

UNCLASSIFIED

The ABC's of RSM: An introduction to the Replacement Sensor Model

September 30, 2006

Michelle Iiyama
John Dolloff
Charles Taylor

BAE SYSTEMS

NS-C3I-082708-0034- These commodities, technology or software were exported from the United States in accordance with the Export Administration Regulations. Diversion contrary to U.S. Law prohibited.

©2006 BAE Systems

National Security Solutions P.O. Box 509009 San Diego, CA 92150-9009 USA
Telephone (858) 675-2600

UNCLASSIFIED

Table of Contents

1.0	Introduction.....	4
2.0	Prelude	4
3.0	What's a Physical Sensor Model?.....	6
3.1	Parts Is Parts.....	10
3.2	What Can You Do With It?	15
3.3	How Does It All Work?	21
3.4	If It's Not Broke, Why Fix It?	21
4.0	What Is RSM?.....	22
4.1	Parts... We Have Them All	24
4.2	One Size Fits All.....	31
4.3	Advantage...RSM.....	32
5.0	Evolution.....	33
5.1	The Journey.....	33
5.2	Seeing Is Believing	34
6.0	Cooking With RSM	37
6.1	Cooking From Scratch	37
6.2	Just Add Water.....	38
7.0	Tributes	42
8.0	Coda	42
9.0	References.....	43

Table of Tables

Table 1: RSM TREs.....	31
Table 2: Original Sensor Model vs. RSM.....	36

Table of Figures

Figure 1: Sample Optical Image	5
Figure 2: Sample Sensors and Their Images.....	6
Figure 3: Image Point / Ground Point Relationship.....	7
Figure 4: Image-To-Ground Height Dependence	9
Figure 5: Effects of Sensor Position Errors	12
Figure 6: Effects of Attitude Errors	13
Figure 7: Effects of Focal Length Errors	14
Figure 8: Mono Extraction and Accuracy Prediction	17
Figure 9: Multi-Image Extraction and Accuracy Prediction.....	18
Figure 10: Sensor Parameter Adjustments.....	20
Figure 11: RSM Flow	23
Figure 12: RSM Ground-To-Image	25
Figure 13: RSM Ground-To-Image Function	26
Figure 14: Application of RSM Adjustable Ground-to-Image Function.....	28
Figure 15: RSM Sample Geo-positioning Tests	35
Figure 16: Integration of RSM Generator.....	39
Figure 17: Integration of RSM Exploiter.....	41

1.0 Introduction

This document provides a light introduction to the Replacement Sensor Model referred to by the acronym, RSM. No prior knowledge of sensor models, or advanced math is required; and in fact, we will avoid the use of equations. The following will describe RSM, its benefits, its development history, and in the end, discuss how to obtain it. For those intrigued to know more, they are invited to follow the references to more detailed documentation.

Note 1.0-1: Throughout this document there will be “little blue boxes” containing notes that offer additional insight and depth. They are not deemed necessary to the basic description.

Advisory: Parts of this document are intended to be somewhat whimsical in nature in order to prevent drowsiness.

2.0 Prelude

Before we can begin talking about a replacement sensor model, we need to understand not only what a sensor model is, but what a sensor is. A sensor, according to the Merriam-Webster dictionary is:

A device that responds to a physical stimulus ([such] as heat, light, sound, pressure, magnetism, or a particular motion) and transmits a resulting impulse (as for measurement or operating a control)

In our case, a sensor is quite simply a fancy camera mounted upon a moving vehicle (platform), and takes pretty pictures (images). Imagine being in an airplane, taking a picture of the ground below (Figure 1).

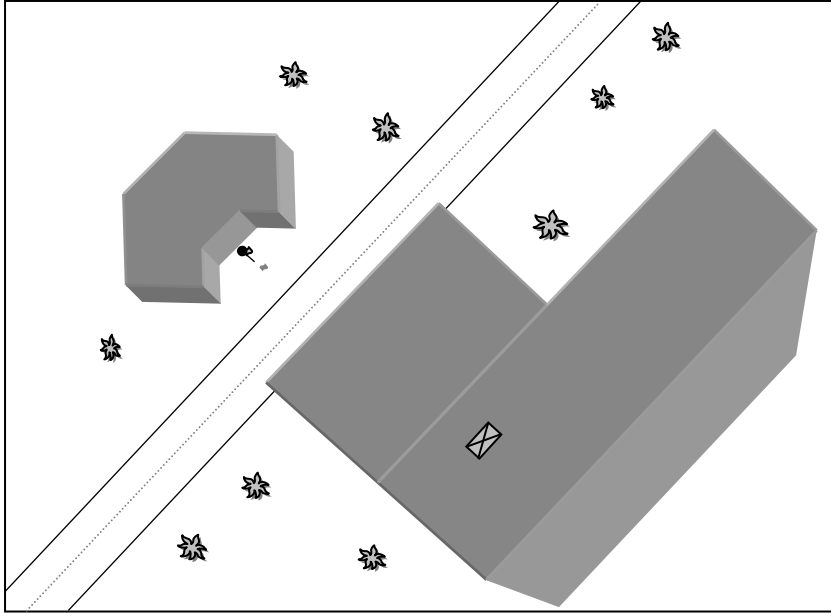


Figure 1: Sample Optical Image

So, in our context a sensor takes pictures. Looking at the picture above, we can see identifiable objects: buildings, a street, a flag pole. But it's just a picture. Suppose we wanted more information from that picture, like the ground coordinates (i.e. ground position or geo-location) of the vent (marked with an 'X') on the rooftop of the larger building. What processes are involved and what do they rely on?

In general, the process of deriving and using information from an image is termed "image exploitation". Furthermore, when the information includes the geo-location of features or objects in an image, the processes involved are more specifically termed "geo-positioning" processes. Targeting with reliable accuracy predictions is an important example of a geo-positioning process. Other examples include: Digital Elevation Model (DEM) generation, site model generation, and change detection. And what do these processes have in common? They as well as all other geo-positioning processes rely upon the physical sensor model as their basic building block.

Before proceeding further, we must first back up a little and note that there are different types of sensors with different resulting images (Figure 2). Not only are there the optical sensors that produce images we recognize as pictures, but there are sensors that produce images with geo-positioning information that are difficult to visually interpret, like Synthetic Aperture Radar (SAR). And still, there are other sensors that measure quantities such as radiometric data. For this discourse, we are concerned with geo-positioning aspects; thus, for introductory ease, we will (in general) limit our discussion to optical sensors producing easily recognizable images. However, we note that although we are discussing optical sensors and images primarily, the Replacement Sensor Model also works for SAR and other types of images, which is discussed later.




Figure 2: Sample Sensors and Their Images

Now, before we go on to discuss the physical sensor model, let's first present a few preliminary definitions that will be used throughout the remainder of this exploration. The term "image point" will be used to represent a specific pixel location on the image. It is specified by two coordinates: (line, sample) or alternatively (row, column). Furthermore, the term "ground point" will be used to represent a specific location on or near the earth's surface. It is specified by three coordinates: (x, y, z) or alternatively (latitude, longitude, height or elevation). Now, on with the adventure...

3.0 What's a Physical Sensor Model?

A sensor's physical sensor model is a mathematically rigorous description of the relationship between the pixel coordinates of the image and the physical ground coordinates they represent. Figure 3 presents a conceptual representation of this relationship.

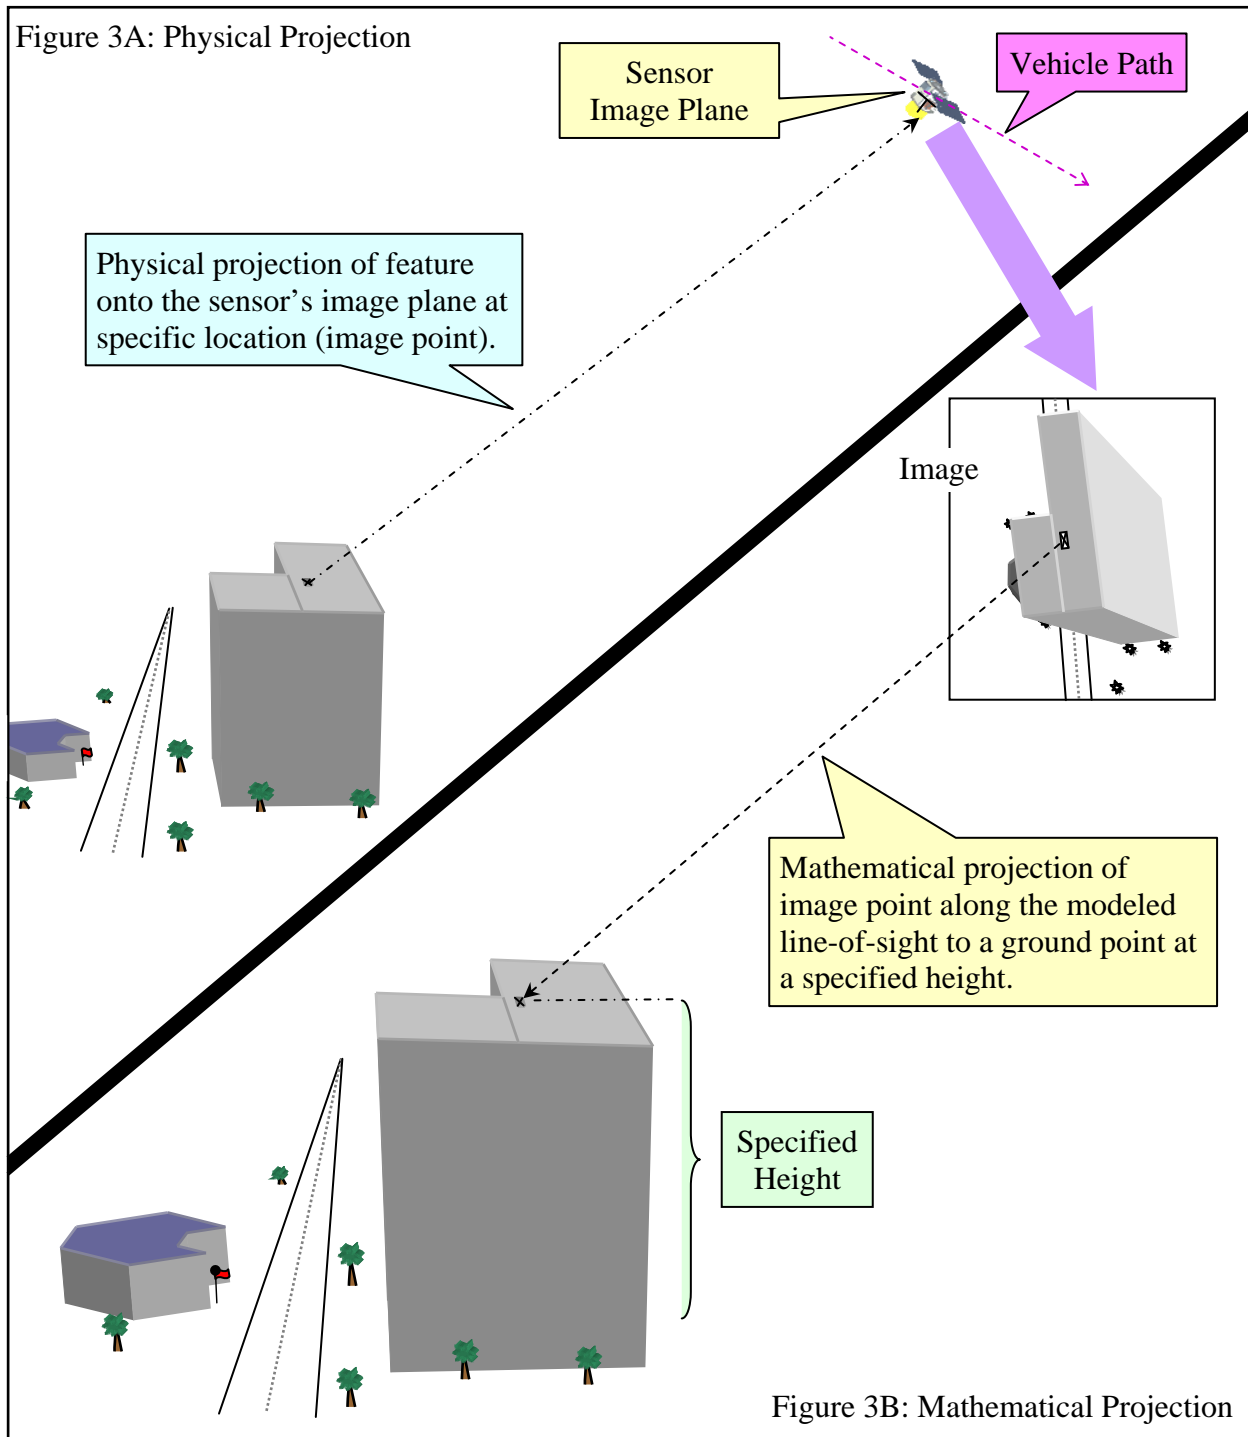


Figure 3: Image Point / Ground Point Relationship

For an optical sensor, the incident radiation from the sun (a.k.a. sunlight) reflects off of objects and features in the scene onto the sensor’s image plane. Let’s term this process a “physical projection”—see Figure 3A, the top portion of Figure 3. An image is then formed, either in (near) real-time or later at a ground processing station. The physical sensor model is then used to relate an image point to corresponding ground points along the line-of-sight. Let’s term this process a “mathematical projection”—see Figure 3B, the bottom portion of Figure 3. A unique

UNCLASSIFIED

ground point is then identified by the specification of an external (user supplied) height, which is typically obtained from a Digital Elevation Model (DEM).

Each sensor has a sensor model that is a physically-based mathematical algorithm that models the line-of-sight (LOS). That means, the algorithm captures and interprets the mechanics of the sensor's physical construction, as well as the physics and relationships of additional information about how the image was taken, and uses this information in a mathematical algorithm to determine (model) the LOS. It is important to note that the physical sensor model can vary greatly from one sensor to the next.

The additional information needed for the physical sensor model is collected and reported as part of the image support data for each image, and includes physical parameters such as: where the sensor was when the image was taken (i.e. the geo-location of the sensor, or sensor position, or ephemeris), the direction the sensor was pointing (i.e. the attitude), focal length (for optical sensors, the distance from the imaging plane to the optical lens), and vehicle velocity (for SAR sensors).

The image support data is usually contained in a file that accompanies the image; whereas the image itself can be thought of as simply a file of "gray-shade" or "gray-count" values per pixel location.

The physical sensor model's image-to-ground function projects an image point along the LOS to a specified height, resulting in the reported ground point. Thus, if we change the specified height, we will get a different ground point. Notice, that in Figure 4, using height 1 would return the "correct" geo-location (ground point) of the vent. However, as an example, specifying an incorrect height 2, the same projection returns the geo-location not of the intended vent, but of the "Embassy" across the street.




Figure 4: Image-To-Ground Height Dependence

In the example above, an incorrect ground point was obtained when the wrong height was specified. However, errors in the reported ground point will always happen operationally because of unavoidable errors in a DEM, unavoidable image measurement errors, and unavoidable errors in the image support data (examples to be given in the following sections). Given all of this, it is undeniable that for most applications, a ground point is not enough information. In addition to the geo-location, users (especially those responsible for launching weapons) need a reliable prediction of how accurate that ground point is. Thus, a complete physical sensor model will provide not only the capability of identifying an image/ground relationship, but also be able to tell us how accurate that image/ground relationship is.

Note 3.0-1: For optical sensors, the modeled line-of-sight (LOS) is also called the image ray. For all types of imaging sensors, the image/ground correspondence model equivalent to the optical LOS is called the imaging locus. For example, the imaging locus for a SAR sensor corresponds to the range-doppler circle.

Note 3.0-2: For appropriate vehicles (e.g., space-borne) the modeled LOS includes corrections for physical phenomena such as atmospheric refraction (i.e., the physical bending of the light as it travels through the atmosphere).

3.1 Parts Is Parts

Every sensor model has, at the least, an image/ground relationship needed for extraction of a target's ground point. However, complete sensor models have three components or "parts": 1) an image/ground relationship, 2) adjustable parameters that affect the image/ground relationship, and 3) an error covariance that corresponds to errors in the adjustable parameters.

Note 3.1-1: It is the adjustable parameters and their error covariances that allow us to generate accuracy predictions necessary for all reliable targeting applications.

3.1.1 Image/Ground Function

The image/ground relationship is needed for extraction of a target's ground coordinates. In order to support all geo-positioning applications, both an image-to-ground function and a ground-to-image function are needed.

The image-to-ground function has already been introduced—see Figure 3 and 4. It is a projection of an image point (pixel location) to the corresponding ground point at a specified height. The ground-to-image function goes the other direction. That is, it's a projection of a ground point to the corresponding image point. Both functions require 3 independent coordinates as input and provide 2 independent coordinates as output. More specifically, the image-to-ground function requires a 2-dimensional image point, e.g. (line, sample), and a specified height as input, and returns the 3-dimensional ground point, e.g. (latitude, longitude, height), whose 3rd dimension is really the height given as an input; while the ground-to-image function requires a 3-dimensional ground point as input, and returns the 2-dimensional image point.

Physical sensor models usually have a direct (explicit) image-to-ground function. Whereas in most cases the inverse function (ground-to-image) is not direct, but can be computed as an

iterative process that repeatedly calls the image-to-ground function (typically 3 times). However, for some sensor models, the direct relationship is the ground-to-image function, in which case the image-to-ground function can be computed by an iterative process that repeatedly calls the ground-to-image function. Thus, for every sensor model, there is always both an image-to-ground and a ground-to-image function, one of which is usually an iterative inverse of the other.

3.1.2 Adjustable Parameters

Now, complete sensor models also have parameters that not only affect the image/ground relationship but are adjustable as well. In a typical optical sensor, these adjustable parameters include the sensor position when the image was taken, its attitude (what direction it was pointing), and its focal length.

Adjustments to these parameters consist of corrections to the initial (*a priori*) parameter values contained in the image support data. If the image support data contains adjustments/corrections to the *a priori* parameters, then we know that the image support data has gone through an “external” adjustment process such as “triangulation” or “image registration”. The use of the adjusted image support data in subsequent target extractions makes the resulting geo-location estimates more accurate.

In addition, adjustable parameters also serve as a mechanism for the statistical description of image support data (adjusted or unadjusted) uncertainty. In particular, an error covariance is defined relative to errors in the adjustable parameter values. It is the use of this error covariance in conjunction with the appropriate partial derivatives that allows various exploitation processes to correctly account for image support data uncertainty (see Sections 3.1.3 and 3.2.1 for more information).

So, how do errors in these adjustable parameters affect the image/ground relationship?

The sensor position affects the image/ground relationship such that if the sensor position is different from where you thought it was, the basic geometry will be off, resulting in a translation error in the image/ground projection. Thus, as depicted in Figure 5, using an incorrect sensor position returns the geo-location not of the “warehouse” (marked with an ‘X’), but of the “hospital” across the street.




Figure 5: Effects of Sensor Position Errors

Note 3.1.2-1: The sensor position is a three-dimensional position, so keep in mind that the position error can be in any or all three directions.

The attitude is the direction the camera was pointing at the time the image was taken. Thus, if the attitude values reported are in error, the sensor was actually pointing in a different direction than where we thought it was, and instead of returning the desired geo-location of the “factory”, we inadvertently get the geo-location of the “day-care center”. These errors are the rotation errors/adjustments commonly referred to as (roll, pitch, and yaw).

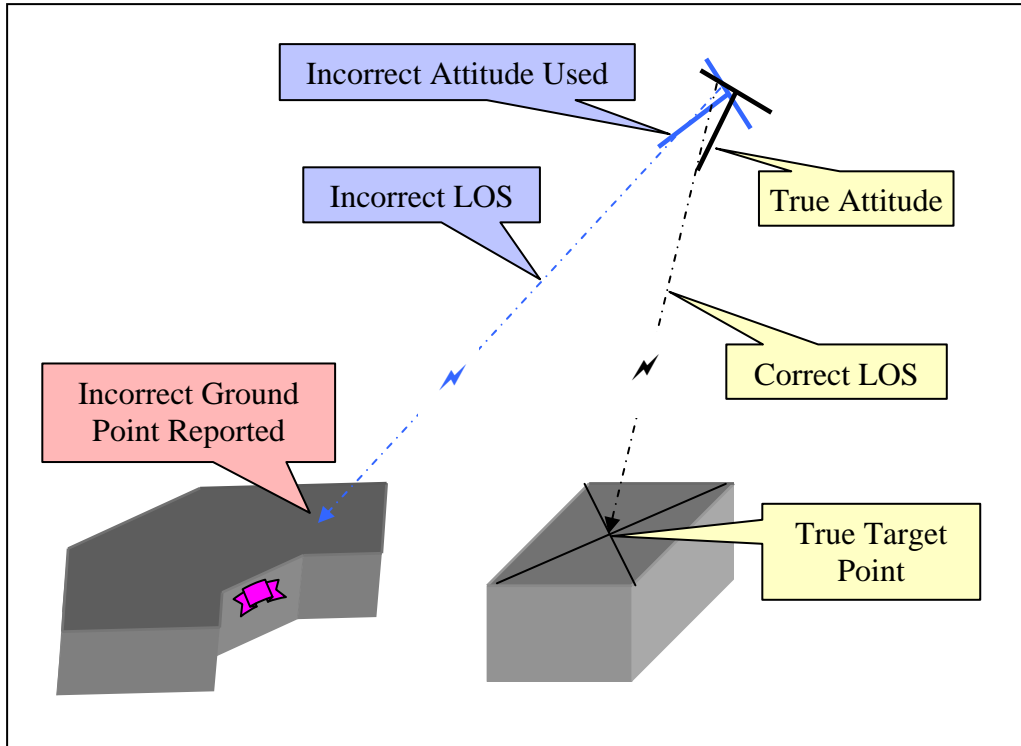


Figure 6: Effects of Attitude Errors

Note 3.1.2-2: The effect of the attitude errors increases as the distance from the sensor to the target increases.

Finally, if the focal length is in error (from heat expansion, contraction, or calibration error) we end up with an incorrect LOS, which is again ultimately realized as an incorrect ground point.




Figure 7: Effects of Focal Length Errors

Note 3.1.2-3: The effect of the focal length error increases as the distance from the sensor to the target increases.

Thus, given that we need to know values for the adjustable parameters (sensor position, attitude, and focal length) in order to determine the ground point, we have demonstrated that if any of those values are in error, the resulting ground point will be in error as well.

But, we know that nothing is ever perfect. Adjustable parameter values are always in error since they correspond to previous measurements, estimates, and predictions, all of which are imperfect by their very nature. In fact, for sensor model applications, these errors are far from insignificant and must be accounted for. Specifically, we need to know how much uncertainty is in each parameter's value. That is, for a complete sensor model we also need an error covariance that corresponds to the adjustable parameters.

3.1.3 Error Covariance

An error covariance is simply a statistical description or representation of both the error (uncertainty) in each parameter of interest, and the correlation of errors between these parameters.

Note 3.1.3-1: If there are n parameters of interest, the corresponding error covariance is represented by an $n \times n$ matrix.

A simple, non-sensor model, scalar (one-dimensional) example: If we measure a distance x to be 5 meters, and the measurement has a standard deviation of error of ± 10 cm then, assuming the underlying error is normally distributed (sorry, a statistical definition necessity), there is a 68% probability that the true distance is somewhere between 4.9 and 5.1 meters. In this simple example, the covariance matrix is a 1×1 matrix (scalar) with a value equal to the variance (i.e. the square of the standard deviation).

As discussed earlier, a complete sensor model includes an error covariance corresponding to the adjustable parameters; which captures the uncertainty in the image support data.

Note 3.1.3-2: In our optical frame sensor example, there are 7 adjustable parameters (3 that define sensor position, 3 that define attitude, and 1 for focal length). Thus, for one image the error covariance is a 7×7 matrix.

3.2 What Can You Do With It?

Very simply put, with a complete sensor model, we can do extractions, adjustments, and accuracy predictions.

Extract, adjust, and predict...what?

The extraction processes are methods for obtaining a geo-location, while the adjustment processes are methods used to refine the image support data used in the extraction process (to get a better ground point estimate). Thus, the extraction and adjustment processes both deal with methods of obtaining a geo-location. However as we pointed out earlier, for most users, knowing a “where” is not sufficient information to make a decision to take action on—users also need to know the accuracy of that ground point estimate. Thus, the following sections will discuss the extraction processes in conjunction with an appropriate accuracy prediction method.

3.2.1 Monoscopic Extraction and Accuracy Prediction

An extraction of ground point information using only one image is termed a mono (monoscopic) extraction. The mono extraction solution corresponds to the intersection of the modeled LOS with a height estimate, which is just the application of the sensor model's image-to-ground function (see Figure 3 above).

More specifically, the mono extraction (image-to-ground function) inputs the image point specified as image coordinates (line, sample), or equivalently (row, column), where the target/object is seen in the image. It then projects the image point along the modeled LOS and intersects it with a specified height provided by an external source, such as a Digital Elevation Model (DEM). This projection results in a ground point, usually specified in geodetic ground coordinates (latitude, longitude, and height or elevation), but on occasion in rectangular ground coordinates, (x, y, z) , depending upon the user's implementation.

The (mono extraction) solution's error covariance is computed by the extraction process using the image support data (adjustable parameter) error covariance projected along the modeled LOS, and the height error covariance obtained from the external source. The resultant solution error covariance, a 3×3 matrix, can be represented graphically as an error ellipsoid; which in turn, can be approximated by two numbers, the 90% horizontal circular error (CE), and the 90% vertical error (LE), (see Figure 8). These two numbers, CE and LE form the basis for accuracy prediction.

Note 3.2.1-1: The extraction process also accounts for the uncertainty in the actual identification and measurement of the object's location in the image. However, the effect of this error is typically much less than the effects of the image support data errors.

Example: Suppose CE is 10 meters and LE is 30 meters. Then, we can say that it is 90% probable that the target's true location is horizontally within 10 meters of the computed solution; and similarly, it is 90% probable that the target's true location is vertically within 30 meters of the computed solution.

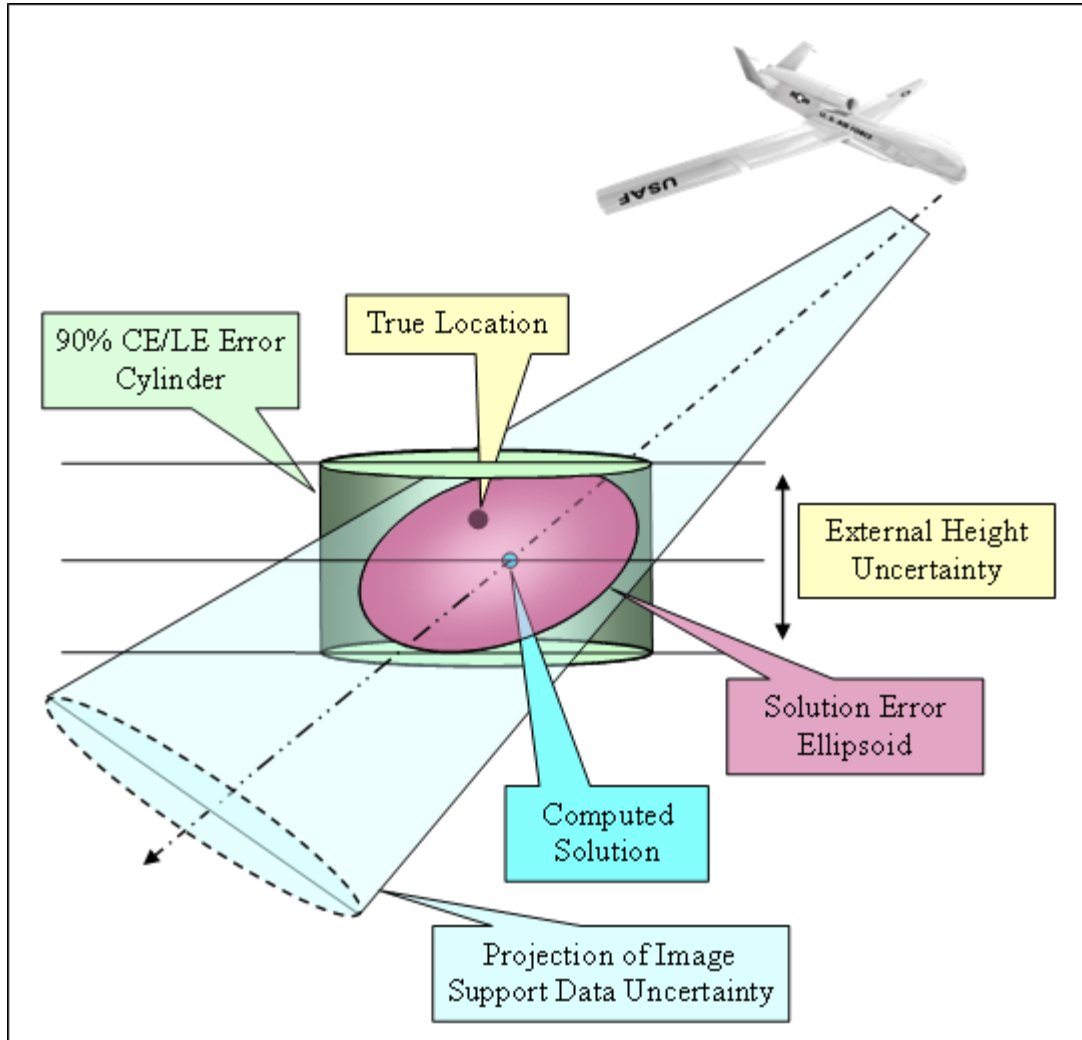


Figure 8: Mono Extraction and Accuracy Prediction

Note 3.2.1-2: The projection of the adjustable parameter error covariance is achieved via partial derivatives. In order to support extraction and accuracy prediction, a complete physical sensor model must support the computation of both partial derivatives of image coordinates (line and sample) with respect to the sensor's adjustable parameters, and the partial derivatives of image coordinates with respect to the target's ground coordinates.

Note 3.2.1-3: The extraction process is also applicable to the mensuration of objects—e.g. we can extract two ends of a runway for a measure of its length. In this case, the applicable accuracy predictions are then relative circular error (RCE) and relative linear error (RLE).

3.2.2 Multi-Image Extraction and Accuracy Prediction

When an object is viewed in multiple images (two or more), the multi-image extraction process is able to not only produce a more accurate solution, but is able to do it without external height information.

In this process, the modeled LOS, or image ray, from each image is generated using the image-to-ground function, and each ray's contribution to the solution is then weighted according to its corresponding image support data uncertainty (error covariance); such that, the smaller the uncertainty, the larger the weight is.

In this case, both the solution ground coordinates and its accuracy prediction are dependent on the error covariance (image support data uncertainty).

The following Figure 9 graphically represents a three image solution and its resultant CE/LE error cylinder or accuracy prediction (underlying solution error ellipsoid not shown).

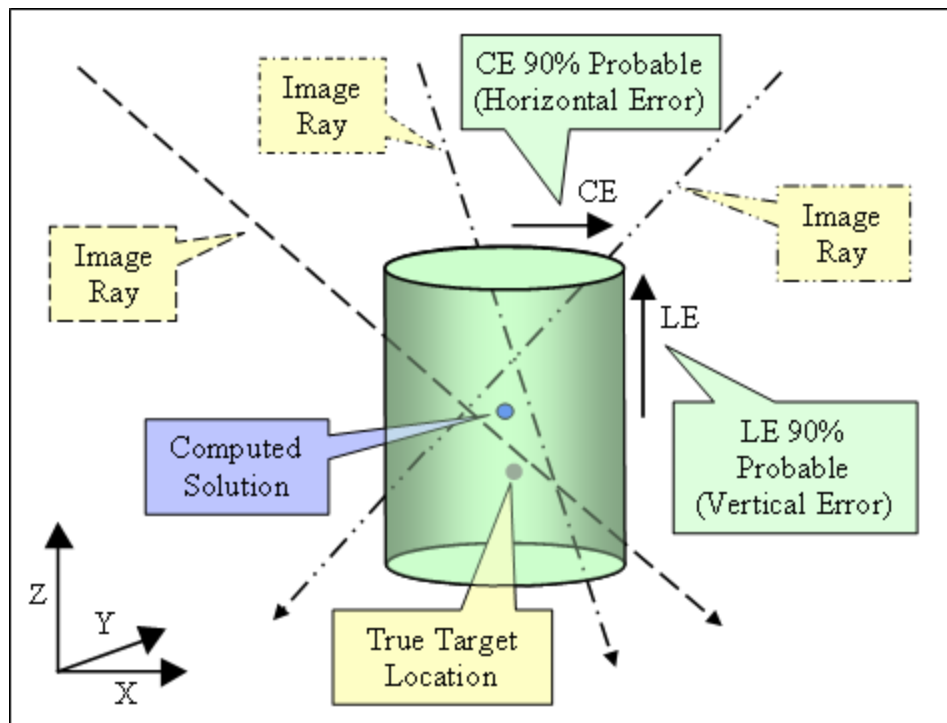


Figure 9: Multi-Image Extraction and Accuracy Prediction

Note: 3.2.2-1: An optimal multi-image extraction solution (algorithm) is presented in the 2004 Manual of Photogrammetry (see reference [1]).

3.2.3 Support Data Adjustment

In an adjustment process such as triangulation, we simultaneously solve for corrections to sensor model adjustable parameters in one or more images—which requires the use of all three components of a complete physical sensor model(s).

The image-to-ground function and image support data (sensor model adjustable parameter) error covariance are used in a manner similar to their use in multi-image extractions. However, where the multi-image extraction process only solves for ground point locations, the adjustment process also solves for corrections to the adjustable parameters for each image and their corresponding (improved) error covariance. These corrections and improved error covariance are then saved in the image support data so that subsequent target (mono and multi-image) extractions can utilize the corrections for more accurate geo-locations, and the improved error covariances for reliable accuracy predictions.

Note 3.2.3-1: If n images are involved in the adjustment process, with m adjustable parameters per image, the *a posteriori* (post solution) error covariance is an $nm \times nm$ matrix and reflects less uncertainty (is smaller) for each individual image, but also contains non-negligible correlations between images.

Figure 10 presents an overview of the adjustment process for the simplest scenario possible: one image (not shown explicitly in figure) using ground control points. Ground control points are photo-identifiable points with pre-determined ground locations and accuracy. The difference between the estimated ground control point locations using the pre-adjusted image support data with their known location (within the supplied accuracy) drives the adjustment process. The resulting accuracy of a post-adjustment extraction is far better than for a pre-adjustment extraction. In the following Figure 10, the resulting accuracy is idealized, as small residual errors always remain.

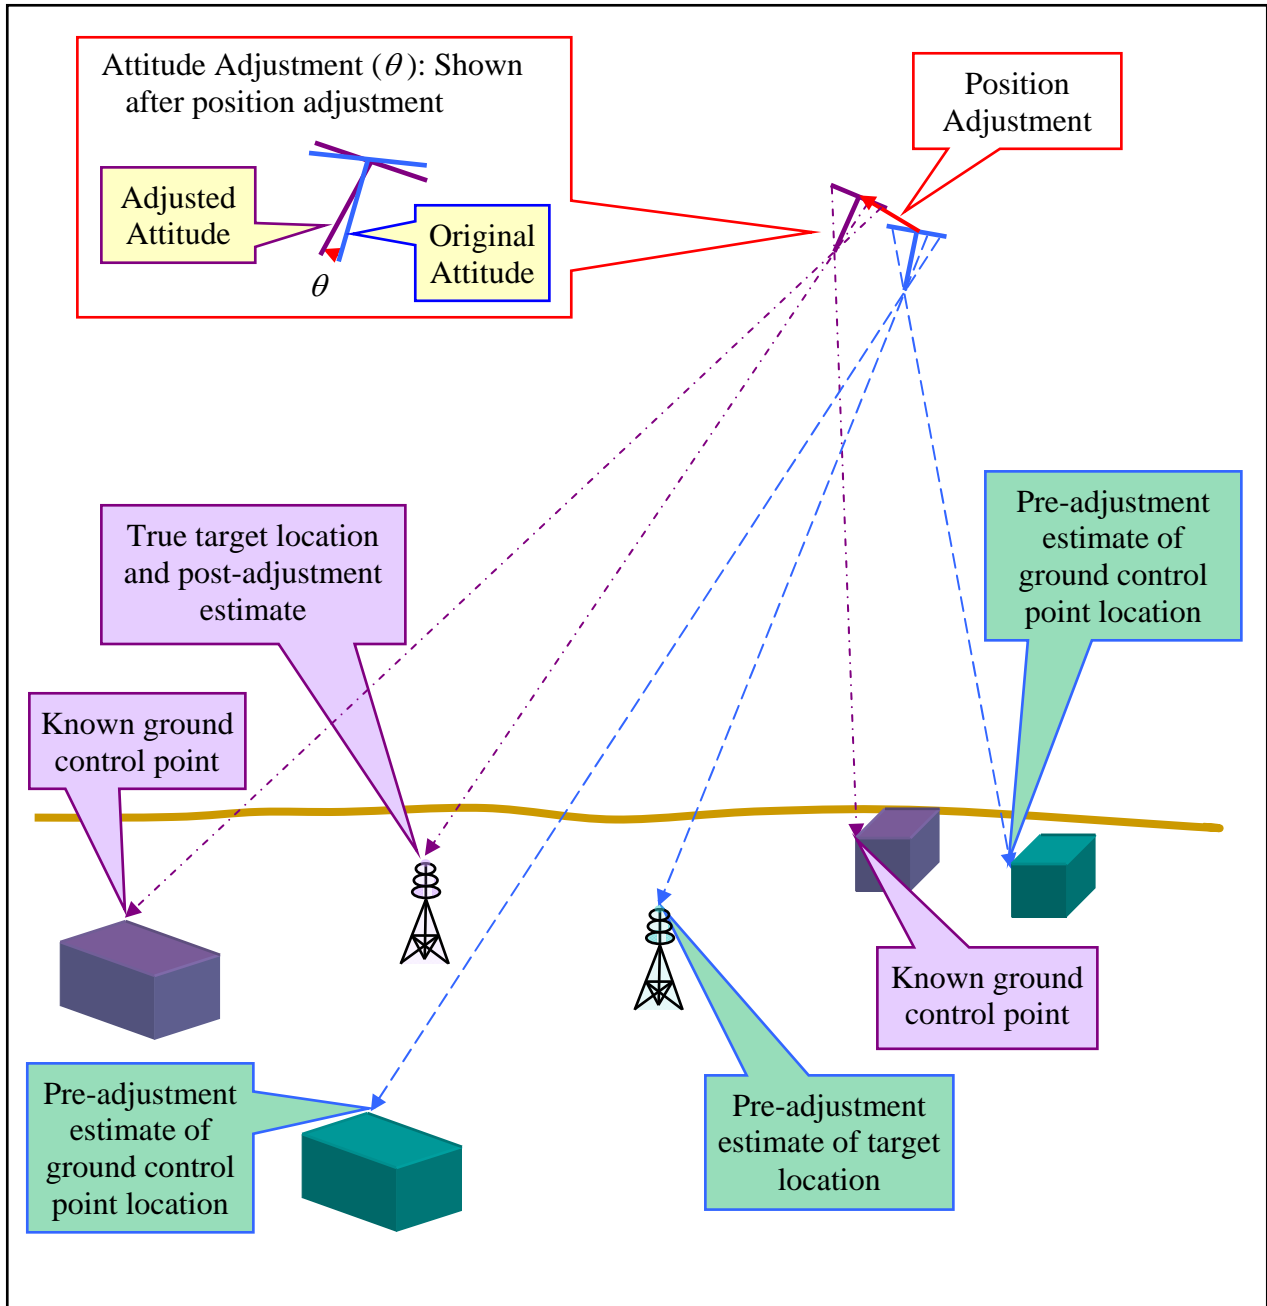


Figure 10: Sensor Parameter Adjustments

Note 3.2.3-2: In addition to the optimal geo-positioning (multi-image extraction) algorithm, the 2004 Manual of Photogrammetry discusses an optimal adjustment algorithm as well.

This concludes our “guided tour” of the three major components of a complete physical sensor model, and their critical roles in geo-positioning: extraction, adjustment, and accuracy prediction.

Note 3.2.3-3: We need to emphasize that a complete physical sensor model doesn't actually perform these geo-positioning processes (extraction, adjustment, and accuracy prediction), higher level applications perform them. However, the applications require the complete physical sensor model in order to do so.

We now discuss the practical issues associated with development, maintenance, and operational readiness of sensor models, and following that, we discuss RSM itself.

3.3 How Does It All Work?

In today's world, a sensor model user (exploitation system) must work with imagery from many different sensors. There are different sensors for different capabilities, most with their own unique image support data format. Of the potentially hundreds of different sensors/image support data formats that exist now and in the near future, a given exploitation system may be interested in tens of them.

When a sensor is designed and put into operation, its sensor model and support data are defined (published), and made available to the user community. Each potential user must then implement and integrate the new sensor model and its typically unique image support data format in their system.

However, even once a sensor is in “full operation”, it is not uncommon for the physical sensor model to change due to unexpected (or even expected but poorly planned) events. These events range from oversights in the initial development of the physical sensor model, to modifications to the sensor itself. Resultant changes to either the physical sensor model or the sensor-unique image support data formats result in an on-going development and maintenance effort on each user's part.

3.4 If It's Not Broke, Why Fix It?

For a user community to effectively develop, test, and maintain current and new sensor models, they must continually allocate a significant amount of money and manpower. Furthermore, during the implementation of any modifications to an operational sensor model or its sensor-unique image support data format, there is an operational readiness issue at hand. That is, until the user's system is modified, tested, and fielded (which could take months) the imagery produced from that sensor may result in wrong and inaccurate information—unusable for anything but possibly pretty pictures.

To help alleviate this on-going financial burden, and address the operational readiness issue, there have been several previous attempts made to develop a general sensor model that is capable of replacing physical sensor models in the user community. However, they were tailored to certain classes of sensors (e.g. space-borne electro-optical sensors), and they were not complete sensor models in that they were missing the adjustability and/or reliable accuracy prediction capabilities.

There is one exception. The Replacement Sensor Model (RSM) is a complete sensor model that combines all three characteristics, and is flexible enough that it can be used to replace virtually all sensor models in the user community.

4.0 What Is RSM?

The Replacement Sensor Model, RSM, is an adjustable non-physical based mathematical model. Its image support data has one format, and is initially populated using the original physical sensor model and its image support data. Once generated, only the RSM sensor model and its image support data is needed for all geo-positioning processes, i.e. extraction, adjustment (triangulation), and accuracy assessment. If the RSM image support data is generated “up-stream” (i.e. the image libraries, or at the ground processing stations), “down-stream” exploitation systems need only implement one sensor model, RSM.

An up-stream process containing an RSM (image support data) “generator” can process either adjusted or un-adjusted original (physical sensor model) image support data, and then make the RSM image support data (RMSD) available to the down-stream user community (see Figure 11). Thus, down-stream users will be able to use RSM in their applications to support their image exploitation (geo-positioning) processes such as optimal target extraction. An RSM “exploiter” uses the RSM sensor model and the RSM image support data to provide all sensor model functionality to a user’s application (geo-positioning process).

The RSM support data received by the down-stream users faithfully reflects the accuracy of the physical sensor model’s image support data from which it was generated. However, if this accuracy is not adequate, adjustments to the RSM image support data can also be performed by the users if they have access to additional control information.

Note 4.0-1: Keep in mind that the RSM generation process is intended for “up-stream” image providers like ground stations, and image libraries who disseminate imagery. Most users are “down-stream” receivers of imagery that need an RSM exploiter, but not a RSM generator. However, if a “down-stream” user makes adjustments to the RSM image support data, and if properly equipped, they can also disseminate the adjusted RMSD to other users.

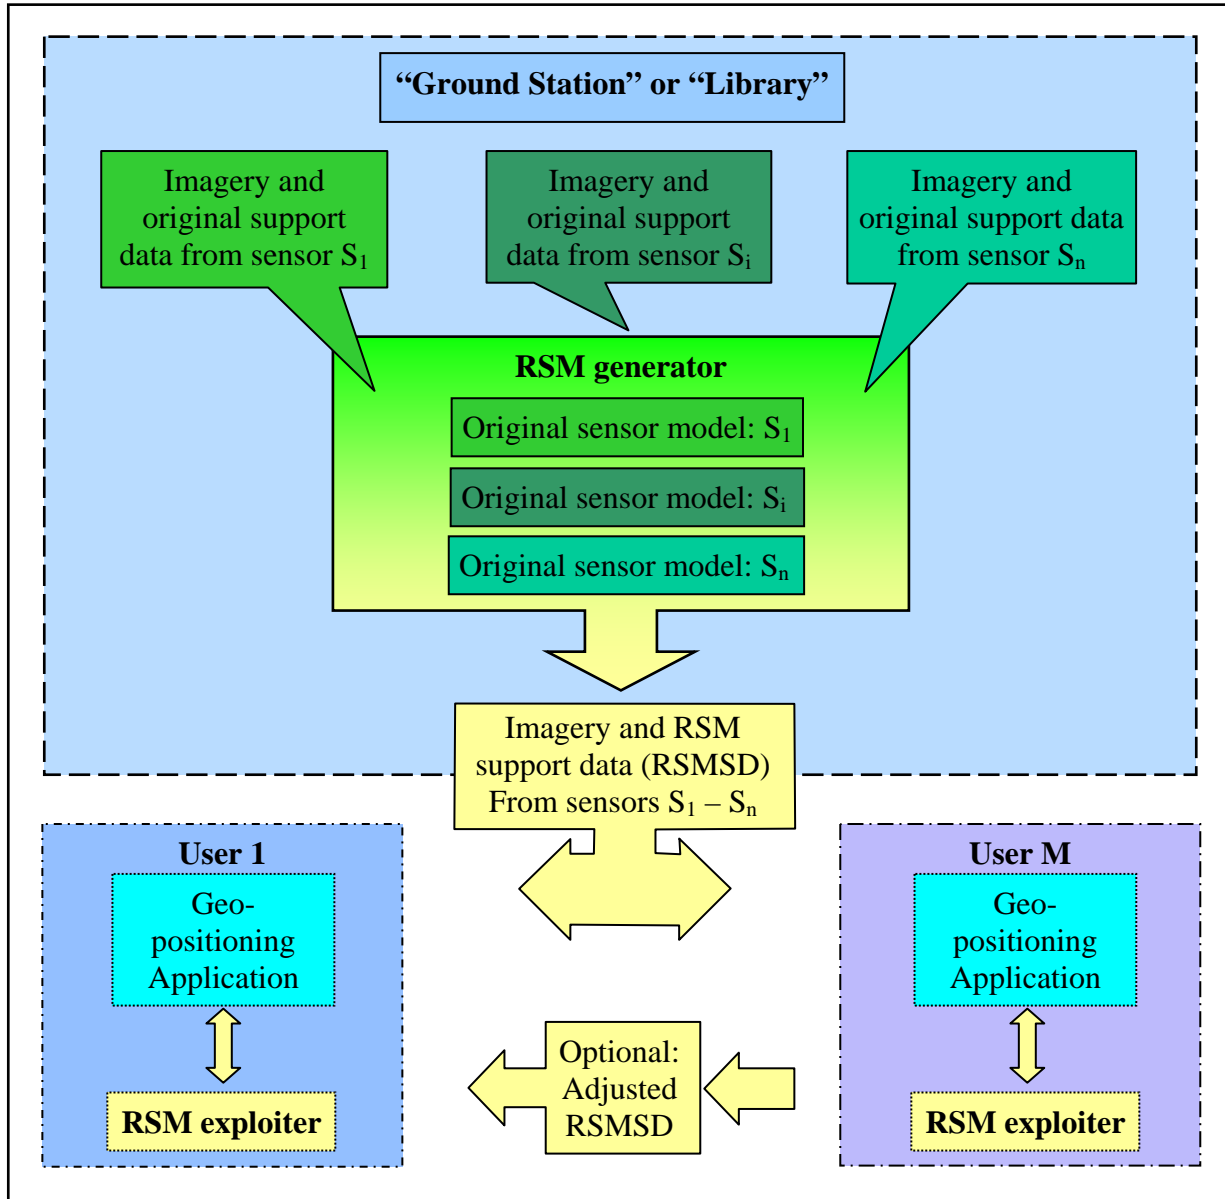


Figure 11: RSM Flow

As will be discussed in more detail, RSM is capable of:

- 1) Replacing a complete physical sensor model ("original" sensor model) for virtually any imaging sensor including: optical (electro-optical) sensors, such as frame and push-broom, SAR sensors, airborne and space-borne sensors (platforms)
- 2) Capturing all of the original sensor model error covariance information in an equivalent form, resulting in virtually identical accuracy predictions
- 3) Providing equivalent mensuration, extraction, and triangulation (support data adjustment) capabilities relative to the original sensor model it is replacing. For example, RSM is able to provide virtually identical multi-image targeting solutions—more about this a little later

Also, since RSM is flexible enough to work for virtually any imaging sensor the RSM support data (RMSD) format is an excellent candidate to become the universal support data format standard for geo-positioning applications.

4.1 Parts... We Have Them All

RSM is a complete sensor model, and has all of the parts we expect a complete original physical sensor model to have. That is, RSM consists of three major components: 1) a ground-to-image function, whose image-to-ground relationship is computed by an iterative inverse, 2) adjustable parameters that affect the ground-to-image function, and 3) an error covariance relative to those adjustable parameters.

4.1.1 Ground-To-Image Function

The RSM uses a flexible non-physical based ground-to-image function. That is, as a mathematical function, it takes in a ground point and returns the corresponding image point. (Recall that an image-to-ground function is also available as an iterative inverse of the ground-to-image function.)

The following Figure 12 is different from the previous figures, such as Figure 3, in that the ground point is directly mapped (functionally) to an image point (ground-to-image), and that there is no physical connection to the sensor as is depicted in Figure 3.

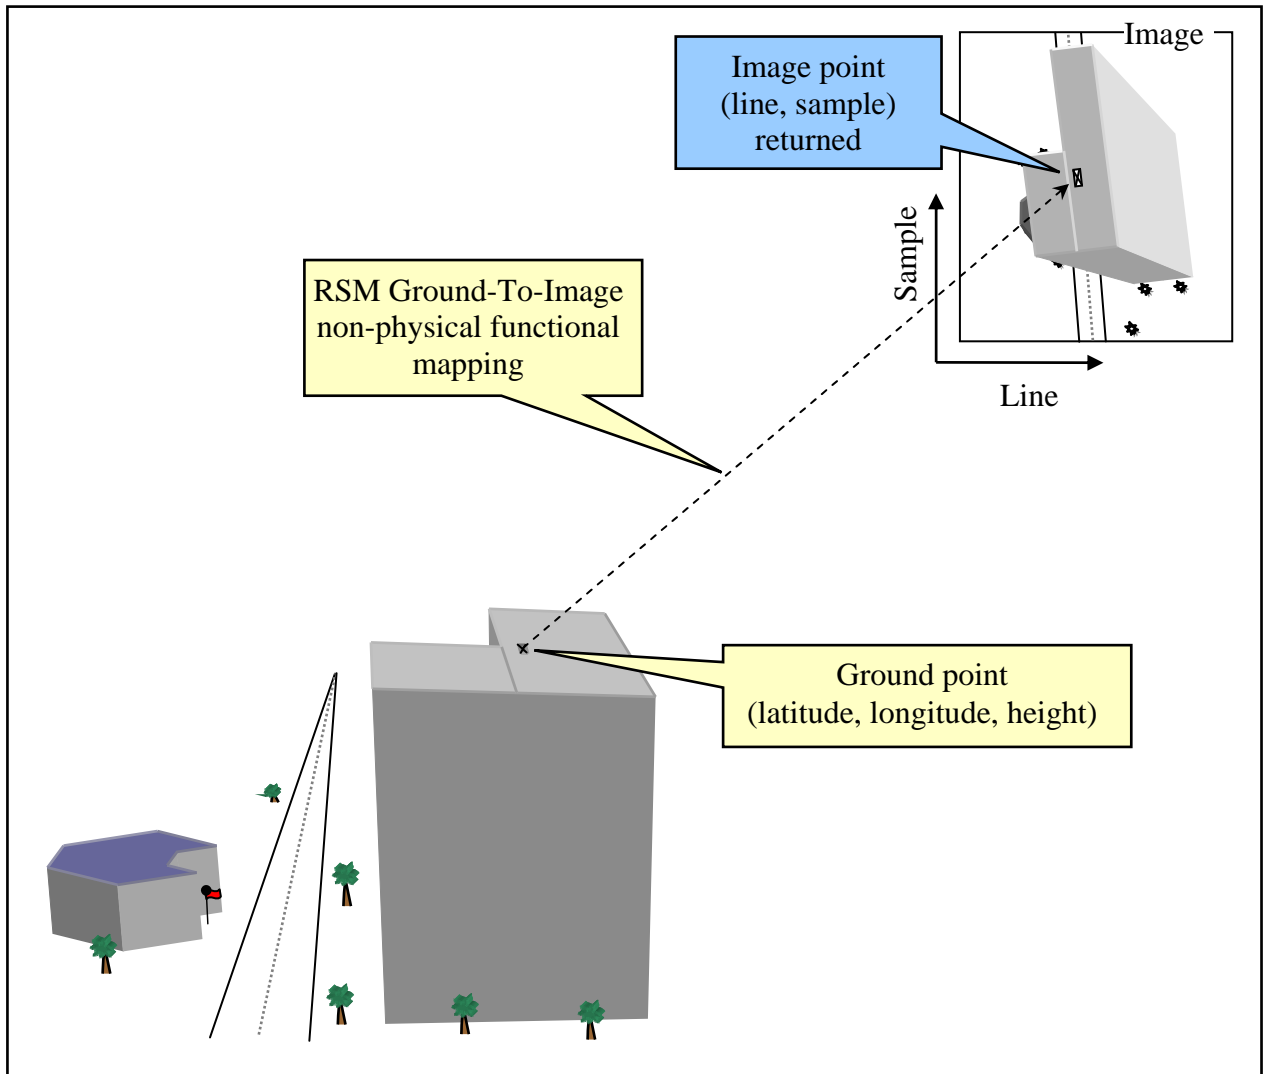


Figure 12: RSM Ground-To-Image

To accomplish this for any sensor, the elements defining the RSM ground-to-image function are generated using the original sensor model's image-to-ground correspondences in an "up-stream" RSM generating process. The RSM generation process then outputs the RSM ground-to-image function (elements) as part of the RSMDS. The RSM ground-to-image function is designed to be either a ground-to-image rational polynomial or a set of grids of ground point-image point correspondences, from which we can interpolate a specific image point for a given ground point, see Figure 13. If the RSM generator builds a polynomial, the RSM ground-to-image function elements are polynomial coefficients, if the RSM generator builds a grid, the function elements are correspondence values.

Warning: sufferers of acute math anxiety should skip the top portion of Figure 13's RSM ground-to-image function depiction, or view it as a piece of squiggly-lined modern art, either way, we apologize for the equational slip.

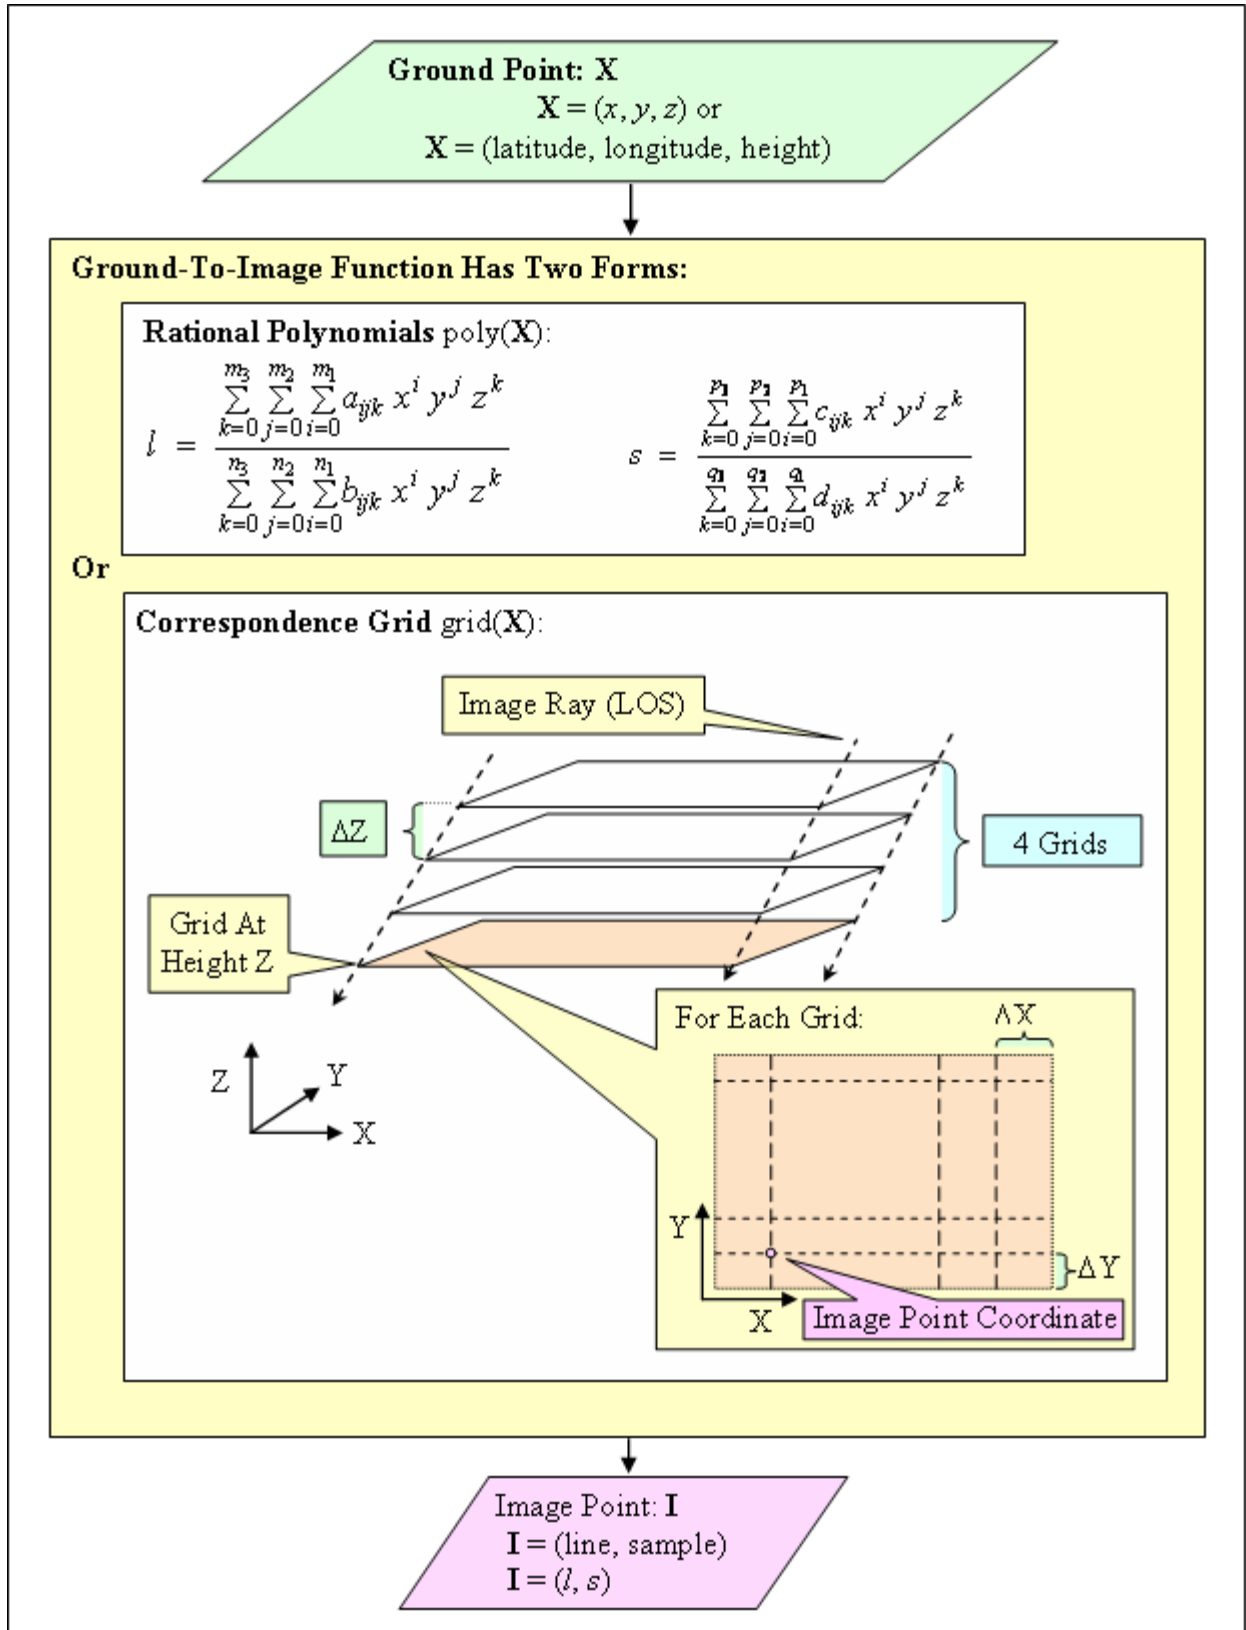


Figure 13: RSM Ground-To-Image Function

UNCLASSIFIED

The RSM rational polynomial function is a robust model able to accommodate a diverse range of polynomial forms independently for both the numerator and denominator of each image coordinate (line and sample), and whose orders are not fixed (for example, at the 3rd degree). For additional fidelity, the RSM rational polynomial function also supports image partitioning, by using multiple polynomials, each of which covers a different section of the overall image.

Note 4.1.1-1: Typically, the rational polynomial's coefficients are generated by a weighted least squares fit to a redundant set of image/ground point correspondences from the original sensor model. These correspondences are generated by calling the original sensor model's image-to-ground function over a grid of (line, sample) locations in the image and at various heights. The corresponding ground locations complete the correspondences.

The other option for the RSM ground-to-image function is a set of correspondence grids from which ground point/image points mappings can be interpolated. This option provides the flexibility needed for accurately matching sensor models when the polynomial fidelity is insufficient, which occurs with some wide field-of-view and some tactical sensors, i.e., imaging geometries that are not nicely modeled by the rational polynomial.

Note 4.1.1-2: Images for which the correspondence grid is needed are determined during RSM generation on a case-by-case basis. For example, even at low altitude (600 m) with a wide field of view (90°) a (large format) frame camera sensor can be modeled well using the RSM polynomial as demonstrated in a study done in conjunction with Purdue University—see reference [4]. This study included nadir (looking straight down) and oblique imaging geometries, with the latter using images that include the horizon.

Note 4.1.1-3: Ok, we lied, there is actually a third RSM ground-to-image function option: using both the polynomials and correspondence grids in a synergistic fashion. That is, the correspondence grids can be used to provide corrections to the RSM polynomial output, and ultimately reduce the RSM image support data bandwidth.

4.1.2 Adjustable Parameters

The RSM adjustable parameters are either generic parameters (A_G) directly affecting the ground-to-image function's input (ground point X), or generic parameters (A_I) directly affecting the ground-to-image function's output (image point I), as depicted in Figure 14. The adjustable parameters transform the ground-to-image function into an "adjustable" ground-to-image function.

The applicable adjustable parameters are selected by the RSM generator from a fixed set in order to replicate the effects of the physical sensor model's adjustable parameters and their error covariance.

The identities of the RSM adjustable parameters selected by the RSM generator are included in the RSMSD. In addition, if an RSM adjustment has taken place, the adjusted (non-zero) values of these adjustable parameters are also included in the RSMSD.

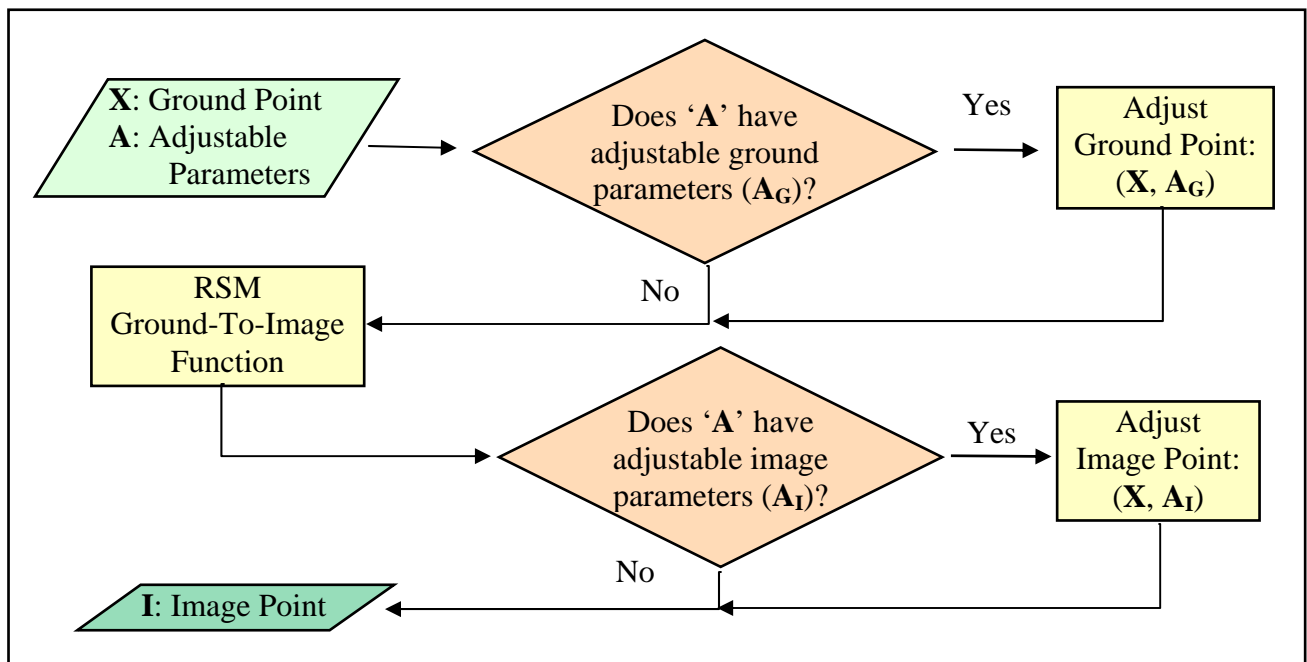


Figure 14: Application of RSM Adjustable Ground-to-Image Function

The number of RSM adjustable parameters selected by the RSM generator (generation algorithm) is usually a few less than the number of corresponding physical sensor model adjustable parameters. For example, the number of RSM adjustable parameters for the optical frame sensor we've been discussing is typically six, but can vary between five and seven depending on the particular sensor and imaging geometry. An electro-optical push-broom sensor might have 15 adjustable parameters in the physical sensor model and 12 for RSM.

Note 4.1.2-1: \mathbf{A}_I contains coefficients of a polynomial-based correction; while \mathbf{A}_G typically contains elements of an affine transformation (translation and small angle rotation). Both their selection and subsequent application are independent of the functional form (polynomial and/or grid) of the RSM ground-to-image function.

4.1.3 Error Covariance

The RSM is designed such that multi-image (also referred to as multi-ray), multi-ground point targeting solutions using RSM are equivalent to multi-image, multi-ground point solutions using the corresponding complete physical sensor models.

Note 4.1.3-1: The term “multi” is general in that there can be anywhere from 1 to n images and 1 to m ground points.

This is possible only because the RSM error covariance is generated in such a way by the RSM generator as to embody the equivalent information contained in the physical sensor model error covariance. This is true whether the error covariance is applicable to one image or a correlated set of images.

Note: 4.1.3-2: If there are a total of n physical sensor model adjustable parameters associated with the image or images, and m RSM adjustable parameters ($m \leq n$), the physical sensor model error covariance is an $n \times n$ matrix and the RSM error covariance is $m \times m$.

This error covariance is utilized to optimally weight multiple measurements from the same image as well as from different, possibly time correlated images; which then yields an optimal solution and provides for reliable solution accuracy predictions.

Note 4.1.3-3: Since RSM is a complete sensor model, the optimal multi-image extraction and adjustment algorithms presented in the 2004 Manual of Photogrammetry (mentioned in Note 3.2.2-1 and Note 3.2.3-2) are applicable for RSM as well.

4.1.4 I Can Do That Too!

Since the RSM is able to provide virtually identical multi-image targeting solutions, including reliable accuracy predictions, it's almost enough to say that the extraction and adjustment processes, and accuracy assessment for RSM are equivalent to the original sensor model's capabilities (see Section 5.2 for more details).

But not quite...

In addition to the extraction and adjustment capabilities previously described for sensor models, the original physical based sensor models can also provide other related capabilities such as: time of image, illumination direction, and the specification of platform position and velocity. Thus, in addition to its three major components, the RSM and its image support data also include an optional time model, illumination model, and platform trajectory model, which completes all of the original physical sensor model functionality.

Note 4.1.4-1: The time model specifies the time a particular image pixel (i.e. image location (line, sample)) was imaged.

Note 4.1.4-2: The illumination model specifies the illumination direction as a function of pixel location within the image. For an electro-optical sensor, the illumination direction corresponds to the direction of the sun; whereas, for a SAR sensor, the direction of radar energy. These illumination directions are needed to support shadow-based mensuration techniques, (i.e., using measurements of shadows to determine an object's height and length, but it's a shady practice).

Note 4.1.4-3: The platform trajectory model supplies the platform trajectory (direction and velocity) as a function of time, useful for various ancillary operational and assessment activities.

...Now we can say the RSM is equivalent to the original physical sensor models in all aspects related to geo-positioning, i.e. extraction, adjustment, error propagation, and mensuration.

4.2 One Size Fits All

The RSM image support data (RSMISD) is contained in the RSM Tagged-Record Extensions (TREs) for the National Imagery Transmission Format (NITF) 2.1, detailed in the reference [2]. Eight TREs have been defined, as shown in Table 1.

Table 1: RSM TREs

CATEGORY: (When included in the RSM TRE Set for an image)		
Description	TRE(s)	Additional Comments
IDENTIFICATION (Always included)		
ID data	RSMIDA	Image identification, coordinate system definitions, etc.
GROUND-TO-IMAGE (Almost always included) One polynomial for entire image is typical (one RSMPCA only) Multiple sections/polynomials and/or multiple sections/grids possible		
Polynomial section ID	RSMPIA	Section identifications in image. Typically one section for entire image. RSMPCI only required if more than one (polynomial) section.
Polynomial coefficients	RSMPCA	Coefficients for ground-to-image polynomial. One RSMPCA per section.
Grid section ID	RSMGIA	Section identifications in image. Typically one section for entire (small) image. RSMGIA only required if more than one (grid) section.
Grid	RSMGGA	Grid of ground point – image point correspondences. One RSMGGA per section.
IMAGE SUPPORT DATA ERROR COVARIANCE (Usually included) One error covariance is typical Both types (one direct and one indirect) error covariance is possible		
Direct covariance	RSMDCA or RSMDCX	Explicit multi-image error covariance. Corresponding image and adjustable parameter identifications.
Indirect covariance	RSMECA or RSMECX	Data to build multi-image error covariance. A priori correlation model used for images from same sensor. Un-modeled error covariance may also be included.
IMAGE SUPPORT DATA CORRECTIONS (Typically not included) Included if RSM image support data has been adjusted		
Adjustable parameter corrections	RSMAPA or RSMAPX	Values of adjustable parameter corrections. Corresponding adjustable parameter identifications. If TRE not included, corrections assumed to equal zero.

Typically, the TRE set for an image consists of three TREs—one each of RSMIDA, RSMPCA, and RSMDCA. The total number of bytes (ASCII characters) for this TRE set is approximately 10k.

Note 4.2-1: The baseline or “A” set of RSM TREs are currently available on the public side of the NTB Web Pages—see reference [2].

Note 4.2-2: The baseline set of RSM TREs have undergone validation by the NITFS Technical Board (NTB) and the Joint Interoperability Test Command (JITC).

Note 4.2-3: As a result of on-going development over the last few years, there is currently an updated set of TREs with improved capabilities that are waiting for adoption. The updated “X” TREs are RSMAPX, RSMDCX, and RSMECX. These latest RSM improvements are documented in the introduction to each updated TRE—see reference [3]. Upon adoption, the “X” TREs will either replace their “A” counterparts or become “B” TREs.

4.3 Advantage...RSM

Now that we’ve discussed RSM details for a while, let’s step back from the trees and look at the forest—where we can discuss some top-level characteristics.

RSM is not just different, it’s revolutionary. Just think how much simpler and more cost effective systems would be if you could just implement one sensor model into an exploitation system, and then only plan on minimal maintenance—if any at all. When new sensors, or modifications to operational sensors, need to be implemented into an exploitation system, there is zero implementation cost—IF the sensor provider or library or anywhere upstream has generated a set of RSM TREs for the new or modified sensor. Furthermore, given the RSM TREs, the new or modified sensor data (for geo-positioning functions) is immediately available, instead of months as it takes now. With RSM, we can now realize no development/integration time or cost associated with the implementation of a new or modified sensor (or sensor data format) into an existing exploitation system.

RSM is also very flexible in that not all of its components (capabilities) need to be utilized for every sensor; however, all components are available in order to support the highest levels of optimal image exploitation. In addition, the RSM also supports proprietary original sensor model development, in that, developers need not provide their detailed original sensor models and image support data to the user community, only the generated RSM image support data.

Thus, the National Geospatial-Intelligence Agency (NGA) is working to have RSM generators in place for government and commercial imaging systems, so that RSM will be universally available—RSM is already available in some U.S. government systems. Furthermore, the NITFS Technical Board (NTB) is working to insure that RSM image support data is an adopted NATO standard and an ISO-compliant international standard, which would complete the revolution. Viva la RSM!

5.0 Evolution

5.1 *The Journey*

Caution: the art of allegory is about to be abused.

...It was another dull and dreary day in 1997 when John Dolloff entered the BAE System's (sensor model) cafeteria. The sign read, "Special Du Jour—replacement sensor model". Can you believe it? It was the same at all cafeterias that day...it was Thursday. Hmm, I guess I'll try it, he thought. It looked good, real good. Taking a huge bite, he gagged. This is like cardboard! Picking up the recipe card entitled "Enjoy Yours At Home," he noticed the recipe was incomplete. 2/3 of the ingredients seemed to be missing. Forever a perfectionist in the kitchen, John went to work. A little here, a little there, adding the missing ingredients, he rebuilt the recipe. Finally, wiping the sweat from his brow...Ah, this is what it's supposed to taste like. Now, to get this recipe into every kitchen...

...all right, a bit of a stretch, but still metaphorically true. In particular:

Prior to RSM, the only replacement sensor models available were tailored to specific types of sensors and were not complete sensor models. However, they served as an important first step, and building upon them, RSM was created. RSM is the only known replacement sensor model that is complete and applicable to virtually all imaging sensors/platforms.

Following its inception, an RSM prototype was developed at BAE Systems. But there was still room for improvements, and RSM needed to be standardized and made available to image developers, distributors, and general users as well.

Enter the NGA.

Over the last six years (2001-2006), with the help (and significant contributions) of Dr. Charles Taylor, RSM has been fully developed under NGA sponsorship. With refinements and verification tests made from the suggestions of Professor Ed Mikhail of Purdue University and additional support received from the National Technical Board (NTB) regarding the specification of RSM image support data, RSM has been brought to fruition.

This development effort has resulted in a mature RSM, documented in the 2004 Manual of Photogrammetry, and in the RSM image support data specification for the NITF 2.1 Tagged Record Extensions (TREs).

5.2 Seeing Is Believing

The claims that the RSM gives virtually identical results in extraction, triangulation, and accuracies, are not lightly made. There have been numerous NGA sponsored studies and validation requirements levied upon RSM and its image support data.

So, what does it mean when we say virtually identical results?

In a series of tests over the last six years, we compared geo-positioning results using the original physical sensor model and its image support data against geo-positioning results using RSM and its image support data (RSMDS). The RSMDS was generated from the original sensor model and original image support data, and all other test conditions were identical. Some of the tests compared results after extraction, some after triangulation (adjustment), and still others after triangulation and extraction (see Figure 15). Both extraction and adjustment tests involved both single images and multiple images.

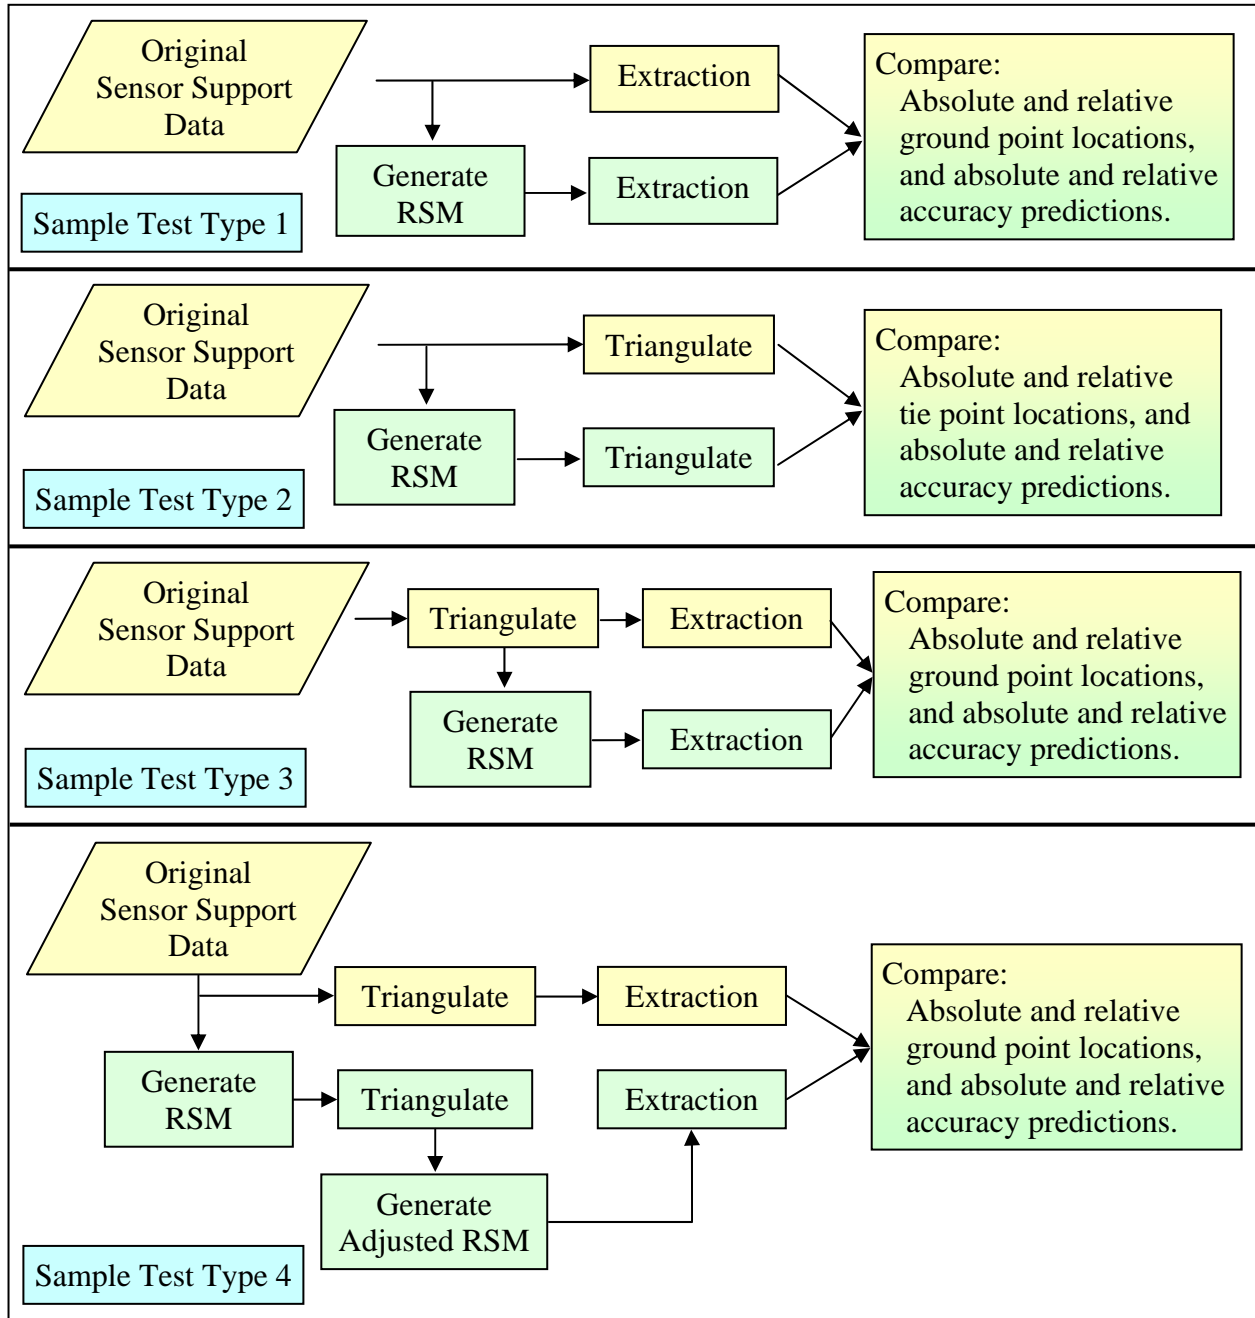


Figure 15: RSM Sample Geo-positioning Tests

From these studies, we have shown that the typical difference in the ground-to-image relationship between the original physical sensor model and its RSM counterpart is less than 1/20 pixel; the difference in (multi-image) target solution coordinates is less than 1/10 meter, and the difference in target absolute and relative accuracy predictions (CE/LE and RCE/RLE) is also less than 1/10 meter (see Table 2).

Similar results are also applicable after support data adjustments of the RSM image support data versus adjustment of the physical sensor model image support data. These adjustment

comparisons correspond to Sample Test Types 2 and 4 in Figure 15 and their results are characterized by “adjustability” in Table 2.

Table 2: Original Sensor Model vs. RSM

Original Sensor Model & Support Data vs. RSM & RSMDS	Difference (typical)	Units
Ground-to-Image	< 0.05	pixel
Image-to-Ground (pixel equivalent)	< 0.05	pixel
Multi-Image Target Solution	< 0.1	meter
Accuracy Prediction (CE/LE and RCE/RLE)	< 0.1	meter
Adjustability	< 0.1	meter

Furthermore, included in these geo-positioning studies, RSM and its support data have been successfully tested and validated for a number of diverse sensors (on both space-borne and air-borne platforms) including: commercial sensors, tactical sensors, EO (optical), SAR (Synthetic Aperture Radar), frame sensors, and push-broom sensors.

Note 5.2-1: little disclaimer:

Through these studies, we have also verified that the original image support data uncertainty can not be excessive in order for any RSM adjustment (triangulation) of the generated RSM image support data to yield virtually the same results as a (hypothetical) original adjustment of the original image support data. (See sample test types 2 and 4, in Figure 15). For example, the position of a space-borne sensor should be known, initially, to within 1000 meters in order to get virtually the same results.

Note 5.2-2: If the original support data uncertainty is excessive, a resection of the original image support data using the original sensor model and external control should first be performed prior to generating RSM image support data, in which case the above disclaimer (Note 5.2-1) is no longer applicable.

Note 5.2-3: This qualifier (Note 5.2-1) regarding original image support data uncertainty is not necessary for RSM-based extraction, only RSM-based adjustment.

6.0 Cooking With RSM

Now that it's clear that you need RSM, the question remains as to how to incorporate it into your system. There are two possibilities. One is to develop your own capability from scratch, which is discussed in the following Section 6.1, and the second is to take advantage of “software modules” that have already been developed and tested, which is discussed in Section 6.2.

However, we'd like to take this moment to reiterate: It has already been pointed out that the RSM model is a general and very flexible sensor model—it had to be in order for it to be applicable to virtually all sensors. But, with this flexibility came complexity. The RSM is not a simple algorithm to develop and implement, at least for RSM generation. Although one can argue that implementing the RSM sensor model from scratch will allow the user to uniquely tailor the deployment of RSM, this must be weighed against the potential development costs, which can be mitigated with the use of the available, fully developed and tested “software modules”.

6.1 Cooking From Scratch

For those who wish to implement RSM capabilities from scratch, you will need to follow the references for more detailed descriptions, as the specific implementation details exceed the scope of this “light” introduction to RSM.

RSM is fully described and documented in the 2004 Manual of Photogrammetry (reference [1]). This treatise includes the algorithms and equations needed to generate RSM support data, such as polynomial coefficients, and the error covariance; as well as, the algorithms and equations needed to exploit the resulting RSM image support data.

Note 6.1-1: Most users are considered down-stream receivers of imagery; in this case, you would only be concerned with implementing the RSM exploitation processes.

Furthermore, there is additional RSM documentation, giving more insight to algorithms, descriptions, and equations specific to the standardized formatting of the RSMSD, contained in the NITF 2.1 RSM TRE definitions currently available on the public side of the NTB Web Pages (reference [2]).

Note 6.1-2: (Recall Note 4.2-3) As a result of recent development, there is also an updated subset of RSM TREs that are waiting for adoption. The latest RSM improvements are documented in the introduction to each new TRE (reference [3]).

Note 6.1-3: Nothing about RSM is proprietary. BAE Systems has published a “Free License Agreement” to (1) assure 3rd parties that they will not infringe BAE Systems’ intellectual property if they build their own RSM exploiter or generator, and (2) gives 3rd parties the right to build their own products which would emulate the standards of RSM.

6.2 *Just Add Water...*

The RSM has been implemented by BAE Systems for the NGA as two application programming interface (API)-driven software modules: the “RSM Generator”, and the “RSM Exploiter”. The initial release of the “RSM Generator” and “RSM Exploiter”, were delivered to NGA at the end of 2004, and a second release in September of 2006.

Included with both the RSM Generator and RSM Exploiter software packages are an installation package, User’s Manual, and the API description (see references [5] and [6]).

Currently, the RSM Generator and Exploiter are being maintained by BAE Systems, under the NGA’s Mensuration Services Program (MSP). Thus, to inquire about obtaining the RSM Generator and/or the RSM Exploiter, contact reference [7]:

Additional information about the RSM Generator and RSM Exploiter can also be found in reference [8].

So, now that we know how to get RSM, which module(s) do I need?

Recall that the RSM generation processes are intended for “up-stream” image disseminators like ground processing stations or image libraries. Thus, if you are disseminating imagery, you are likely in need of the RSM Generator module. However, most users are “down-stream” imagery receivers, and thus only need the RSM Exploiter module.

Note 6.2-1: For users who are also adjusting RSM image support data (see Figure 15, Sample Test Type 4), an RSM Generator “lite” module is also needed if those adjustments are to be disseminated to others in the user community.

6.2.1 RSM Generator

RSM provides “one-stop” shopping for all your sensor model needs. However, in-order for the store to be fully stocked, the RSM Generator (factory) must have access to the original sensor

model for each sensor of interest and corresponding image support data for each image of interest.

For a given image, the corresponding original sensor model and image support data are held by a “sensor model object”, which is basically just a software “container” that supplies sensor model functionality initialized by the appropriate (and accompanying) image support data. The RSM Generator utilizes original sensor model objects in order to generate corresponding RSM image support data as depicted in Figure 16.

The RSM Generator uses the standardized Community Sensor Model (CSM) interface, which supports “plug-in” architecture (see reference [9] for more CSM details) to “communicate” with the sensor model objects provided by the higher-level application. If the higher-level application that hosts the RSM Generator uses an interface other than CSM, a translator must be provided or developed.

Thus, with access to the original sensor model objects, the RSM Generator is able to generate RSM image support data and populate the RSM TREs for incorporation into an NITF image by the upper-level application, thereby allowing the down-stream user community to use the RSM sensor model for all their geo-positioning needs.

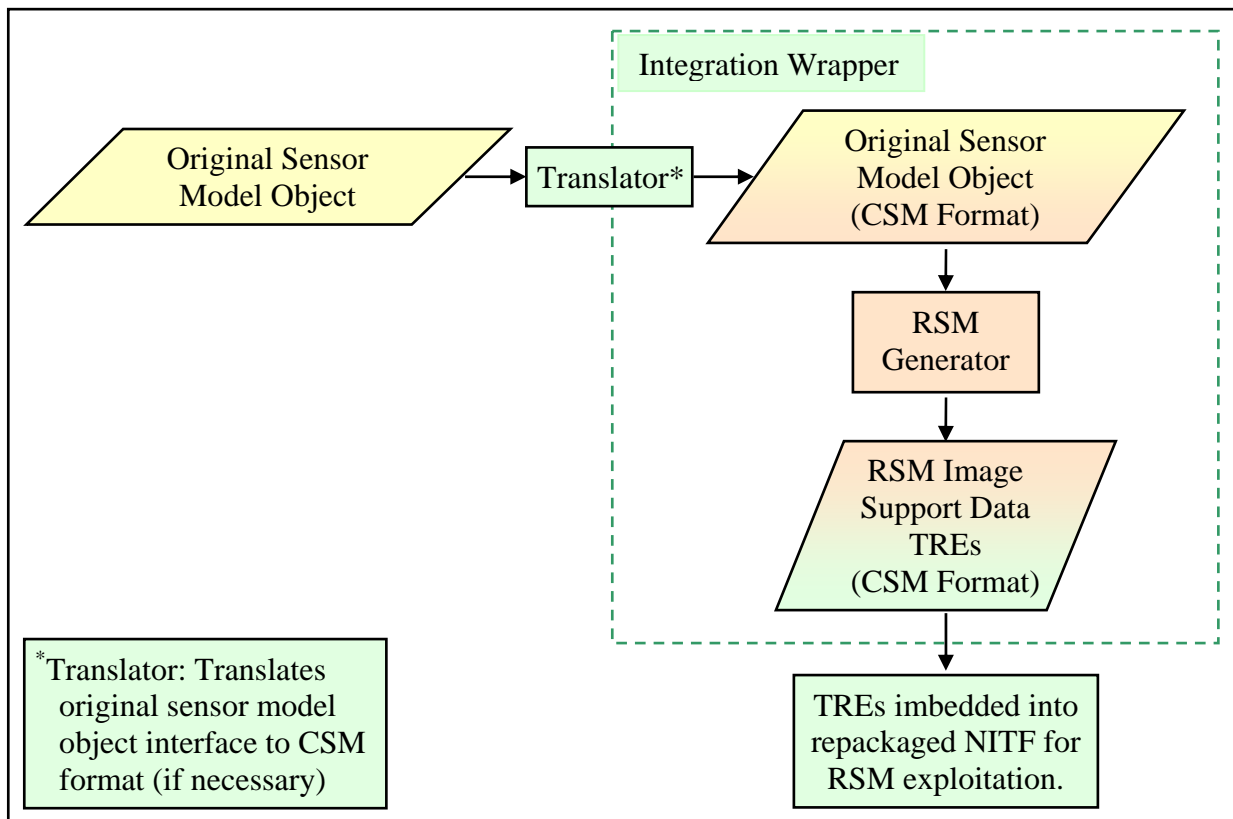


Figure 16: Integration of RSM Generator

The RSM Generator works automatically, and is designed to select and populate the best combination of TREs for many types of sensors/images.

More specifically, the RSM Generator automatically selects the appropriate form for the RSM ground-to-image function, and the appropriate set of RSM adjustable parameters and corresponding error covariance applicable for a specific physical sensor model(s) and its image support data. The RSM Generator then automatically produces the corresponding RSM image support data in standardized NITF 2.1 RSM TRE format.

The RSM Generator produces the appropriate subset of the 8 available TREs (see Section 4.2). Of course, the TREs produced must be compatible with the original sensor model and its image support data. For example, if the original sensor model has a ground/image relationship but not adjustable parameters or error covariance, only RSM TREs associated with the ground/image relationship will be produced.

The resultant RSM TREs are then typically repackaged with the original NITF tags in the NITF image by the upper-level application. By imbedding the new RSM TREs into the original set of NITF tags, there is never a loss of data (e.g., data not directly related to geo-positioning); and down-stream users also have the choice of using the original image support data (if they have the original physical sensor model implemented), or the RSM image support data.

Note 6.2.1-1: Figure 16 is a notional representation for integrating the RSM Generator into an upper-level application host. For details on integration, see references [5], [6] and [9].

6.2.2 RSM Exploiter

As described above, the RSM Exploiter is integrated by down-stream users, and does not need access to the original sensor model or its support data, all it needs are the RSM TREs.

Thus, given any set of RSM TREs retrieved from the NITF image by the upper-level application, the RSM Exploiter automatically provides all corresponding sensor model functionality to an exploitation application through the use of RSM sensor model objects (see Figure 17). The Exploiter also uses the standardized CSM interface to communicate with the sensor model object; thus, if necessary, a translator to and/or from CSM must be provided or developed.

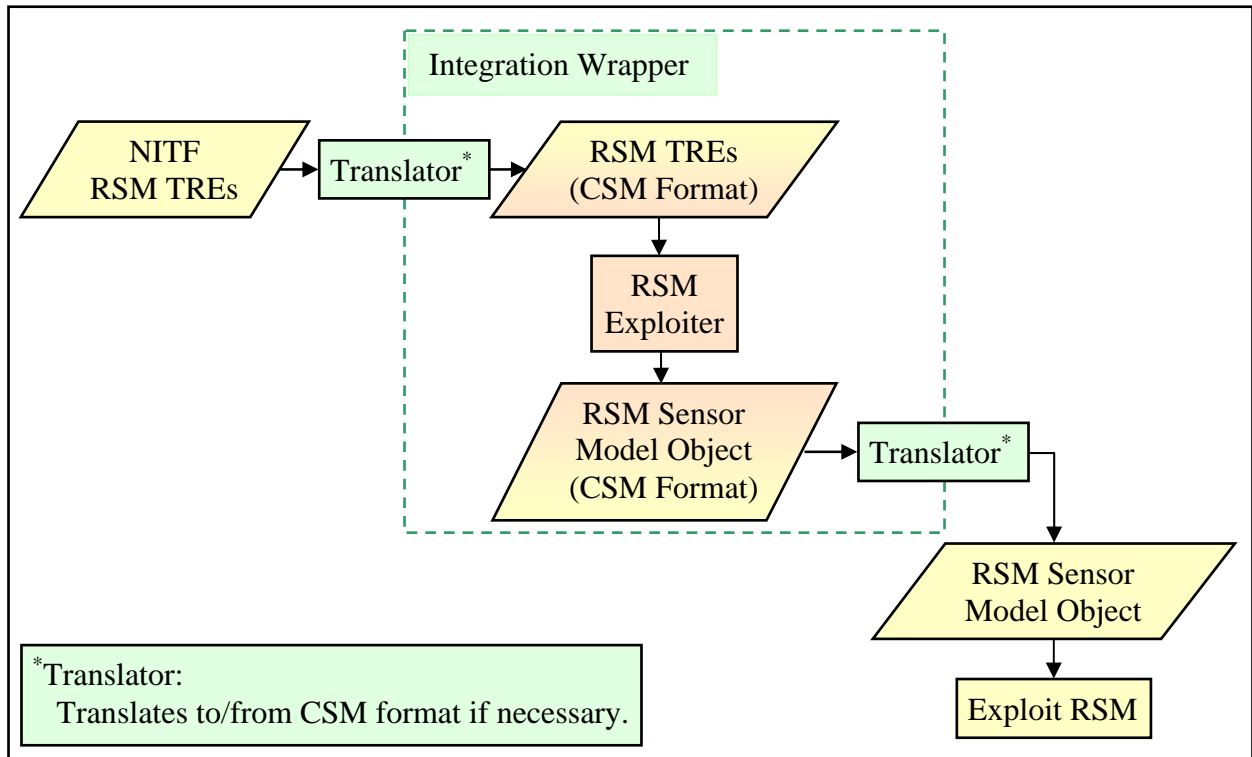


Figure 17: Integration of RSM Exploiter

But, what exactly do we mean by the RSM Exploiter “provides all sensor model functionality”?

It means, that when a user (application program) ...

- ... asks for an image point, and provides a ground point, the RSM Exploiter returns the corresponding image point.
- ... asks for a ground point, and provides an image point and specific height, the RSM Exploiter returns the corresponding ground point.
- ... asks for the partial derivatives of image coordinates with respect to adjustable parameters and provides a ground point, the RSM Exploiter returns the corresponding partial derivatives.
- ... asks for the multi-image support data error covariance for this image and another image, the RSM Exploiter returns the corresponding multi-image error covariance.
- ... and so on,

... everything needed for optimal extraction, mensuration, adjustment, and accuracy prediction.

Note 6.2.2-1: Figure 17 is a notional representation for integrating the RSM Exploiter into an upper-level application host. For details on integration, see references [5], [6] and [9].

7.0 Tributes

We would like to take this opportunity to thank the National Geospatial-Intelligence Agency (NGA) for this document's overall charter and for their support during its preparation as a part of JTW contract # NMA201-01-C-0028.*

*Any mistakes in the use and possible abuse of attempted humor are solely the responsibility of the authors.

8.0 Coda

*It's a small world after all...*and RSM will help make our world smaller, and more accessible.

Well, on that note, we hope you've found this "introduction to RSM" helpful, insightful, and dare we say mildly entertaining? If you're intrigued about RSM, and want to learn more, then we've succeeded, and invite you to continue on to the references.

Thank you for sharing your time with us, and we hope to hear from you soon. In the immortal words of Dale Evans...

"Happy trails to you..."

9.0 References

- [1] Dolloff, John "Replacement Sensor Models", Chapter 11.3 in *Manual of Photogrammetry*, Fifth edition, Chris McGlone, editor, ASPRS, 2004.
- [2] RSM Support Data (RMSD) definitions are currently available on the public side of the NTB Web Pages: <http://www.gwg.nga.mil/ntb/index.html>, in particular:
Dolloff, John, and C. Taylor, *RSM Tagged Record Extensions Spec for NITF 2.1*, 2004, http://www.gwg.nga.mil/ntb/coordinationitems/RSM_NITF_TRE%27s_delivery_July_23_04.pdf.
- [3] Dolloff, John, and C. Taylor, *Proposed Updates to RSM TREs: Extension to the RSM Tagged Record Extensions Spec for NITF 2.1*, September 30, 2006.*[†]
- [4] Taylor, Charles, J. Dolloff, M. Iiyama, and E. Mikhail, *RSM Extraction and Adjustment of Large Field of View Frame Imagery*, NGA September 30, 2006.*[†]
- [5] BAE Systems, *JTW RSM Generator/Exploiter 2.1 User's Manual*, 2006.*
- [6] BAE Systems, *JTW RSM Generator/Exploiter 2.1 Application Program Interface (API)*, 2006.*
- [7] NGA's Mensuration Services Program (MSP) contact for RSM requests:
Scott Krenzke (scott.h.krenzke@nga.mil).
- [8] BAE Systems Geospatial eXploitation (Special) Products:
http://socetset.com/products/products_git.htm.
- [9] *Community Sensor Model (CSM) Technical Requirements Document (TRD)*, Version 2.0, 2005. *

* References [3] through [6] and [9] are available via reference [7].

[†] References [3] and [4] are also available for download via reference [8].