Improving a Template-Based Classifier in a SAR Automatic Target Recognition System by Using 3-D Target Information

Shawn M. Verbout, William W. Irving, and Amanda S. Hanes

■ In this article we propose an improved version of a conventional template-matching classifier that is currently used in an operational automatic target recognition system for synthetic-aperture radar (SAR) imagery. This classifier was originally designed to maintain, for each target type of interest, a library of 2-D reference images (or *templates*) formed at a variety of radar viewing directions. The classifier accepts an input image of a target of unknown type, correlates this image with a reference template selected (by matching radar viewing direction) from each target library, and then classifies this image to the target category with the highest correlation score. Although this algorithm seems reasonable, it produces surprisingly poor classification results for some target types because of differences in SAR geometry between the input image and the best-matching reference image. Each reference library is indexed solely by radar viewing direction, and is thus unable to account for radar motion direction, which is an equally important parameter in specifying SAR imaging geometry. We correct this deficiency by incorporating a model-based reference generation procedure into the original classifier. The modification is implemented by (1) replacing each library of 2-D templates with a library of 3-D templates representing complete 3-D radar-reflectivity models for the target at each radar viewing direction, and (2) including a mathematical model of the SAR imaging process so that any 3-D template can be transformed into a 2-D image corresponding to the appropriate radar motion direction before the correlation operation is performed. We demonstrate experimentally that the proposed classifier is a promising alternative to the conventional classifier.

An automatic target recognition (ATR) system is an integrated collection of algorithms designed to process sensor measurements so that targets can be efficiently detected and identified. The algorithms that comprise an ATR system are applied on a computer and are organized so that human intervention is not required.

An important component of any ATR system is its classifier. The function of the classifier is to take input measurements that represent detected targets and categorize these inputs according to target type. The classifier is designed with the assumption that each input belongs to one and only one category from a predetermined set (e.g., *tank*, *truck*, *gun*), and that the input has certain observable characteristics that aid in its assignment to this category. The classifier output

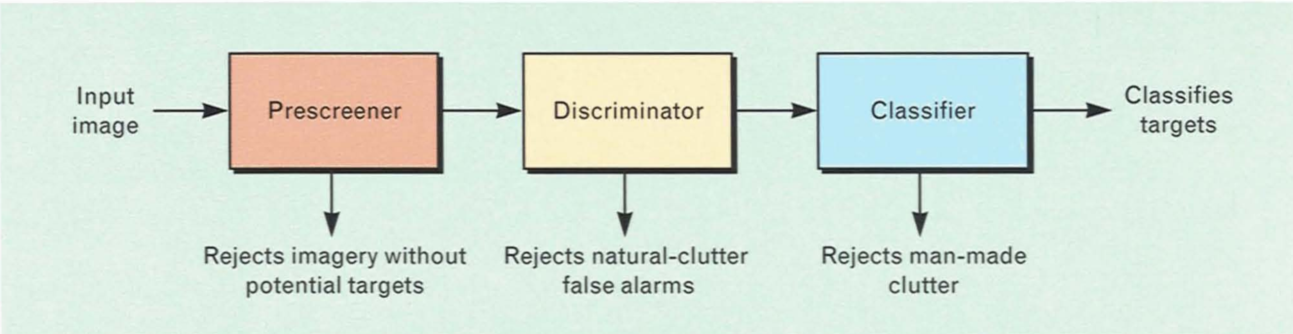

FIGURE 1. Block diagram of the three-stage SAR automatic target recognition system developed by Novak. The input consists of SAR imagery representing many square kilometers of terrain and potentially containing several targets of interest; the output consists of locations and classification labels for these targets. This article proposes an improved version of the classifier stage.

corresponding to each input is an estimate of the correct category label, based on the observable characteristics of the input.

Of course, the issues involved in building a classifier vary according to the kinds of sensor measurements being processed. In this article, we are concerned exclusively with a classifier whose inputs are 2-D synthetic-aperture radar (SAR) images. In particular, we analyze and suggest extensions to the structure of the classifier in the SAR ATR system developed by Leslie M. Novak (see the article entitled "Performance of a High-Resolution Polarimetric SAR Automatic Target Recognition System" in this issue). This ground-based system has been designed to operate in an off-line, experimental setting. It has been rigorously tested over the past five years, and it is one of the first systems of its kind to process large quantities of actual SAR data.

Novak's SAR ATR system is conveniently decomposed into a sequence of three processors: a detector (or *prescreener*), a discriminator, and a classifier (see Figure 1). The detector searches through imagery representing many square kilometers of terrain, and outputs a collection of regions of interest centered at possible target locations. (Each region of interest is a subimage extracted from the original SAR dataset; collectively, all of the regions of interest comprise only a small fraction of this original dataset.) The discriminator applies further processing to distinguish between two kinds of regions of interest: those containing man-made objects (i.e., either targets or man-made clutter) and those containing natural clutter. All regions of interest that appear to contain natural clutter are discarded. Finally, the classifier assigns each remaining region of interest to a predefined target category, or to a none-of-the-above category if the region of interest appears to contain man-made clutter.

Although Novak's ATR system is usually applied to the multiclass target identification problem (i.e., the problem in which two or more kinds of targets must be distinguished), for convenience we consider only the one-class problem in this article. In the one-class problem only one kind of target is of interest; therefore, the classifier output reduces to a simple yes/no decision that indicates whether a target of this kind is present in the region of interest. Figure 2 illustrates the input-output operation of a one-class classifier.

Novak's classifier, which we refer to as the *baseline* classifier, uses a conventional template-matching al-

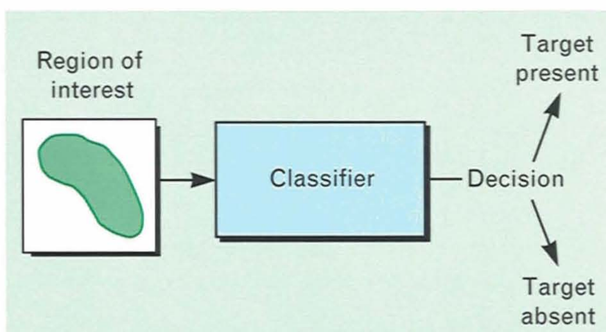

FIGURE 2. Illustration of the input-output operation of the baseline one-class classifier. The input is a subimage, or region of interest, extracted from the original SAR dataset; the output is a decision indicating whether a target of interest is present.

gorithm. The one-class version of this classifier is implemented in the following way. For a particular target of interest, the classifier has a database of stored reference images, each formed by using a different radar viewing direction. The reference image whose associated radar viewing direction best approximates that of the incoming test image (i.e., the incoming region of interest) is called up from the database and is correlated with the test image to generate a match score. If this score exceeds a predetermined threshold, the classifier declares that a target of interest is present.

The template-matching algorithm is attractive because it is readily implemented on a computer and it has an intuitively pleasing structure. For a database formed by using a typical imaging configuration, however, the classifier produces poor results for some target types. This fact is not surprising, because the system was originally designed to process images formed with a fixed SAR geometry, whereas in the most commonly used imaging configurations the SAR geometry is continually changing. In this article, we seek to generalize the structure of the classifier to account for variability in SAR geometry.

SAR geometry can be characterized as a function of two parameters—radar viewing direction and radar motion direction. The baseline classifier does not account for radar motion direction, however, and is therefore equipped with an incomplete set of reference images in its database. The baseline classifier was designed with the assumption that the radar viewing direction is the only parameter that can be varied to produce different images of a target.

In reality, the direction of radar motion is an equally important parameter in defining the SAR imaging geometry, which implies that two images formed with the same radar viewing direction, but with different radar motion directions, will look different. Even though the same physical target scatterers are illuminated in both cases, the 3-D scatterer positions become mapped to two different 2-D SAR image locations. Because the baseline classifier ignores the direction-of-motion parameter, it often correlates a test image and a reference image that are formed with different SAR imaging geometries. These differences in imaging geometry cause the test image and reference image to have dissimilar characteristics, and consequently to have a low correlation score.

An immediate solution to this problem is to include in the classifier database additional reference images formed by using the SAR geometries that are not currently represented. This solution is undesirable, however, because it would require a costly data collection, and it would also increase the storage requirements for the database by roughly an order of magnitude.

In this article we describe a more elegant solution for improving classifier performance than the mere tenfold augmentation of the reference set described above. This new solution, which maintains the traditional template-matching engine, calls for two major modifications to the baseline classifier: (1) the replacement of the present set of 2-D reference images with a set of 3-D templates, and (2) the incorporation of a mathematical model of the SAR imaging process so that any 3-D template can be appropriately transformed to synthesize a 2-D reference image for the correlation operation. Later in this article we describe a novel method for creating 3-D templates from currently existing 2-D target images.

The body of the article is divided into three major sections. In the first section we describe in detail how the baseline classifier works. In the second section we demonstrate the problem with this classifier and explain why this problem exists. In the third section we describe specifically how we can modify the baseline classifier to improve its overall performance. Finally, we summarize the key points of the article and suggest directions for future work toward improving classification performance in a SAR ATR system.

How the Baseline Classifier Works

The algorithm used by the baseline classifier is described schematically in Figure 3. As shown in this figure, the input to the classifier consists of two components. The first input component is a 2-D test image representing a region of interest from the original SAR dataset. As mentioned above, this image has passed through the first two stages of the ATR system (i.e., the detection and discrimination stages) and thus contains an object that appears sufficiently targetlike to be considered for classification. The second input component is a pair of angle values that

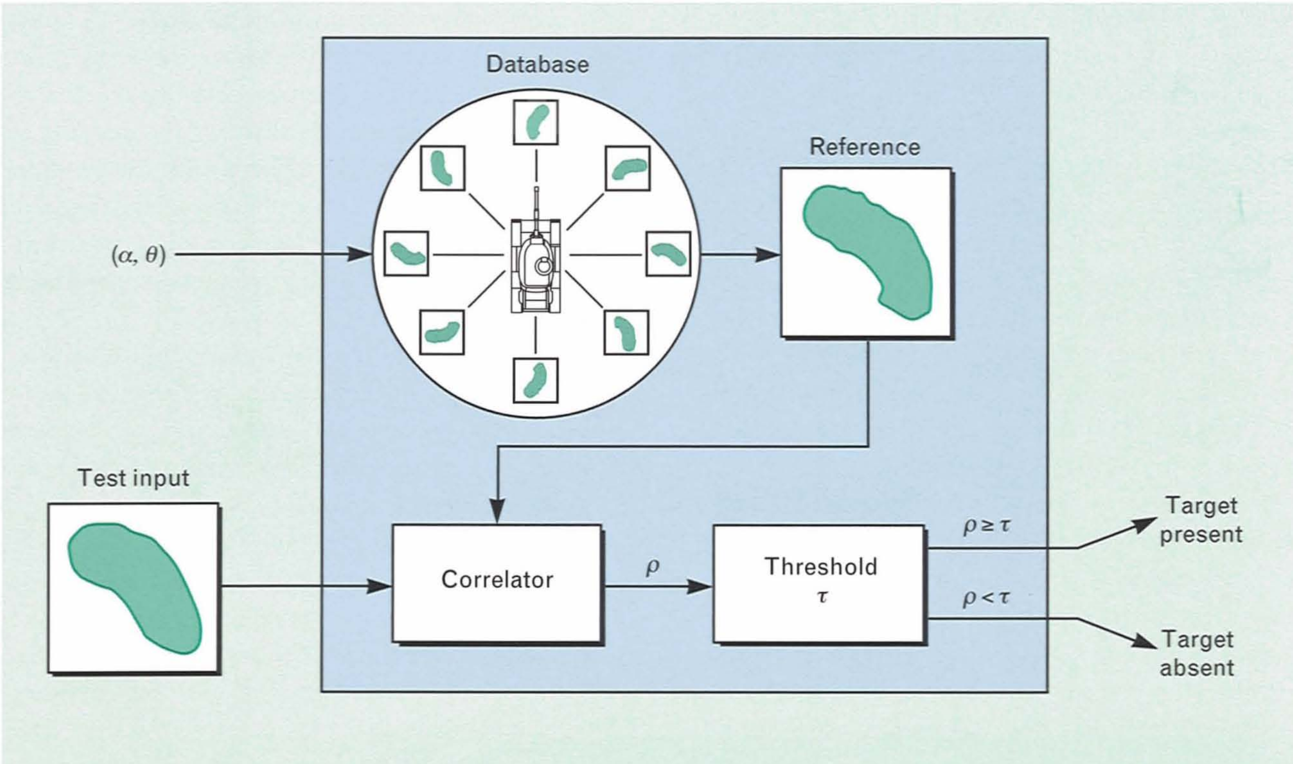

FIGURE 3. Schematic description of the algorithm used by the baseline one-class classifier. The classifier uses the aspect angle α and the depression angle θ to select a 2-D reference image from the database. This reference image is then correlated with the input test image; if the correlation score ρ is greater than or equal to the threshold τ, the *target present* decision is declared.

define the radar viewing direction with respect to the imaged object. These values are estimates of the angles α (the *aspect*) and θ (the *depression*), which are defined pictorially in Figure 4.

Because the database is conveniently indexed according to these two radar viewing angles, the classifier can readily select the reference image whose aspect and depression are closest to the input estimates of α and θ computed for the test image. Once the appropriate reference image is selected, it is scaled so that the sum of the squares of its pixel values is equal to unity; the test image is also scaled in this way. Next, the normalized test and reference images are correlated to yield a correlation score ρ whose value is between 0 and 1.

This correlation operation is mathematically defined in the following way. Let us assume that the test and reference images are equal in size, each having M cells in the range dimension and N cells in the cross-range dimension. Let the function $T(\cdot,\cdot)$ be defined for integer values of its arguments such that $T(m,n)$ is equal to the amplitude of the test image at the range and cross-range location (m,n) for $1 \leq m \leq M$ and $1 \leq n \leq N$, and is equal to zero for all other values of m and n. Let the function $R(\cdot,\cdot)$ be defined analogously with respect to the reference image. Then the correlation score ρ for the two images is defined by

$$\rho = \max_{i,j}\left\{\frac{1}{s}\sum_{m=1}^{M}\sum_{n=1}^{N} T(m,n)\, R(i+m, j+n)\right\},$$

where s is the overall normalization factor given by

$$s = \sqrt{\sum_{m=1}^{M}\sum_{n=1}^{N}\left[R(m,n)\right]^2}\sqrt{\sum_{m=1}^{M}\sum_{n=1}^{N}\left[T(m,n)\right]^2}.$$

As shown in Figure 3, the classifier declares that a target is present in the test image only if ρ is greater than or equal to the preselected threshold τ.

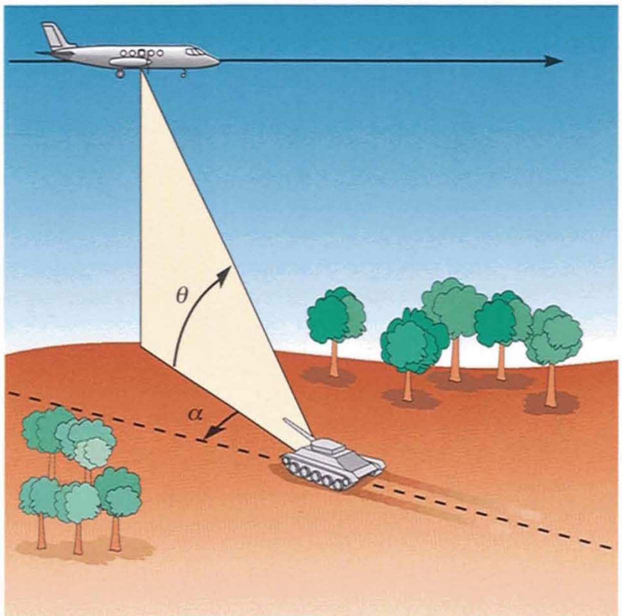

FIGURE 4. Pictorial definition of the aspect angle α and the depression angle θ. These angles specify the radar viewing direction with respect to the imaged object.

We can more clearly understand the fundamental problem with the baseline classifier by analyzing how the target reference images are generated for the classifier database. Each reference image is formed from data collected by the Lincoln Laboratory millimeter-wave airborne radar [1]. Once a target of interest is deployed in an open area, the data are collected by using a special mode of the radar known as *spotlight* mode, which is illustrated in Figure 5. In this mode, the airplane flies in a straight line at constant altitude, and the radar antenna is steered continuously so that it always points in the direction of the target.

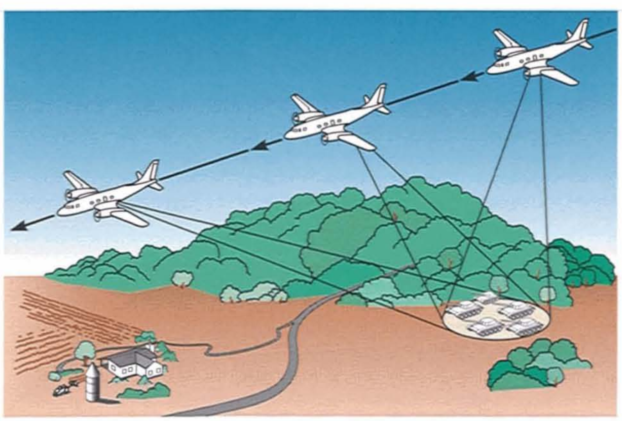

FIGURE 5. Imaging configuration for spotlight-mode SAR. In this mode the airplane moves in a straight line at a constant altitude, while the antenna is steered continuously so that it always points at a fixed patch of terrain.

With the radar beam illuminating the target like a spotlight throughout the flight, a new image of the target can be formed approximately every degree of azimuth. Each new image can be used as a reference for the classifier database. A set of reference images representing 360° of aspect coverage is created by flying four such linear paths to view the target from all sides, as shown in Figure 6.

Although spotlight mode is not the only mode that can be used to generate reference images, it is the most convenient and efficient mode for imaging a target at a variety of radar viewing angles. To see why this statement is true, consider the database of spotlight imagery that can be generated by flying the basic pattern shown in Figure 6 at a sequence of increasing altitudes. Clearly, if the difference between successive flight-path altitudes is small enough, then the database will contain a representative image of the target that is close to any desired aspect-angle and depression-angle pair. Moreover, this complete coverage is obtained without ever having to move the target. In spite of the many advantages to using this kind of data-collection procedure, there is a serious deficiency associated with it. This deficiency is analyzed in detail in the next section.

Why the Baseline Classifier Needs Improvement

Now that we have discussed the method used to generate target reference images for the database, we are better equipped to analyze why the baseline classifier can make a gross error in categorizing an input test image. In this section we explain how such misclassifications occur, even though the database is densely populated with target reference images from all desired radar viewing directions.

We begin by using Figure 7 to demonstrate what is wrong with the baseline classifier. Figure 7(a) shows an optical photograph of an M48 tank, and Figures 7(b), 7(c), and 7(d) show three simulated SAR images of the tank. The SAR images are color coded with a scale that makes a gradual transition from black (low intensity) to green (medium intensity) to white (high intensity). In each SAR image, the front part of the tank is pointing toward the upper left corner of the image. All three SAR images shown in this figure were formed by using the same aspect angle and depression angle. In other words, the same scatterers on the target were illuminated by the radar from the same viewing direction for each image. Note, however, that the images look dramatically different. This phenomenon contradicts the key design assumption that fixing the radar viewing direction uniquely specifies the SAR image of the target.

To understand how the existence of three such images affects the classifier, let us assume that Figure 7(b) is a stored reference image and that Figure 7(d) is an incoming test image. Because the two images were formed by using exactly the same radar viewing directions, the image in Figure 7(b) would be chosen as the reference image most likely to match the test image. But because the two images are so dissimilar, their correlation score would be low, and consequently the test image could be erroneously labeled as containing no target.

The only difference between the SAR imaging configurations used to generate the images in Figures 7(b), 7(c), and 7(d) was the direction in which the radar was moving with respect to the viewing direction. This change alone is sufficient to yield SAR images that look quite different, and yet the direction-of-motion parameter has been completely ignored in the design of the baseline classifier.

To see how this parameter directly affects the appearance of a SAR image, we devote much of this section to the description and application of a widely used mathematical model of the SAR imaging process. In particular, we model the SAR transformation as a *projection* of the 3-D distribution of target scatterers onto a 2-D image plane, and we demonstrate the usefulness of this model by a simple example.

The projection model is conceptually important

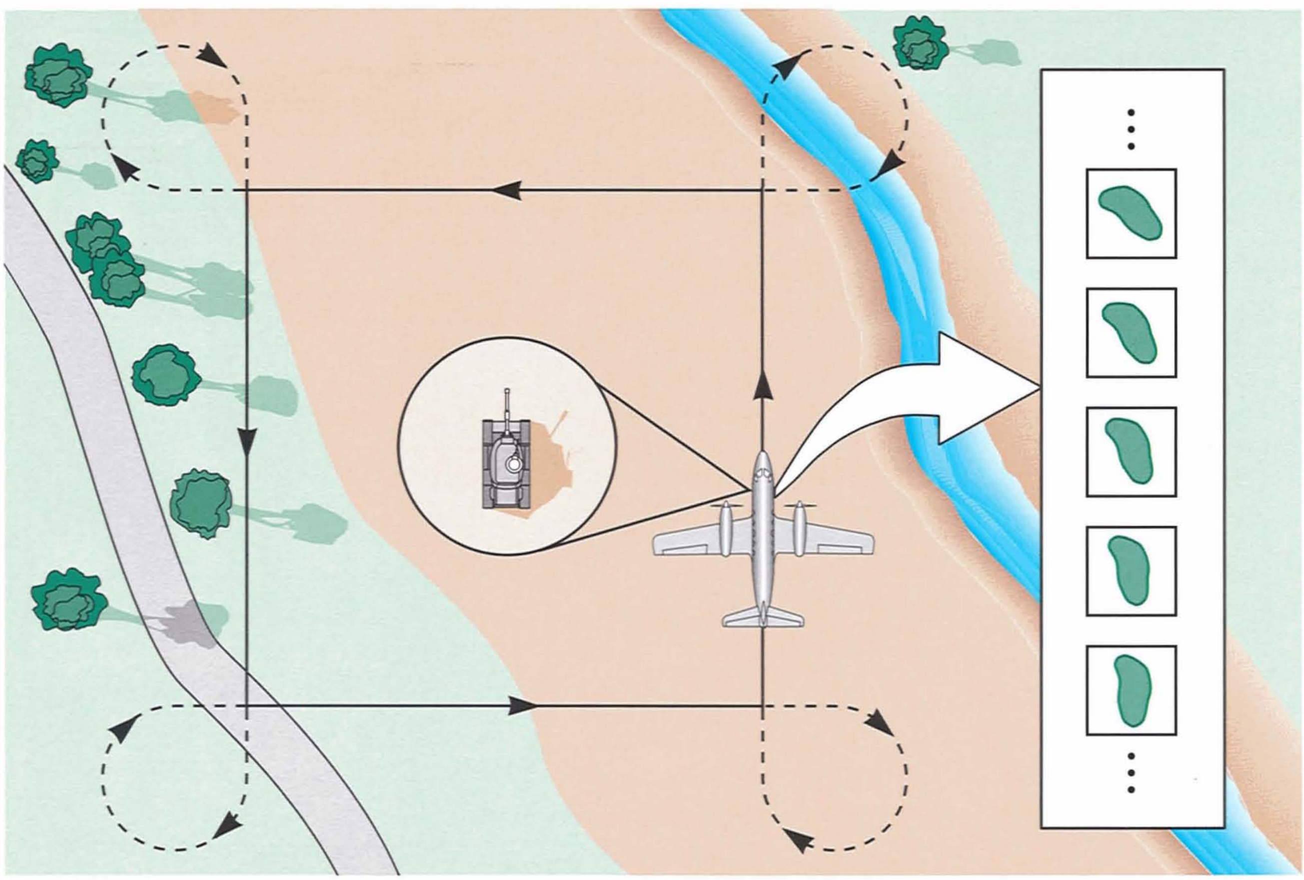

FIGURE 6. Top view of the flight path used to create a set of target reference images representing 360° of aspect coverage. A sequence of the generated reference images is shown notionally at right.

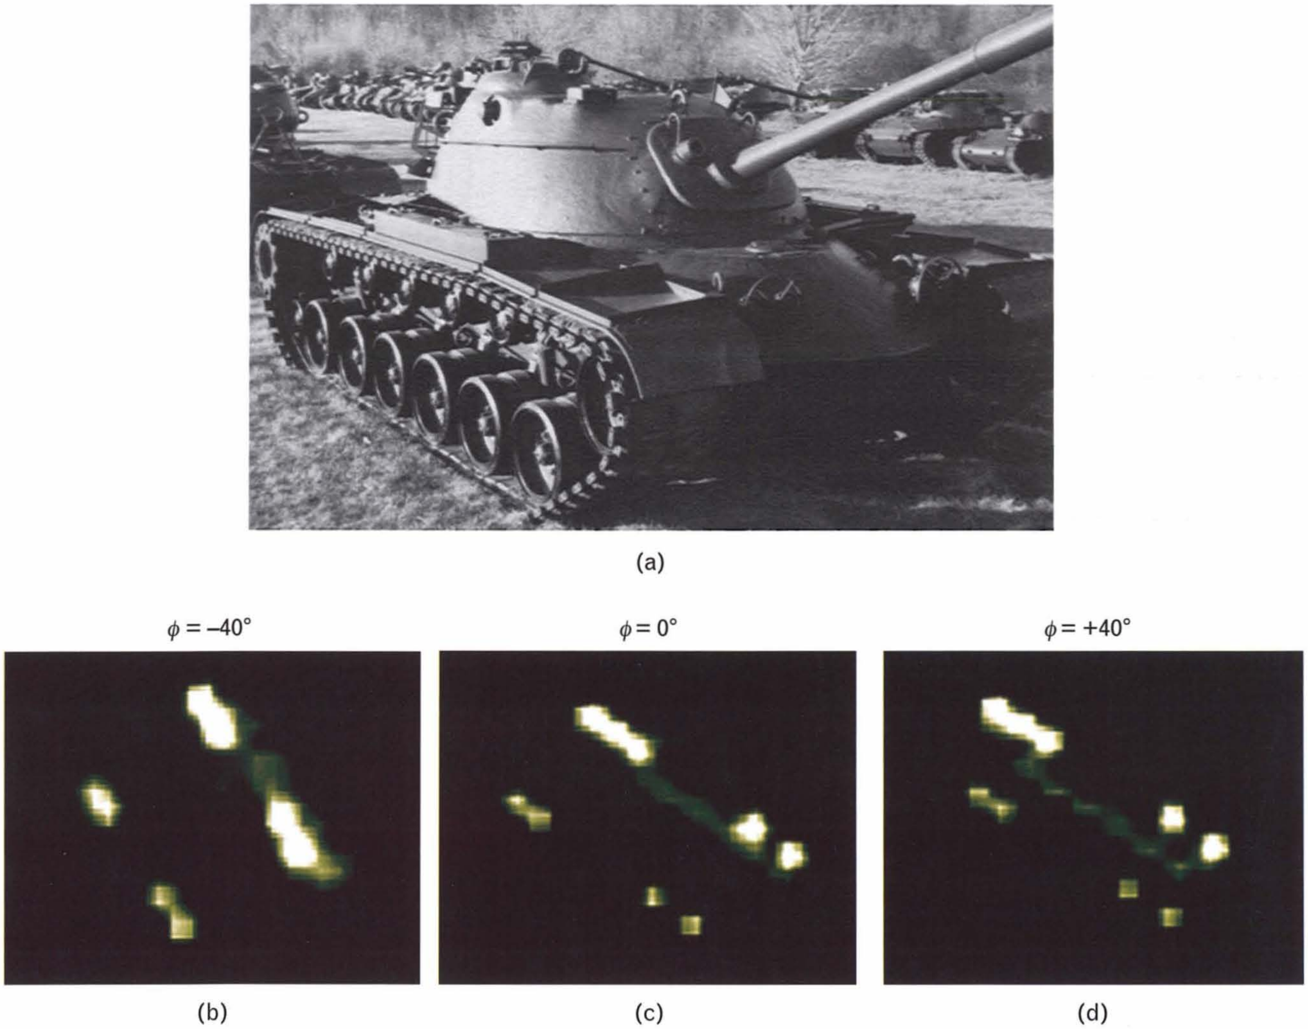

FIGURE 7. (a) Photograph of an M48 tank. The computer-generated images in (b), (c), and (d) are simulated SAR images of an M48 tank, all created with the same aspect and depression angles but with different squint angles ϕ, where ϕ specifies the direction in which the radar is moving with respect to the viewing direction. The SAR images are color coded with a scale that makes a gradual transition from black (low intensity) to green (medium intensity) to white (high intensity). In each SAR image, the front part of the tank is pointing toward the upper left corner of the image. (The SAR images were generated with the SARTOOL signature simulation software developed by The Analytical Sciences Corporation.)

for the remainder of the article, particularly in the final section in which we incorporate this model into an improved version of the baseline classifier. We now prepare to introduce the projection model with some fundamental definitions associated with the SAR imaging process.

Description of SAR Imaging Geometry

Figure 8 illustrates the basic elements that define the spotlight SAR imaging geometry. In this figure, we see the airborne radar as it moves in a straight line while its antenna is steered to illuminate a fixed ground location known as the *aimpoint.* We have imposed a mathematical structure on this geometry by using a Cartesian coordinate system (known as the *world coordinate system*) whose origin coincides with the aimpoint, and whose coordinate locations (x, y, z) are measured in terms of the unit basis vectors $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$, and $\hat{\mathbf{z}}$ (shown in blue). In this system, we use the convention that $\hat{\mathbf{y}}$ points in the direction of radar motion. Also, note that in the vicinity of the aimpoint we model the local earth surface as a *ground plane* defined by the equation $z = 0$.

The line that passes through the radar position and

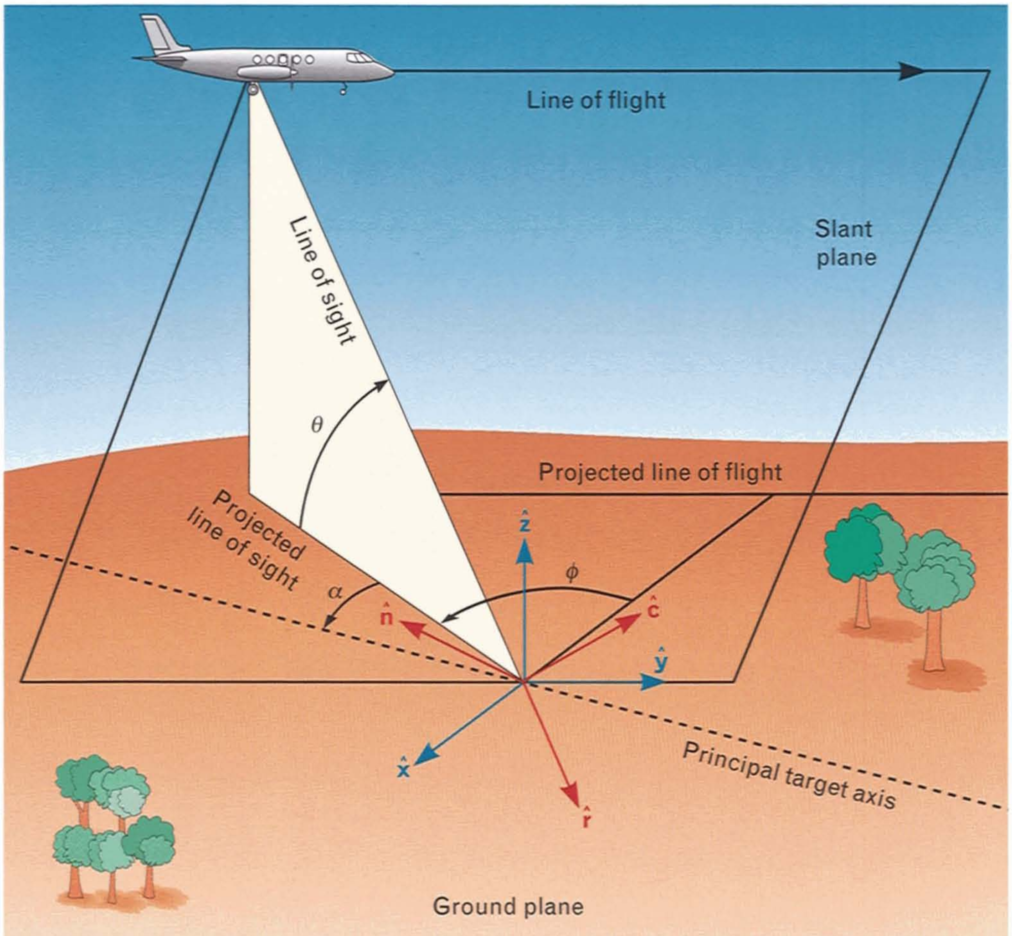

FIGURE 8. Illustration of the basic elements that define the spotlight SAR imaging geometry. Two coordinate systems are represented that share a common origin coinciding with the aimpoint. The world coordinate system (in blue) is defined by the unit vectors $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$, and $\hat{\mathbf{z}}$. The radar coordinate system (in red) is defined by the unit vectors $\hat{\mathbf{r}}$, $\hat{\mathbf{c}}$, and $\hat{\mathbf{n}}$.

coincides with the direction of radar motion is called the *line of flight*, while the line that intersects both the radar position and the radar aimpoint is called the *line of sight*. The plane defined by these two lines is called the *slant plane*; the projections of these two lines downward onto the ground plane are called, respectively, the *projected line of flight* and the *projected line of sight*.

Throughout our discussion, we often refer to the three angles α, θ, and ϕ, which are also shown in Figure 8. The angles α and θ were introduced earlier in the text in Figure 4, but we define them more precisely here. The aspect α is the angle between the principal target axis in the ground plane and the projected radar line of sight. The depression θ is the angle between the radar line of sight and the ground plane. The angle ϕ, which is known as the *squint*, is defined as the geometric complement of the angle between the projected line of flight and the projected line of sight. Note that the two angles α and θ determine the radar viewing direction, whereas ϕ determines the direction of radar motion relative to the viewing direction.

Each of these definitions is useful as we develop our mathematical abstraction of the SAR imaging process. Specifically, we use the above definitions to describe the SAR transformation as a projection of a 3-D distribution of target scatterers onto the 2-D slant plane. In the appendix we give a detailed justification for representing the SAR transformation this way, based on the actual physical quantities measured by a SAR. In the next subsection, however, we omit this justification and merely provide a concise mathematical description of our SAR imaging model. Following this description, we give an application of our SAR imaging model in the form of a simple visual example.

A Mathematical Model for SAR Imaging

We begin our description of the SAR transformation by introducing the *radar coordinate system* (shown in red) in Figure 8. The origin of this Cartesian coordinate system coincides with the aimpoint, and the coordinate locations are represented in terms of the unit basis vectors $\hat{\mathbf{r}}$ (the *range vector*), $\hat{\mathbf{c}}$ (the *cross-range vector*), and $\hat{\mathbf{n}}$ (the *slant-plane normal vector*). This coordinate system can be defined in terms of the basic elements of the spotlight SAR geometry defined above.

We begin with the observation that the range vector $\hat{\mathbf{r}}$, which points in the direction of the radar line of sight, can be expressed in world coordinates as

$$\hat{\mathbf{r}} = \begin{bmatrix} \cos\phi\cos\theta \\ \sin\phi\cos\theta \\ -\sin\theta \end{bmatrix}.$$

To check that this expression is correct, the reader can easily verify the following three properties of $\hat{\mathbf{r}}$: (1) $\hat{\mathbf{r}}$ is unit length, (2) the projection of $\hat{\mathbf{r}}$ onto the x-y plane is rotated counterclockwise by the angle ϕ with respect to the x-axis, and (3) $\hat{\mathbf{r}}$ is tilted downward by the angle θ with respect to the x-y plane.

By using the vector $\hat{\mathbf{r}}$ and the world-coordinate basis vector $\hat{\mathbf{y}} = [\,0 \;\; 1 \;\; 0\,]^T$, both of which lie in the slant plane, we can construct the slant-plane normal vector $\hat{\mathbf{n}}$ with the cross-product formula

$$\hat{\mathbf{n}} = \frac{\hat{\mathbf{r}} \times \hat{\mathbf{y}}}{\left\|\hat{\mathbf{r}} \times \hat{\mathbf{y}}\right\|} = k \begin{bmatrix} \sin\theta \\ 0 \\ \cos\phi\cos\theta \end{bmatrix},$$

where k is the normalizing constant required to make $\hat{\mathbf{n}}$ a unit-length vector. The value of k is given by

$$k = \frac{1}{\sqrt{\sin^2\theta + \cos^2\phi\cos^2\theta}}. \qquad (1)$$

Because the cross-range vector $\hat{\mathbf{c}}$ must be perpendicular to both $\hat{\mathbf{r}}$ and $\hat{\mathbf{n}}$, it is constructed by using the formula

$$\hat{\mathbf{c}} = \hat{\mathbf{n}} \times \hat{\mathbf{r}} = k \begin{bmatrix} -\sin\phi\cos\phi\cos^2\theta \\ \cos^2\phi\cos^2\theta + \sin^2\theta \\ \sin\phi\cos\theta\sin\theta \end{bmatrix}.$$

From this coordinate-system construction, we see that the vectors $\hat{\mathbf{r}}$ and $\hat{\mathbf{c}}$ form an orthonormal basis for the slant plane, just as the vectors $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ form an orthonormal basis for the ground plane.

Now we consider imaging a point reflector at location $\mathbf{p} = [\,p_x \; p_y \; p_z\,]^T$ in the world coordinate system. We can express this point in the radar coordinate system by using the standard dot product to project the point $\mathbf{p}$ onto each of the unit basis vectors $\hat{\mathbf{r}}$, $\hat{\mathbf{c}}$, and $\hat{\mathbf{n}}$. The resulting vector $\mathbf{q}$ can be written in radar coordinates as

$$\mathbf{q} = \begin{bmatrix} q_r \\ q_c \\ q_n \end{bmatrix} = \begin{bmatrix} \mathbf{p}\cdot\hat{\mathbf{r}} \\ \mathbf{p}\cdot\hat{\mathbf{c}} \\ \mathbf{p}\cdot\hat{\mathbf{n}} \end{bmatrix}.$$

According to our basic model for the SAR transformation, we must now project the 3-D vector $\mathbf{q}$ onto the 2-D slant plane to obtain its location in the SAR image. We can do this projection by retaining the first two components of $\mathbf{q}$ and neglecting the third component, because the entire third dimension of the radar coordinate system becomes collapsed in the projection process. This procedure gives the slant-plane coordinates of the original point reflector as

$$\begin{bmatrix} q_r \\ q_c \end{bmatrix} = \begin{bmatrix} p_x\cos\phi\cos\theta + p_y\sin\phi\cos\theta - p_z\sin\theta \\ \\ -kp_x\sin\phi\cos\phi\cos^2\theta \\ \quad + kp_y(\cos^2\phi\cos^2\theta + \sin^2\theta) \\ \quad + kp_z\sin\phi\cos\theta\sin\theta \end{bmatrix}. \qquad (2)$$

We can use the above expression for the range and cross-range coordinates of a point to show mathematically that when the point is imaged at a fixed aspect angle and depression angle, but at different squint angles, it will appear at different SAR image locations.

To demonstrate this concept, we conduct two imaging experiments in which we keep the aspect and depression angles constant but allow the squint angle to vary. In particular, for the first imaging experiment we use the values $\alpha = 0$, $\theta = \pi/4$, and $\phi = 0$; for the second experiment we use the values $\alpha = 0$, $\theta = \pi/4$, and $\phi = \pi/4$. In each experiment, once these angles have been fixed, we consider imaging a point **p** that lies on the principal target axis in the ground plane at a distance of one unit from the aimpoint. Based on the diagram in Figure 8, this point must have the coordinates

$$\mathbf{p} = \begin{bmatrix} p_x \\ p_y \\ p_z \end{bmatrix} = \begin{bmatrix} \cos\phi \\ \sin\phi \\ 0 \end{bmatrix}. \tag{3}$$

For the first imaging experiment (with the angles $\alpha = 0$, $\theta = \pi/4$, $\phi = 0$), we can readily verify from Equation 1 that $k = 1$, and from Equation 3 that

$$\mathbf{p} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}.$$

We can now compute the slant-plane coordinates of **p** by substituting these numerical values into Equation 2. This computation yields the 2-D slant-plane location $\mathbf{s}_1$ associated with the first set of imaging angles; this location is given by

$$\mathbf{s}_1 = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ 0 \end{bmatrix}.$$

For the second imaging experiment (with the angles $\alpha = 0$, $\theta = \pi/4$, $\phi = \pi/4$), we find that $k = 2/\sqrt{3}$, and

$$\mathbf{p} = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \\ 0 \end{bmatrix}.$$

As before, we compute the slant-plane coordinates of **p** by substituting these values into Equation 2. This yields the 2-D slant-plane location $\mathbf{s}_2$, given by

$$\mathbf{s}_2 = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{6}} \end{bmatrix}.$$

For the two different squint angles above, the range coordinate of the imaged point remains constant but the cross-range coordinate changes dramatically. This observation is an example of the more general result that, given a fixed radar viewing direction, a change in the squint angle causes the cross-range coordinate of a point to change. Thus the above example provides quantitative proof that the slant-plane location of a point is not uniquely determined by the aspect and depression angles alone.

In the next subsection, we give a simple qualitative example that visually demonstrates the effects of the squint angle on the appearance of a SAR image, and thus demonstrates the importance of incorporating information about the squint angle into the baseline classification algorithm.

SAR Imaging Example

Figure 9(a) shows a perspective view of a simple object that is being imaged by an airborne SAR. The object consists of a square grid of point reflectors (shown in blue) in the ground plane, and one additional point reflector (shown in red) above the ground plane and directly over the center of the grid. Figure 9(b) shows a top view of the same imaging configuration. From this top view, we can see that the grid of point reflectors is perfectly aligned with the projected radar line of sight; we arbitrarily define this orientation to correspond to a 0° aspect angle. The object is also being imaged at a 0° squint angle, because the radar is looking in a direction perpendicular to the line of flight.

Figure 9(c) shows the same imaging configuration once again, but from a viewing direction perpendicular to the slant plane. Thus we see the projection of the object onto the slant plane, which (according to our mathematical SAR model) corresponds directly to the result produced by the SAR imaging process.

Because of the projection operation, the grid of point reflectors (in blue) appears foreshortened in the vertical dimension, and the point reflector above the ground (in red) appears just above the grid. Finally, Figure 9(d) shows an image-sized portion of the slant-plane projection (displayed according to the convention that range increases in the downward direction) that represents the SAR image of the object at a 0° aspect angle and a 0° squint angle.

Figure 10(a) shows a perspective view of the same object being imaged with a different SAR geometry. For this example, we assume that the slant plane has been adjusted so that the depression angle matches that of the previous example. From the top view shown in Figure 10(b) we can see that the aspect angle has not changed (i.e., the grid is still aligned with the projected radar line of sight), but that the squint angle has changed from 0° to 45°.

Figure 10(c) shows the object projected onto the slant-plane under this new imaging configuration. The grid of point reflectors (in blue), which appeared as a diamond from the top view, now appears as a foreshortened diamond in the vertical dimension because of the projection operation; the additional point reflector above the ground plane (in red) appears over the upper corner of this diamond because of its height. Figure 10(d) shows an image-sized portion of the slant-plane projection (again displayed according to

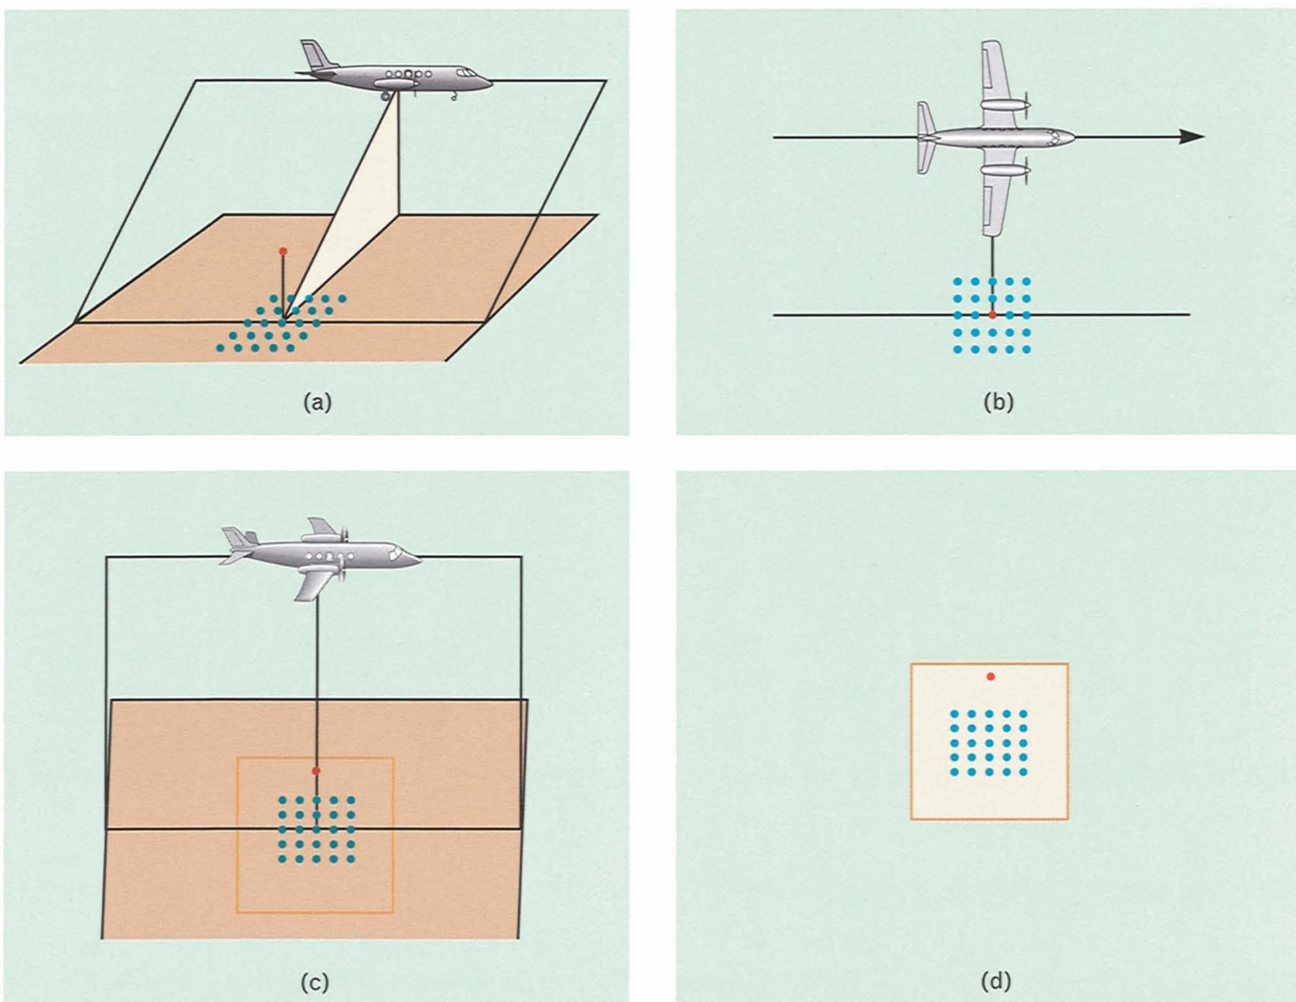

FIGURE 9. Illustration of SAR imaging as a projection (broadside case). The collection of point reflectors being imaged is shown from (a) perspective view, (b) top view, and (c) slant-plane view. (d) The resulting SAR image can be interpreted as an image-sized portion of the slant-plane projection, as indicated by the orange outline in part *c*.

the convention that range increases in the downward direction) that represents the SAR image of the object at a 0° aspect angle and a 45° squint angle.

The primary difference between this SAR image and the one shown in Figure 9(d) is that the point reflectors now appear shifted in the cross-range dimension. The shift for each reflector is not, however, a simple function of the range of the reflector, as it may appear at first glance. Rather, the shift is a function of the 3-D location of the reflector, which is demonstrated by the large shift of the point reflector above the ground plane. This shift has caused the reflector to move out of alignment with the middle column of the grid, which can be seen by comparing Figures 9(d) and 10(d).

Sensitivity of Baseline Classifier to Changes in Squint

The simple geometric examples given in Figures 9 and 10 show that we can produce two different images of an object by using two different squint angles for a fixed set of aspect and depression angles. From a qualitative standpoint, these differences adversely affect the performance of the baseline classifier, because the classifier has only one reference image for each aspect-angle and depression-angle pair. In this section, we describe an experiment that demonstrates quantitatively that the classification statistic used by the baseline classifier—the correlation score—changes significantly as a function of squint angle for a fixed aspect-angle and depression-angle pair.

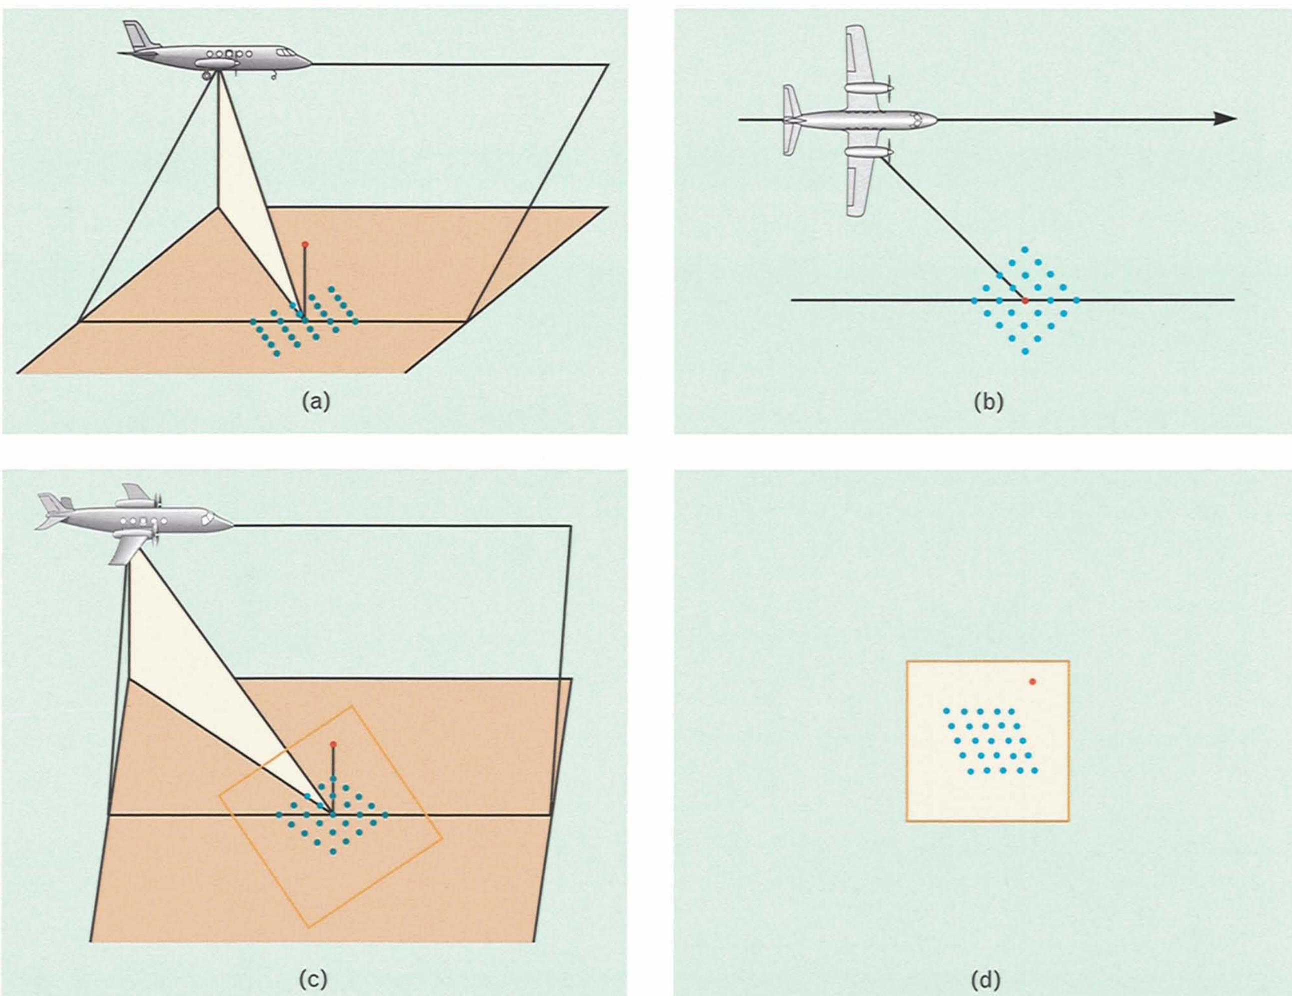

FIGURE 10. Illustration of SAR imaging as a projection (forward-looking case). The sequence of figures—(a) perspective view, (b) top view, (c) slant-plane view, and (d) resulting SAR image—corresponds directly to the sequence shown in Figure 9.

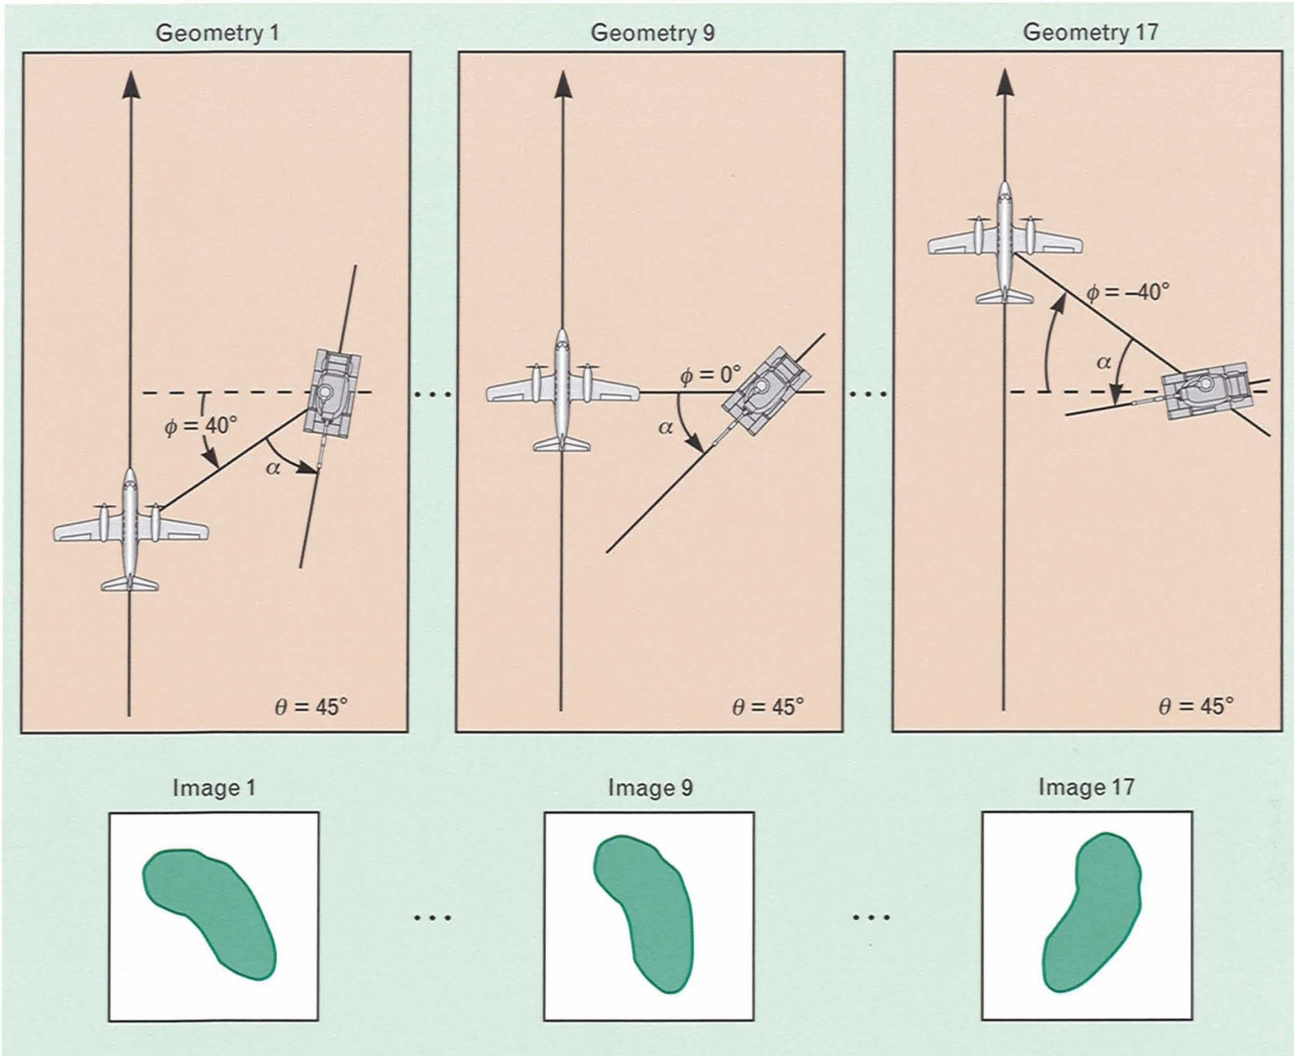

FIGURE 11. Depiction of the imaging configurations used to quantify the robustness of the baseline classifier with respect to variations in squint angle. The test images, which are shown notionally below their respective configuration diagrams, were all formed by using the same aspect angle ($\alpha = 45°$) and depression angle ($\theta = 45°$), but each had a unique squint angle ϕ (a multiple of 5° in the range from –40° to +40°).

We conducted the experiment by using a target-signature simulation package from The Analytic Sciences Corporation known as SARTOOL [2], which was designed to model the dominant electromagnetic characteristics of a target. We used the SARTOOL model of an M48 tank oriented such that both the aspect and depression angles were fixed at 45°. We carried out the experiment by using a single reference image, which was created at a 0° squint angle, and 17 test images, which were created at squint angles ranging from –40° to +40° in 5° increments. Figure 11 illustrates the imaging configurations we used to generate these test images. (Note that the aspect angle and the depression angle are the same for each configuration.) The experiment consisted of correlating the reference image with each of the test images to generate a plot of correlation score versus squint angle. This plot is shown as a solid line in Figure 12.

Of course, the classifier gave perfect performance (i.e., correlation score ρ = 1.0) for the test image formed at a 0° squint angle, because the reference image was formed at this same squint angle. The correlation scores progressively decline, however, as the squint angle for the test set varies in either direction from 0°. At the extremes of –40° and +40°, the correlation scores are approximately 0.5. This result

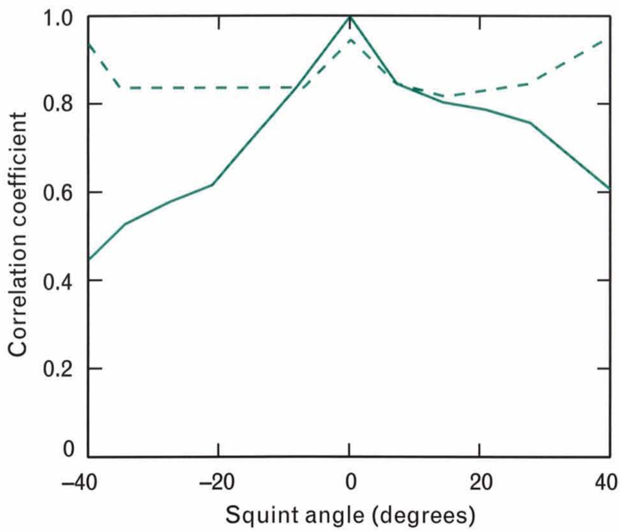

FIGURE 12. Plots of correlation score versus squint angle for the baseline classifier (solid line) and the proposed classifier (dashed line). The test images used for the experiments were those described in Figure 11. The 2-D template used by the baseline classifier and the 3-D template used by the proposed classifier both had aspect and depression angles matching those of the test images.

suggests that the performance of the baseline classifier is sensitive to changes in squint angle. Thus, to improve the baseline classifier we must account for the effects of squint (in addition to the already recognized effects of aspect and depression) on the process of SAR image formation.

How to Develop a Better Classifier

Our analysis in the previous section suggests that we can improve the performance of the baseline classifier by taking into account the effects of both radar viewing direction and radar motion direction on the SAR imaging process. In this section we propose a new classifier that maintains the conventional template-matching engine, but calls for two major modifications to the baseline classifier: (1) the replacement of the present set of 2-D reference images with a set of 3-D templates, and (2) the incorporation of our mathematical model of SAR imaging as a projection so that any 3-D template can be transformed appropriately to synthesize a 2-D reference image for the correlation operation.

We begin by giving a definition of a 3-D template, and we then describe how a 3-D template is transformed into a 2-D reference image. In addition, we present a novel technique for creating 3-D templates from our existing database of 2-D reference images. Finally, in a continuation of the experiment discussed in the previous section, we show that the proposed classifier is much more robust with respect to changes in squint angle than the baseline classifier.

Description of 3-D Templates

A 3-D template is a finely sampled 3-D grid of points representing the volume occupied by the target of interest, in which each grid point corresponds to a scatterer on the target. Each 3-D template is associated with a distinct radar viewing direction, specified by the aspect angle α and the depression angle θ. We thus index the collection of 3-D templates by these radar viewing angles, and we let $T_{\alpha\theta}$ denote the template corresponding to the particular pair (α, θ). The value stored at each point in the template $T_{\alpha\theta}$ represents the radar reflectivity of the scatterer at that point, when the target is illuminated from the direction corresponding to (α, θ). To prepare for the development that follows, we assume that the template $T_{\alpha\theta}$ contains K grid points. The location of the jth point in the 3-D grid is denoted by $\mathbf{p}_j$, and the radar-reflectivity value stored at this point is denoted by A_j, for $j = 1, \ldots, K$.

To transform $T_{\alpha\theta}$ into a 2-D reference image, we use our projection model of the SAR imaging process. Specifically, we project the points in the template $T_{\alpha\theta}$ onto the slant plane defined by the depression angle θ and the squint angle ϕ, to yield a reference image that we denote by $I_{\alpha\theta\phi}$. Let $I_{\alpha\theta\phi}(m, n)$ be the value at the range/cross-range location (m, n) in this reference image. The relation between the values in the template $T_{\alpha\theta}$ and the reference image value $I_{\alpha\theta\phi}(m, n)$ is given by

$$I_{\alpha\theta\phi}(m,n) = \sum_{j \in Q_{\alpha\theta\phi}(m,n)} A_j \,,$$

where $Q_{\alpha\theta\phi}(m, n)$ is the set of indices specified by

$$Q_{\alpha\theta\phi}(m,n) = \left\{ j \;\middle|\; \begin{array}{l} \mathbf{p}_j \text{ projects to location} \\ (m,n) \text{ in the SAR image} \\ \text{corresponding to } (\alpha,\theta,\phi) \end{array} \right\}.$$

Because a SAR image is composed of discrete pixels in the range and cross-range dimensions, the location (m, n) actually corresponds to a locus of points in the slant plane. Let Δ_r and Δ_c be the range and cross-range pixel spacing intervals, respectively, associated with the image. Then any slant plane location (q_r, q_c) such that

$$\begin{aligned} r_m &\le q_r < r_m + \Delta_r \\ c_n &\le q_c < c_n + \Delta_c , \end{aligned}$$

(where r_m and c_n are appropriate constants) is mapped to SAR image location (m, n). Thus $\mathbf{p}_j$ projects to the SAR image location (m, n) if

$$\begin{aligned} r_m &\le \mathbf{p}_j \cdot \hat{\mathbf{r}} < r_m + \Delta_r \\ c_n &\le \mathbf{p}_j \cdot \hat{\mathbf{c}} < c_n + \Delta_c , \end{aligned}$$

where $\hat{\mathbf{r}}$ and $\hat{\mathbf{c}}$ are the unit range and cross-range vectors in the radar coordinate system that was defined earlier.

Description of Proposed Classifier

Figure 13 illustrates the algorithm used by the proposed version of the baseline classifier (note the similarity between this figure and Figure 3). The classifier uses the aspect angle α and the depression angle θ to select a 3-D template from the database. The points in this 3-D template are projected onto the slant plane specified by the squint angle ϕ to produce a 2-D image, which is then correlated with the input test image. If the correlation score exceeds the threshold τ, the target of interest is declared to be present.

Note that the new classifier continues to use the conventional template-matching engine, so that the overall structure of the algorithm is unchanged. The

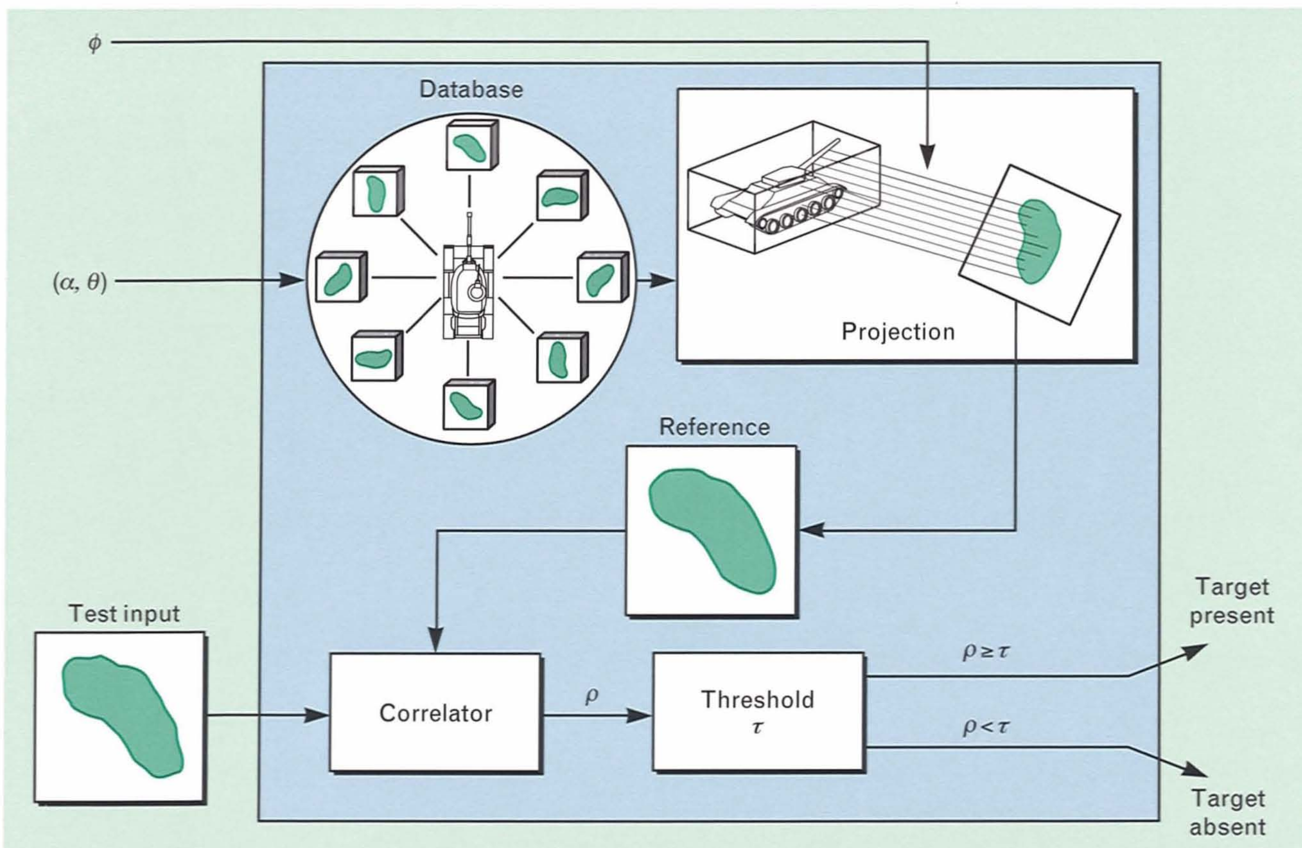

FIGURE 13. Schematic description of the algorithm used by the proposed one-class classifier. The classifier uses the aspect angle α and the depression angle θ to select a 3-D template from the database. This 3-D template is projected onto the slant plane specified by the squint angle ϕ to produce a 2-D image, which is then correlated with the input test image. If the correlation score ρ is greater than or equal to the threshold τ, the *target present* decision is declared.

principal modification to the proposed classifier is that the original database of 2-D reference images has been replaced by a database of 3-D templates. In addition, the proposed classifier is equipped with a processor that transforms a 3-D template into a 2-D image.

Creation of 3-D Templates

Let us now consider how a 3-D template can be created from the existing set of 2-D reference images. For a target of interest, we wish to construct a 3-D template corresponding to a particular aspect-angle and depression-angle pair. To perform this construction we require two or more SAR images of the target, all formed by using the fixed aspect and depression angles of the template, but each formed by using a different squint angle. Recall that each of these SAR images represents a projection of the 3-D distribution of target scatterers onto a 2-D slant plane. Our goal is to use the information contained in these projections to reconstruct the locations and amplitudes of the target scatterers.

We fix the locations of the scatterers $\mathbf{p}_1, \ldots, \mathbf{p}_K$ so that they represent a uniform sampling of a parallelepiped that is approximately the size of the target. Once these scatterer locations are fixed, we then solve for the unknown amplitudes A_j corresponding to the points $\mathbf{p}_j$. Mathematically, we formulate the problem of determining the A_j values in the following way. Let us assume we have L actual SAR images $I_1, \ldots, I_L$ formed at squint angles $\phi_1, \ldots, \phi_L$, respectively. Corresponding to this sequence of actual images, we let $\hat{I}_1, \ldots, \hat{I}_L$ be a sequence of synthetic images formed from the template amplitude values. By using our projection model for the SAR imaging process, the value in the ith synthetic image corresponding to the range/cross-range location (m, n) is computed by

$$\hat{I}_i(m,n) = \sum_{j \in Q_i(m,n)} A_j \,, \tag{4}$$

where $Q_i(m,n)$ is the set of indices specified by

$$Q_i(m,n) = \left\{ j \,\middle|\, \begin{array}{l} \mathbf{p}_j \text{ projects to location } (m,n) \text{ in the} \\ \text{SAR image corresponding to } \phi_i \end{array} \right\}.$$

Let the total mean-square difference between the set of synthetic images and the set of actual images be given by

$$\varepsilon(A_1, \ldots, A_K) = \sum_{i=1}^{L} \left\{ \sum_{m=1}^{M} \sum_{n=1}^{N} \left[\hat{I}_i(m,n) - I_i(m,n) \right]^2 \right\}, \tag{5}$$

where M and N are the range and cross-range dimensions, respectively, of the SAR images. (Note that the dependency of ε on each A_j enters Equation 5 implicitly through Equation 4.)

We can now cast the template construction problem as a multivariable minimization problem. Specifically, we compute the values of $A_1, \ldots, A_K$ by solving

$$\text{Minimize } \varepsilon(A_1, \ldots, A_K)$$
$$\text{such that } A_j \geq 0, \text{ for } j = 1, 2, \ldots, K.$$

To determine an optimal solution, we begin by assigning to each unknown amplitude A_j an initial amplitude that represents our best *a priori* estimate of the actual radar reflectivity at that point. In the absence of *a priori* knowledge, we assign a random initial value to each A_j. Figure 14 illustrates the iterative procedure we use to compute the template amplitude values.

The points $\mathbf{p}_j$ are projected onto each of the L slant planes (each slant plane corresponds to an actual SAR image supplied to the algorithm), which results in a sequence of synthetic images that can be compared to the actual images. The total squared error is computed from these two sets of images (synthetic and actual) and the amplitude values are adjusted such that this total error is reduced. The iteration then cycles through the stages of synthetic image formation, error computation, and amplitude adjustment. The procedure is terminated when the total squared error is less than some prespecified tolerance.

Many standard gradient-descent techniques are available for implementing this iterative minimization; for more details on these techniques see the book by D.G. Luenberger [3].

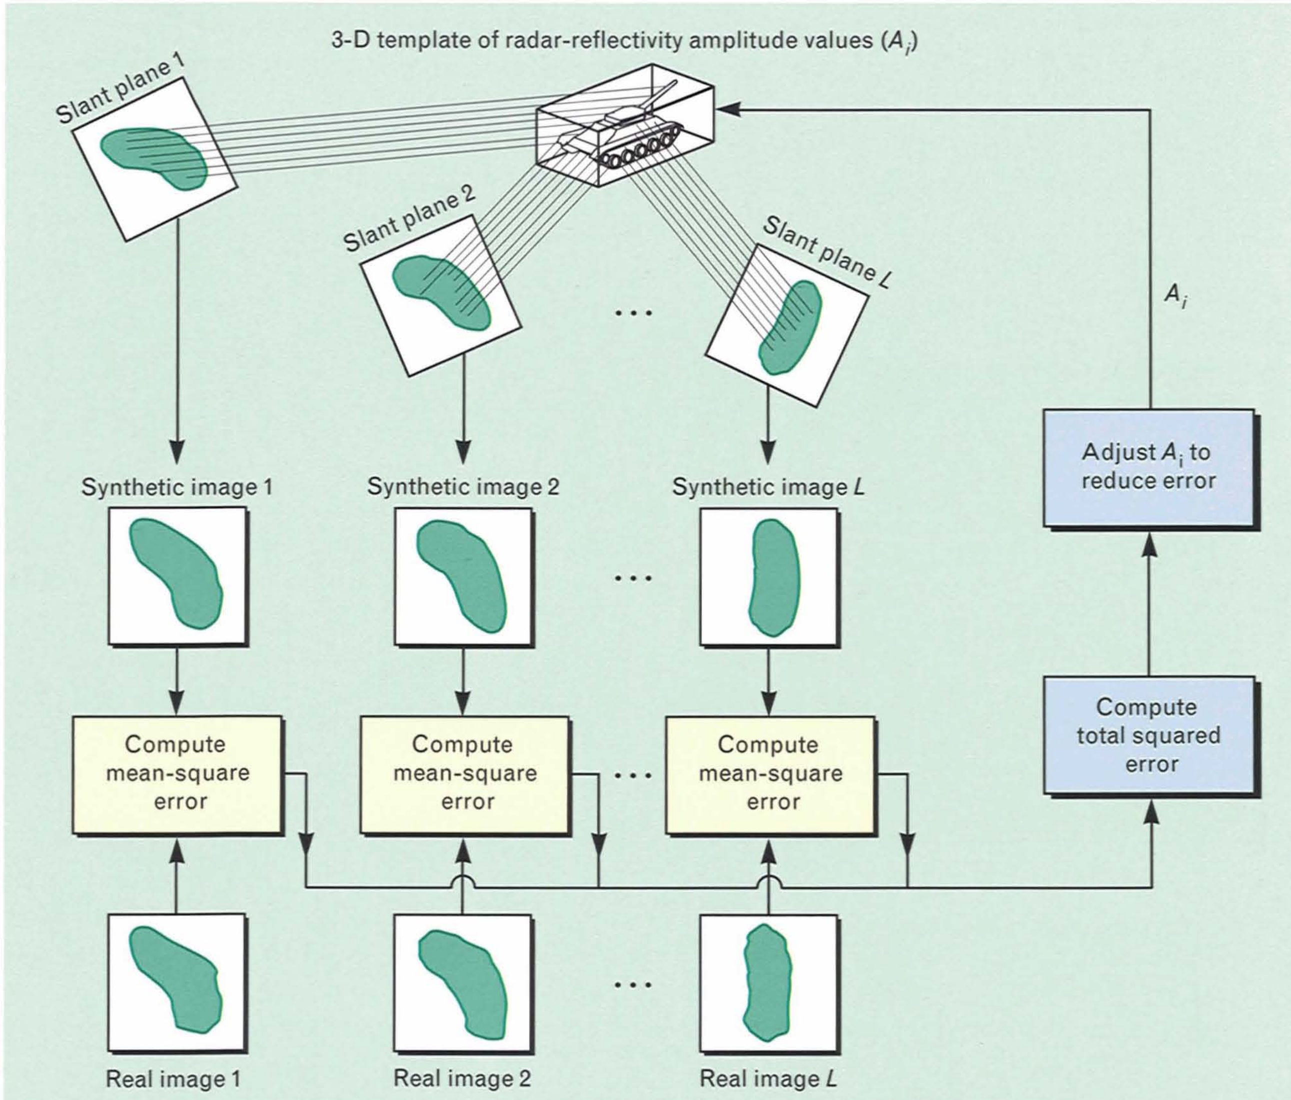

FIGURE 14. Illustration of the 3-D template-creation procedure. The procedure begins with a random assignment of radar-reflectivity values at points in the 3-D template; thereafter the procedure becomes an iterative refinement process. The amplitudes are projected onto L slant planes (each slant plane corresponds to an actual SAR image supplied to the algorithm), which results in a sequence of synthetic images that can be compared to the actual images. The total squared error is computed from these two sets of images (synthetic and actual), and the amplitude values are adjusted such that this total error is reduced. The iteration then cycles through the stages of synthetic image formation, error computation, and amplitude adjustment. The procedure is terminated when the total squared error is less than some prespecified tolerance.

Sensitivity of Proposed Classifier to Changes in Squint

Earlier we described an experiment with the baseline classifier that provided quantitative proof that the correlation score varies significantly as a function of squint angle for a fixed aspect-angle and depression-angle pair. Because the correlation score is the classification statistic used by the baseline classifier, overall performance is extremely sensitive to changes in squint angle. This section summarizes a continuation of the experiment in which we measure the sensitivity of the proposed classifier to changes in squint angle.

For this experiment we used the same set of 17 test images described earlier. Recall that these images were formed by using the SARTOOL model of the M48 tank oriented such that both the aspect angle and depression angle were fixed at 45°. These images were created with squint angles ranging from –40° to +40°

in 5° increments. Let us denote the ith image in this set by I_i, with the corresponding squint angle ϕ_i given by the expression

$$\phi_i = -40 + 5(i - 1), \; i = 1, \ldots, 17.$$

Our 2-D reference image, which was created at a squint angle of 0°, was replaced by a 3-D template corresponding to an aspect angle of 45° and a depression angle of 45°. We constructed this 3-D template by applying the algorithm described in the previous section to three SARTOOL images formed at squint angles of –40°, 0°, and 40°. These three images are displayed in Figures 7(b), 7(c), and 7(d), respectively.

The result of this 3-D template construction is interesting to observe. Recall that the value stored at each point in the template is the radar reflectivity of the scatterer at that location, when the target is imaged from the given viewing direction. Figure 15(a) shows a simple solids model that represents the basic features and dimensions of the M48 tank; Figure 15(b) shows the same model overlaid with the most significant radar-reflectivity values contained in the template. Note that there are significant radar returns from the front right fender, as well as from the turret, the track region, and the front left portion of the tank.

In our previous experiment we correlated the reference image with each of the 17 test images to generate a plot (denoted by the solid line in Figure 12) of correlation score versus squint angle. In this experiment we first transformed the 3-D template into a sequence of reference images, which we denote by $\hat{I}_1, \ldots, \hat{I}_{17}$, corresponding to squint angles $\phi_1, \ldots, \phi_{17}$. We then correlated I_i with $\hat{I}_i$, for $i = 1, \ldots, 17$, to obtain the dashed line in Figure 12.

We observe in this plot that the highest correlation scores occur at the squint angles –40°, 0°, and +40°. This result is not surprising because the template was constructed by using images formed at these three squint angles. The average correlation score obtained by using the 3-D template is approximately 0.85, which is much higher than the average correlation score obtained by using the baseline classifier in the previous experiment. Moreover, the scores do not change significantly as the squint angle varies in either direction from 0°, which suggests that the proposed classifier is more robust with respect to changes in squint angle than the baseline classifier.

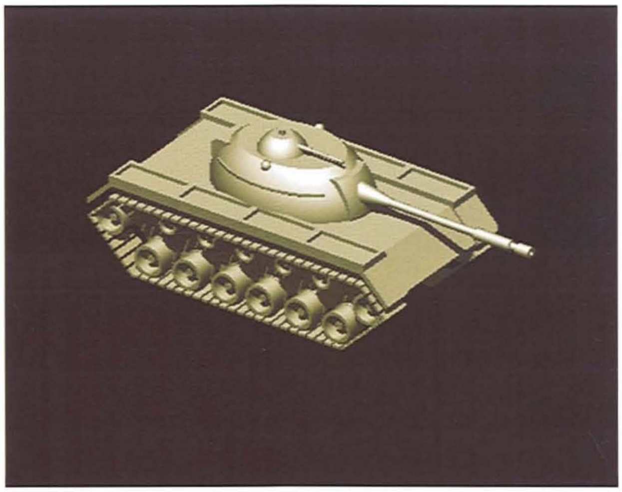

(a)

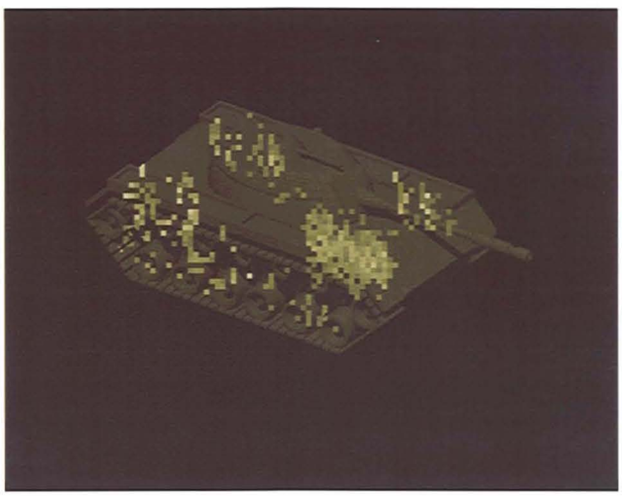

(b)

FIGURE 15. (a) Simple solids model of an M48 tank; (b) the same model of the tank with an overlay of the most significant radar-reflectivity information contained in the 3-D template for the tank. (The overall intensity of the underlying solids model has been reduced to give emphasis to the radar-reflectivity information.) Note that there are significant radar returns from the front right fender, the turret, the track region, and the front left portion of the tank.

Summary

In this article, we have analyzed and suggested improvements to a conventional template-matching classifier currently used in an operational ATR system.

This conventional classifier uses a collection of 2-D SAR reference images to represent a full range of radar viewing directions for a prespecified set of targets. For each target category, the input image is correlated with the reference image that was formed from the most similar radar viewing direction; the input is then classified to the category with the highest correlation score.

Although this algorithm seems reasonable, we found that it produces surprisingly poor classification results for some target types. We explained these poor results by using a simple mathematical model of the SAR imaging process. As our model reveals, radar motion direction is as important as radar viewing direction in specifying SAR imaging geometry. Thus two target images formed with the same radar viewing direction but different radar motion directions can appear quite different. Because the conventional classifier does not explicitly account for radar motion direction, its performance is degraded.

Accordingly, we have proposed and demonstrated an improved version of the conventional template-based classifier that accounts for both direction parameters. In our improved classifier, each 2-D image from the reference library is replaced by a 3-D template so that more target scattering information is available at each viewing direction. As in the conventional classifier, the reference image is selected on the basis of radar viewing direction; by using the mathematical SAR imaging model the improved classifier then transforms the selected 3-D template to a 2-D image whose radar motion direction matches that of the input image.

After comparing the experimental correlation scores between the original 2-D template-based classifier and the improved 3-D template-based classifier, we conclude that the new classifier is significantly more robust with respect to changes in squint angle.

Acknowledgments

The authors gratefully acknowledge the encouragement and support of Leslie M. Novak and Gerald B. Morse. We also thank Michael C. Burl for many technical discussions on the topic of SAR imaging. This research was sponsored by the Advanced Research Projects Agency.

REFERENCES

1. J.C. Henry, "The Lincoln Laboratory 33-GHz Airborne Polarimetric SAR Imaging System," *IEEE National Telesystems Conf., Atlanta, GA, 26–27 March 1991*, p. 353.
2. S. Zabele, S. Bachinsky, B. Myers, A. Stiehl, and R. Pinto, *Signature Prediction User's Manual, version 7.1* (The Analytic Science Corporation, Reading, MA, 1990).
3. D.G. Luenberger, *Linear and Nonlinear Programming* (Addison-Wesley, Reading, MA, 1984).

APPENDIX: MODELING SPOTLIGHT SAR IMAGING AS A PROJECTION

THROUGHOUT THE TEXT of this article we modeled the SAR imaging process as a projection of the 3-D distribution of target scatterers onto a 2-D slant plane. We relied heavily on this projection model as we analyzed problems with the baseline classifier and developed improvements to it. In this appendix we provide justification for using the projection model, and we state the conditions under which this model is valid.

Our strategy for justifying the projection model consists of four main steps. In the first step, we construct the basis vectors for the radar coordinate system and perform the projection operation on a point reflector to obtain approximate expressions for the SAR image location of the reflector. In the second step, we build a foundation for analyzing the projection approximations by writing the exact nonlinear expressions for the physical quantities that are measured by a SAR when imaging the point reflector. In the third step, we expand these nonlinear expressions into first-order Taylor series in the vicinity of the radar aimpoint, and then observe that the resulting linear approximations are identical to the original projection approximations. Finally, in the fourth step, we quantify the accuracy of the projection model by deriving simple bounds on the approximation error.

We begin by obtaining expressions for the projected location of a point reflector in the radar slant plane. In keeping with the notation we established for Figure 8, we define the position of the point reflector in world coordinates as

$$\mathbf{p} = \begin{bmatrix} p_x \\ p_y \\ p_z \end{bmatrix}.$$

In addition, we define the time-dependent sensor position in world coordinates as

$$\mathbf{s}(t) = \begin{bmatrix} s_x(t) \\ s_y(t) \\ s_z(t) \end{bmatrix}.$$

Because the sensor is moving at a fixed altitude parallel to the y-axis in Figure 8, we explicitly remove the time dependency from the first and third coordinates of $\mathbf{s}(t)$ by setting $s_x(t) \equiv s_x$ and $s_z(t) \equiv s_z$.

With the sensor and point-reflector positions defined, we can now construct the basis vectors for the radar coordinate system. Recall from the main text that we originally expressed the basis vectors $\hat{\mathbf{r}}$, $\hat{\mathbf{c}}$, and $\hat{\mathbf{n}}$ in terms of the imaging angles θ and ϕ. In this section, we reconstruct the same basis vectors $\hat{\mathbf{r}}$, $\hat{\mathbf{c}}$, and $\hat{\mathbf{n}}$, but we express them in a form that is more convenient and more useful for our derivations. In particular, rather than using fixed angles from a single imaging geometry, we express these vectors in a time-dependent form in terms of the sensor coordinates.

At time t, the range vector $\hat{\mathbf{r}}(t)$ can be constructed by using the formula

$$\hat{\mathbf{r}}(t) = -\frac{\mathbf{s}(t)}{\|\mathbf{s}(t)\|} = \frac{-1}{\sqrt{s_x^2 + s_y^2(t) + s_z^2}} \begin{bmatrix} s_x \\ s_y(t) \\ s_z \end{bmatrix}.$$

Also, as before, the slant-plane normal vector $\hat{\mathbf{n}}(t)$ can be constructed by using the formula

$$\hat{\mathbf{n}}(t) = \frac{\hat{\mathbf{r}}(t) \times \hat{\mathbf{y}}}{\|\hat{\mathbf{r}}(t) \times \hat{\mathbf{y}}\|} = \frac{1}{\sqrt{s_x^2 + s_z^2}} \begin{bmatrix} s_z \\ 0 \\ -s_x \end{bmatrix}.$$

(Note that this normal vector is constant, because the radar slant plane does not change with time.) Finally, the cross-range vector $\hat{\mathbf{c}}(t)$ is determined by the cross product of the other two vectors, as given by

$$\hat{\mathbf{c}}(t) = \hat{\mathbf{n}}(t) \times \hat{\mathbf{r}}(t)$$

$$= \frac{1}{\sqrt{s_x^2 + s_z^2}\sqrt{s_x^2 + s_y^2(t) + s_z^2}} \begin{bmatrix} -s_x s_y(t) \\ s_x^2 + s_z^2 \\ -s_y(t) s_z \end{bmatrix}.$$

The projection of the point $\mathbf{p}$ onto each of these basis vectors yields the new vector

$$\mathbf{q}(t) = \begin{bmatrix} q_r(t) \\ q_c(t) \\ q_n(t) \end{bmatrix}$$

in radar coordinates. In particular, the range and cross-range coordinates of the point $\mathbf{p}$ are given by

$$q_r(t) = \mathbf{p} \cdot \hat{\mathbf{r}}(t)$$

$$= \frac{-1}{\sqrt{s_x^2 + s_y^2(t) + s_z^2}} \left[p_x s_x + p_y s_y(t) + p_z s_z \right]$$

and

$$q_c(t) = \mathbf{p} \cdot \hat{\mathbf{c}}(t)$$

$$= \frac{1}{\sqrt{s_x^2 + s_z^2}\sqrt{s_x^2 + s_y^2(t) + s_z^2}}$$

$$\times \left[-p_x s_x s_y(t) + p_y (s_x^2 + s_z^2) - p_z s_y(t) s_z \right].$$

Having obtained these projection approximations for the range and cross-range coordinates of the point $\mathbf{p}$, we now seek expressions for the actual quantities measured by a SAR with respect to the location of $\mathbf{p}$. Specifically, these quantities are (1) the *relative range* of the point reflector (i.e., the difference between the distance from the sensor to the point reflector and the distance from the sensor to the aimpoint), and (2) a scaled version of the *relative range rate* of the point reflector (i.e., the rate of change of the relative range with respect to time). We can express the relative range of the point reflector as

$$r[\mathbf{s}(t), \mathbf{p}] = \left\| \mathbf{s}(t) - \mathbf{p} \right\| - \left\| \mathbf{s}(t) \right\|.$$

In addition, by differentiating the above expression with respect to time, we can write the relative range rate $\dot{r}[\mathbf{s}(t), \mathbf{p}]$ of the point reflector as

$$\dot{r}[\mathbf{s}(t), \mathbf{p}] = \frac{\partial}{\partial t} r[\mathbf{s}(t), \mathbf{p}]$$

$$= \dot{s}_y(t) \left[\frac{s_y(t) - p_y}{\left\| \mathbf{s}(t) - \mathbf{p} \right\|} - \frac{s_y(t)}{\left\| \mathbf{s}(t) \right\|} \right].$$

In reality, the relative range rate is rarely used in its raw form as it appears above, because it is so highly dependent on both the speed of the sensor and the distance of the sensor from the aimpoint. Rather, these undesirable dependencies on the absolute sensor velocity and position are usually removed through preprocessing, so that the cross-range dimension of the resulting SAR images is normalized, and SAR images created under different imaging conditions can be directly compared. Thus we introduce two simple time-dependent corrections to $\dot{r}[\mathbf{s}(t), \mathbf{p}]$ (one correction for the absolute sensor velocity, and the other correction for the absolute sensor position) to obtain the *compensated relative range rate*, which we denote by $\dot{r}_c[\mathbf{s}(t), \mathbf{p}]$.

To compute the correction for sensor motion, we begin by decomposing the sensor velocity vector into two velocity components in the slant plane, with one component along the radar line of sight, and the other component orthogonal to the radar line of sight. We then note (under the assumption that the sensor is far from the aimpoint) that the relative range rate is affected only by the component of the sensor velocity vector that is orthogonal to the radar line of sight. (This statement is true because the velocity component along the radar line of sight, considered separately, induces exactly the same range rate on all points in the imaging area, resulting in a relative range rate of zero for these points.) The sensor speed in the direction orthogonal to the radar line of sight can be expressed as

$$\dot{s}_y^{\perp}(t) = \left(-\frac{\sqrt{s_x^2 + s_z^2}}{\left\| \mathbf{s}(t) \right\|} \right) \dot{s}_y(t).$$

Because the sensor speed appears in the numerator of the expression for relative range rate, the above compensation for sensor speed will appear in the denominator of the overall range-rate correction term.

To obtain the correction for sensor position, we note that the denominator of each term in the expression for relative range rate is on the order of $\|\mathbf{s}(t)\|$. Thus a suitable compensation for the distance of the sensor from the aimpoint is simply this factor $\|\mathbf{s}(t)\|$, which will appear in the numerator of the overall range-rate correction term. By applying both of the computed corrections to the original expression for relative range rate, we can express the compensated relative range rate as

$$\begin{aligned}\dot{r}_c[\mathbf{s}(t), \mathbf{p}] &= \left(\frac{\|\mathbf{s}(t)\|}{\dot{s}_y^{\perp}(t)}\right) \dot{r}[\mathbf{s}(t), \mathbf{p}] \\ &= \frac{-\|\mathbf{s}(t)\|^2}{\sqrt{s_x^2 + s_z^2}} \left[\frac{s_y(t) - p_y}{\|\mathbf{s}(t) - \mathbf{p}\|} - \frac{s_y(t)}{\|\mathbf{s}(t)\|}\right].\end{aligned}$$

Having produced explicit expressions for relative range and compensated relative range rate, we rewrite them more suggestively as range and cross-range measurements by using the notation

$$\tilde{q}_r(t) = r[\mathbf{s}(t), \mathbf{p}]$$

and

$$\tilde{q}_c(t) = \dot{r}_c[\mathbf{s}(t), \mathbf{p}].$$

We now show that these actual radar measurements $\tilde{q}_r(t)$ and $\tilde{q}_c(t)$ are well approximated by the previously computed projections $q_r(t)$ and $q_c(t)$, respectively. We begin by separately expanding the expressions for $\tilde{q}_r(t)$ and $\tilde{q}_c(t)$ into Taylor series around the radar aimpoint (i.e., around $\mathbf{p} = \mathbf{0}$), retaining only the first-order terms. For the range component, this procedure yields

$$\begin{aligned}\tilde{q}_r(t) \approx\ & r[\mathbf{s}(t), \mathbf{p}]\Big|_{\mathbf{p}=\mathbf{0}} + p_x \cdot \frac{\partial r[\mathbf{s}(t), \mathbf{p}]}{\partial p_x}\bigg|_{\mathbf{p}=\mathbf{0}} \\ & + p_y \cdot \frac{\partial r[\mathbf{s}(t), \mathbf{p}]}{\partial p_y}\bigg|_{\mathbf{p}=\mathbf{0}} + p_z \cdot \frac{\partial r[\mathbf{s}(t), \mathbf{p}]}{\partial p_z}\bigg|_{\mathbf{p}=\mathbf{0}},\end{aligned}$$

which reduces to

$$\begin{aligned}\tilde{q}_r(t) &\approx 0 + \frac{-1}{\sqrt{s_x^2 + s_y^2(t) + s_z^2}}\left[p_x s_x + p_y s_y(t) + p_z s_z\right] \\ &= \mathbf{p} \cdot \hat{\mathbf{r}}(t).\end{aligned}$$

For the cross-range component, the Taylor series expansion yields

$$\begin{aligned}\tilde{q}_c(t) \approx\ & \dot{r}_c[\mathbf{s}(t), \mathbf{p}]\Big|_{\mathbf{p}=\mathbf{0}} + p_x \cdot \frac{\partial \dot{r}_c[\mathbf{s}(t), \mathbf{p}]}{\partial p_x}\bigg|_{\mathbf{p}=\mathbf{0}} \\ & + p_y \cdot \frac{\partial \dot{r}_c[\mathbf{s}(t), \mathbf{p}]}{\partial p_y}\bigg|_{\mathbf{p}=\mathbf{0}} + p_z \cdot \frac{\partial \dot{r}_c[\mathbf{s}(t), \mathbf{p}]}{\partial p_z}\bigg|_{\mathbf{p}=\mathbf{0}},\end{aligned}$$

which reduces to

$$\begin{aligned}\tilde{q}_c(t) &\approx 0 + \frac{1}{\sqrt{s_x^2 + s_z^2}\sqrt{s_x^2 + s_y^2(t) + s_z^2}} \\ &\quad \times \left[-p_x s_x s_y(t) + p_y(s_x^2 + s_z^2) - p_z s_z s_y(t)\right] \\ &= \mathbf{p} \cdot \hat{\mathbf{c}}(t).\end{aligned}$$

Note that these linear Taylor series expansions are identical to the original projections $q_r(t)$ and $q_c(t)$, indicating that $q_r(t)$ and $q_c(t)$ are good approximations to $\tilde{q}_r(t)$ and $\tilde{q}_c(t)$ in the vicinity of the radar aimpoint.

Error Analysis

Let us now quantify the error incurred by using these linear approximations instead of the actual measurements. To keep the analysis concise, we examine only the error in the relative range approximation. To simplify notation, we arbitrarily choose a time t_0 and remove the explicit time dependency of variables by setting $\mathbf{s} \equiv \mathbf{s}(t_0)$, $\tilde{q}_r \equiv \tilde{q}_r(t_0)$, and $q_r \equiv q_r(t_0)$. In addition, for convenience we express the length of $\mathbf{p}$ as a fraction δ of the length of $\mathbf{s}$ (i.e., $\|\mathbf{p}\| = \delta\|\mathbf{s}\|$), and we define β to be the angle between $\mathbf{p}$ and $\mathbf{s}$.

With these definitions, we can rewrite the expression for the true relative range measurement $\tilde{q}_r$ by using the law of cosines to yield

$$\begin{aligned}\tilde{q}_r &= \sqrt{\|\mathbf{s}\|^2 + \|\mathbf{p}\|^2 - 2\|\mathbf{s}\|\|\mathbf{p}\|\cos\beta} - \|\mathbf{s}\| \\ &= \sqrt{\|\mathbf{s}\|^2 + \delta^2\|\mathbf{s}\|^2 - 2\delta\|\mathbf{s}\|^2\cos\beta} - \|\mathbf{s}\| \\ &= \|\mathbf{s}\|\left(\sqrt{1 + \delta^2 - 2\delta\cos\beta} - 1\right).\end{aligned}$$

Also, we can rewrite the expression for the relative range approximation q_r by using the identity

$$\mathbf{p}\cdot\mathbf{s} = \|\mathbf{p}\|\|\mathbf{s}\|\cos\beta$$

to yield

$$\begin{aligned}q_r &= -\frac{\|\mathbf{p}\|\|\mathbf{s}\|\cos\beta}{\|\mathbf{s}\|} \\ &= -\delta\|\mathbf{s}\|\cos\beta.\end{aligned}$$

Because q_r always underestimates $\tilde{q}_r$, we can write the absolute approximation error e simply as

$$\begin{aligned}e &= \tilde{q}_r - q_r \\ &= \|\mathbf{s}\|\left(\sqrt{1 + \delta^2 - 2\delta\cos\beta} - 1 + \delta\cos\beta\right).\end{aligned}$$

For a fixed value of δ, the error reaches its maximum value at the angle $\beta = \cos^{-1}(\delta/2)$. Substituting this angle back into the formula for e yields the upper bound

$$e_{\max} = \frac{\delta^2}{2}\|\mathbf{s}\| = \frac{\delta}{2}\|\mathbf{p}\|.$$

For the specific case of data collection with the Lincoln Laboratory millimeter-wave sensor, δ is typically no larger than 0.05, so the error incurred by using the projection approximation for a given point $\mathbf{p}$ is no more than 2.5% of the distance of $\mathbf{p}$ from the aimpoint.

The derivation of the error bound for the cross-range approximation is similar in spirit to the derivation given above, but it is much more lengthy and tedious, and hence is omitted.

SHAWN M. VERBOUT
is an assistant staff member in the Surveillance Systems group. His research interests are in the development of algorithms for signal detection, estimation, and classification. He received a B.S. degree in mathematics from Illinois State University, and he is currently pursuing a master's degree in electrical engineering at MIT as a member of the Lincoln Laboratory Staff Associate Program. Shawn has been at Lincoln Laboratory since 1987.

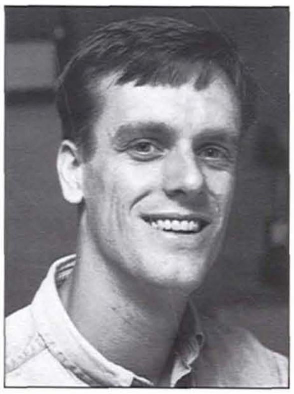

WILLIAM W. IRVING
is an associate staff member in the Surveillance Systems group. His research interests are in the areas of decentralized Bayesian detection theory and multiresolution stochastic modeling. He received S.B., S.M., and E.E. degrees in electrical engineering from MIT, where he is currently pursuing a Ph.D. degree as a member of the Lincoln Laboratory Staff Associate Program. Bill has been at Lincoln Laboratory since 1987.

AMANDA S. HANES
is an assistant staff member in the Surveillance Systems group. Her research interests are in the detection, discrimination, and classification of stationary targets. She received a B.S. degree in mathematics from Tufts University, and she also worked at the Johns Hopkins University Applied Physics Laboratory through the Professional Summer Student Program. Amanda has been at Lincoln Laboratory since 1989.