OR, 2003.

[5] M. Veth and J. Raquet, "Fusing Low-Cost Image and Inertial Sensors for Passive Navigation," *Journal of the Institute of Navigation,* vol. 54, no. 1, pp. 11-20, 2007.

[6] J. Kim and S. Sukkarieh, "Real-time implementation of airborne inertial-SLAM," *Robotics and Autonomous Systems,* vol. 55, pp. 62-71, 2007.

[7] G. Dissanayake, S. Huang, Z. Wang and R. Ranasinghe, "A Review of Recent Developments in Simultaneous Localization and Mapping," in *6th International Conference on Industrial and Information Systems*, Sri Lanka, 2011.

[8] Z. Zhang, Y. Huang and C. Li, "Monocular Vision Simultaneous Localization and Mapping using SURF," in *7th World Congress on Intelligent Control and Automation*, Chongquing, China, 2008.

[9] S. Zhang, L. Xie and M. Adams, "An Efficient Data Association Approach to Simultaneous Localization and Map Building," in *IEEE International Conference on Robotics and Automation*, New Orleans, LA, 2004.

[10] A. D. Weaver, C. Larson, J. Raquet and M. Veth, "Developing a Framework for Image-based Integrity," in *Proceedings of the Institute of Navigation GNSS, 22nd International Technical Meeting of the Satellite Division*, Savannah, GA, 2009.

[11] J. Mitchell, S. Cryan, D. Strack, L. Brewster, M. Williamson, R. Howard and A. Johnston, "Automated Rendezvous and Docking Sensor Testing at the Flight Robotics Laboratory," in *Proceedings of IEEE Aerospace Conference*, 2006.

[12] J. Goddard, W. Jatko, R. Ferrell and S. Gleason, "Robust Pose Determination for Autonomous Docking," in *ANS Sixth Topical Meeting on Robotics and Remote Systems*, 1995.

[13] A. Tsai, P. Gibbens and R. Stone, "Terminal Phase Vision-Based Target Recognition and 3D Pose Estimation for a Tail-Sitter, Vertical Takeoff and Landing Unmanned Air Vehicle," in *Proceedings of the First Pacific Rim conference on Advances in Image and Video Technology*, 2006.

[14] C. Larson, J. Raquet and M. Veth, "Developing a Framework for Image-based Integrity," in *Proceedings of the Institute of Navigation GNSS, 22nd International Technical Meeting of the Satellite Division*, Savanah, GA, 2009.

[15] S. Calhoun, J. Raquet and J. Curro, "Flight Test Evaluation of Image Rendering Navigation for Close-Formation Flight," in *Proceedings Institute of Navigation GNSS*, 2012.

[16] M. Petovello, S. Pullen, J. Syrjarinne and L. Wirola, "Quantifying the performance of navigation systems and standards for assisted-GNSS," *GNSS Solutions,* September/October 2008.

[17] *DO-208: Minimum Operational Performance Standards (MOPS) for Airborne Supplemental Navigation Equipment Using Global Positioning System (GPS),* Washington, D.C.: RTCA-159, 1991.

[18] *DO-245A: Minimum Aviation System Performance Standards for the Local Area Augmentation System (LAAS),* Washington, D.C.: RTCA SC-159, WG-4, 2004.

[19] J. Rife, S. Pullen, P. Enge and B. Pervan, "Paired Overbounding for Nonideal LAAS and WAAS Error Distributions," *IEEE Transactions on Aerospace and Electronic Systems,* vol. 42, no. 4, pp. 1386-1395, 2006.

[20] J. W. Harris, Handbook of Mathematics and Computational Science, New York: Springer-Verlag, 1998.

[21] [Online]. Available: http://en.wikipedia.org/wiki/File:Standard_deviation_diagram.svg.

[22] P. Maybeck, Stochastic Models, Estimation, and Control, Vol. 2, R. Bellman, Ed. Academic Press, 1982.

[23] D. Lowe, "Object Recognition from Local Scale-Invariant Features," in *Proceedings of the International Conference on Computer Vision*, 1999.

[24] H. Bay, T. Tuytelaars and L. Van Gool, "SURF: Speeded Up Robust Features," in *Proceedings of the 9th European Conference on Computer Vision*, 2006.

[25] A. Jain, Fundamentals of Digital Image Processing, New Jersey: Prentice Hall, 1989.

[26] L. Davis, "A survey of Edge Detection Techniques," *Computer Graphics and Image Processing,* vol. 4, 1975.

[27] I. Misra, M. Moorthi, D. Dhar and R. Ramakrishnan, "An Automatic Satellite Image Registration Technique Based on Harris Corner Detection and Random Sample Consensus (RANSAC) Outlier Rejection Model," in *1st International Conference on Recent Advances in Information Technology*, 2012.

[28] C. Harris and M. Stephens, "A Combined Corner and Edge Detector," in *Alvey Vision Conference*, 1988.

[29] K. Ohba and K. I., "Detectability, Uniqueness, and Reliability of Eigen Windows for Stable Verification of Partially Occluded Objects," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. 19, no. 9, 1997.

[30] L. B. Mehmedovic, I. Petrovic and E. Ivanjko, "Hough Transform based Correction of Mobile Robots Orientation," in *IEEE International Conference on Industrial Technology*, 2004.

[31] A. Sharf and A. Shamir, "Feature-sensitive 3D Shape Matching," in *Proceedings of the Computer Graphics International*, 2004.

[32] P. Hwang and R. Brown, "From RAIM to NIOAIM: A New Integrity Approach to Integrated Multi-GNSS Systems," *InsideGNSS,* pp. 24-33, May-June 2008.

[33] I. Gordon and D. Lowe, "Scene Modelling, Recognition, and Tracking with Invariant Image Features," in *International Symposium on Mixed and Augmented Reality*, 2004.

[34] J. Friedman, J. Bentley and R. Finkel, "An Algorithm for Finding Best Matches in Logarithmic Expected Time," vol. 3, pp. 209-226, 1977.

[35] J. Beis and D. Lowe, "Shape Indexing Using Approximate Nearest-Neighbour Search in High-Dimensional Spaces," *Proceedings of the IEEE,* 1997.

[36] M. Fischler and R. Bolles, "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography," *Communications of the ACM,* vol. 24, no. 6, 1981.

[37] L. Cheng, H. Hu, Y. Wang and M. Li, "A New Method for Remote Sensing Image Matching by Integrating Affine Invariant Feature Extraction and RANSAC," in *3rd International Congress on Image and Signal Processing*, 2010.

[38] T. Kim and Y. Im, "Automatic Satellite Image Registration by Combination of Matching and Random Sample Consensus," *IEEE Transactions on Geoscience and Remote Sensing,* vol. 41, no. 5, 2003.

[39] S. Yang, C. Wang and C. Chang, "RANSAC Matching: Simultaneous Registration and Segmentation," in *IEEE International Conference on Robotics and Automation*, Anchorage, AK, 2010.

[40] W. Pratt, Digital Image Processing, 2nd Ed, New York: John Wiley & Sons, 1991.

[41] N. Iyer, S. Jayanti, K. Lou, Y. Kalyanaraman and K. Ramani, "Three-Dimensional Shape Searching: State-of-the-Art Review and Trends," *Computer Aided Design,* vol. 37, pp. 509-530, 2005.

[42] Z. Li, K. Hou and H. Li, "Similarity Measurement Based on Trigonometric Function Distance," in *1st International Symposium on Pervasive Computing and Applications*, 2006.

[43] J. Krumm, "Eigenfeatures for Planar Pose Measurement of Partially Occluded Objects," in *Proceeding of the IEEE Conference on Computer Vision and Pattern Recognition*, 1996.

[44] S. Ando, Y. Kusachi and K. Arakawa, "Appearance Based Pose Estimation of 3D Object Using Support Vector Regression," in *Proceeding of the IEEE Conference on Image Processing*, 2005.

[45] L. Iocchi, D. Mastrantuono and D. Nardi, "A Probabilistic Approach to Hough Localization,"

in *Proceedings of the IEEE Conference on Robotics and Automation*, Seoul, Korea, 2001.

[46] C. Yang, S. Lai and L. Chang, "Reliable Image Matching via Modified Hausdorff Distance with Normalized Gradient Consistency Measure," in *IEEE 3rd International Conference on Information Technology*, 2005.

[47] F. Chang, Z. Chen, W. Wencheng and L. Wang, "The Hausdorff Distance Template Matching Algorithm Based on Kalman Filter for Target Tracking," in *Proceedings of the IEEE International Conference on Automation and Logistics*, 2009.

[48] G. Xu and W. Yang, "A Corner Feature Matching Model Based on Improved Hausdorff Distance," in *2nd International Conference on Intelligent Computation Technology and Automation*, 2009.

[49] G. Xu and W. Yang, "A Composite Model Based on Shape for Fast Image Matching," in *International Asia Conference on Informatics in Control, Automation and Robotics*, 2009.

[50] C. Olson, "A Probabilistic Formulation for Hausdorff Matching," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1998.

[51] W. Wells, "Statistical Approaches to Feature-Based Object Recognition," *International Journal of Computer Vision,* vol. 21, pp. 64-98, 1997.

[52] C. Olson, "Probabilistic Self-Localization for Mobile Robots," *IEEE Transactions Robotics and Automation,* vol. 16, no. 1, pp. 55-66, 2000.

[53] C. Olson, "Maximum Likelihood Image Matching," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. 24, no. 6, 2002.

[54] C. Olson, "Subpixel Localization and Uncertainty Estimation Using Occupancy Grids," in *Proceedings of the IEEE International Conference on Robotics and Automation*, Detroit, MI, 2009.

[55] B. Silverman, Density Estimation for Statistics and Data Analysis, New York: Chapman and Hall, 1986.

[56] H. Poor, An Introduction to Signal Detection and Estimation, 2nd Edition, Dowden & Culver, 1994.

[57] A. Whalen, Detection of Signals in Noise, Academic Press, Inc., 1971.

[58] J. Babuad, A. Witkin, M. Baudin and R. Duda, "Uniqueness of the Gaussian Kernel for Scale-Space Filtering," *IEEE Transactions on Pattern Analysis and Machine Learning,* vol. 8, no. 1, 1986.

[59] R. Brandt, Y. Wang, A. Laub and S. Mitra, "The Recognition of Shape in Binary Images using a Gradient Classifier," *IEEE Transactions on Systems, Man, and Cybernetics,* vol. 19, no. 6, 1989.

[60] R. Gonzalez and R. Wood, Digital Image Processing, 2nd Ed., Prentice Hall, 2002.

[61] Z.Zhang, "Flexible Camera Calibration by Viewing a Plane from Unknown Orientations," in *IEEE Proceedings Computer Vision and Pattern Recognition*, 1999.

[62] O. S. J. Heikka, "A Four-Step Camera Calibration Procedure with Implicit Image Correction," in *IEEE Proceedings Computer Vision and Pattern Recognition*, 1997.

[63] S. Verhagen, *The GNSS integer ambiguities estimation and validation,* 2005.

[64] W. K. Pratt, Digital Image Processing, 2nd ed., New York: John Wiley & Sons, 1991.

[65] *System Requirements Document (SRD) for Joint Precision Approach and Landing System (JPALS) Increment 1,* United States Naval Air Systems Command PMA213, 2010.

[66] *System Requirements Document (SRD) for Joint Precision Approach and Landing System (JPALS) Increment 2,* United States Naval Air Systems Command PMA213, 2010.

[67] G. Box and G. Tiao, Bayesian Inference in Statistical Analysis, Wiley-Interscience, 1992.

[68] G. Cassella and R. Berger, Statistical Inference, 2nd Ed., Duxbury Thomson Learning, 2002.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
24-03-2016	Doctoral Dissertation	Sept 2009 - Mar 2016

4. TITLE AND SUBTITLE
Integrity Determination for Image Rendering Vision Navigation

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Calhoun, Sean M., Contractor

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Air Force Institute of Technology,
Graduate School of Engineering and Management (AFIT/EN)
2950 Hobson Way
WPAFB OH 45433-7765

8. PERFORMING ORGANIZATION REPORT NUMBER
AFIT-ENG-DS-16-M-251

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Intentionally Left Blank

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A. APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

13. SUPPLEMENTARY NOTES
The U.S. Government retains access rights to this work, and has the right to disseminate this work for government purposes

14. ABSTRACT
This research addresses the lack of rigorous and quantitative integrity approaches for vision navigation systems. Two unique approaches are developed under this research for determining the uncertainty and integrity for a vision based, precision relative navigation system. The first technique is based on developing a discrete detection scheme for determining when large excursions from the desired operation envelope occur. The second approach utilizes a generalized Bayesian inference approach, in which a full probability density function determination of the state estimate is realized.

15. SUBJECT TERMS
navigation, integrity, image processing, vision navigation, image-aided navigation

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dr. John F. Raquet, AFIT/ENG
					19b. TELEPHONE NUMBER *(Include area code)*
U	U	U	UU	152	(937) 255-3636 x4580