
Non-Cooperative Identification of Ships with Electrooptical Data

Jerrold Baum, Eleanor Tung, and Steven Rak

■ A multidimensional sensor suite consisting of a laser radar and a passive-infrared (IR) sensor has been proposed and evaluated for detecting and identifying ships at long ranges from an airborne platform. The passive-IR sensor would detect targets by taking advantage of the high target-to-background contrast and the sensor's ability to track over a wide field of view. The higher information content of the ship's range profile would then be exploited for target identification.

A two-part approach was developed to evaluate the proposed system. The first part concentrated on creating synthetic signatures of naval vessels under a variety of controlled sensor operating characteristics and target scenarios. Two neural networks were used to classify the synthetic signatures. In the second part of the evaluation, active and passive-IR measurements of two naval vessels were taken with an existing airborne multidimensional sensor system. These missions demonstrated that measured range profiles were similar to the synthetic profiles. Our results, from both synthetic and measured data, indicate that range-profile and passive-IR signatures complement each other in covering all viewing aspects for long-range ship classification.

ACTIVE AND passive-infrared (IR) sensors combined with neural network classifiers have been proposed to detect and identify ships at long ranges (approximately 100 km) from an airborne platform (Figure 1). To measure the ship features, S. Marcus has designed a noncooperative target identification (NCTI) sensor suite [1, 2]. The multidimensional sensor suite consists of a nonimaging laser radar range profiler and an imaging passive-IR detector array. The passive-IR sensor detects targets by taking advantage of high target-to-background contrast and by utilizing its ability to track over a wide field of view. The higher information content of the active range profile is then exploited for target identification.

To evaluate the proposed system, we developed a two-part approach. In the first part, we created synthetic signatures of naval vessels because of the expense and difficulty of obtaining measured data from targets under a variety of controlled conditions and

scenarios. We created the synthetic signatures by modifying ship models in software to include the operating characteristics of the proposed active and passive-IR sensors. The signatures could then be evaluated under a variety of controlled sensor operating characteristics and target scenarios for automatic ship identification. To classify the synthetic active and passive-IR signatures, two neural network processors were studied: a multilayered neural network (MNN) and the Adaptive Clustering Network (ACN) classifier. The MNN is a feedforward multilayered network that constructs an internal representation by learning from training data. Once training is completed, the learned internal representation enables the classification of unknown input patterns. The ACN is an unsupervised classifier that determines its own categories based on an association threshold and learning rate.

Synthetic data are not, of course, a substitute for measured data. Thus the second part of the system




FIGURE 1. Ship detection/classification scenario for the noncooperative target identification (NCTI) sensor suite. The sensor suite contains an imaging passive-infrared (IR) detector array that scans across a $12^\circ \times 12^\circ$ field of regard (FOR) and a nonimaging laser radar range profiler that scans across a smaller $2^\circ \times 3^\circ$ FOR. The horizon is roughly 250 km from the airplane.

evaluation focused on collecting data from two ships of different naval classes. An airborne laser radar system took multidimensional measurements—passive IR and active range and intensity—of the naval vessels. A major purpose of the missions was to determine whether the range profiles of the two ships at similar ranges and viewing aspects were distinct enough to be used for identification. An equally important goal was to demonstrate that the measured range profiles were similar to the synthetic ship signatures. Such a similarity would validate the neural network classification results, which were based on synthetic imagery.

Theory

We generated synthetic laser radar range signatures and passive-IR images of naval vessels to evaluate the potential performance of the NCTI sensor design for each sensor domain. The active target signatures were generated for a nonimaging laser radar range sensor, which accumulated range-extent information as it scanned across its field of regard (FOR). Passive-IR target signatures were generated in a coarse imaging mode over a larger FOR. This section summarizes the sensor design, then discusses the statistics and as-

sumptions that are contained in the models for the two sensors.

NCTI Sensor Suite

The NCTI sensor suite was designed to detect and identify ships and aircraft at long ranges (Figure 1). In a coarse imaging mode, the passive-IR sensor detects man-made targets against the natural clutter of the ambient sky or sea within a $12^\circ \times 12^\circ$ FOR. The sensor exploits the assumed positive target-to-background contrast of these scenarios. Then, in a nonimaging mode, the pulsed laser radar sensor provides highly resolved range-extent information within a smaller $2^\circ \times 3^\circ$ FOR. The laser radar measures the laser pulse intensity reflected from an extended target. Sufficient detector dwell time is allowed for the transmitted pulse to traverse the target's extent and obtain a range profile of the object.

The sensor suite, described in detail by Marcus in References 1 and 2, was designed to fit in a modified Low Altitude Navigation Targeting Infrared for Night (LANTIRN) pod for mounting on a surveillance aircraft. This design limited the size of the system, notably restricting both the active and passive-IR sensors to share a common 15-cm aperture.

Because of the operational requirement that the range of the sensor suite be 100 km, Marcus selected a CO_2 laser for the active sensor. This choice limited the sensor's instantaneous field of view (IFOV)—i.e., the FOV of a single detector element—to $150 \mu\text{rad}$ for diffraction-limited optics. For the particular CO_2 laser chosen, the theoretical performance was calculated as a carrier-to-noise ratio (CNR) of 32 dB at a range of 100 km for naval targets. This CNR value was sufficient to achieve 1-m range precision with a 30-nsec pulsewidth. A simple 2×2 array was chosen for the active system detector to be scanned through the sensor's $2^\circ \times 3^\circ$ FOR, as shown in Figure 1.

For the passive-IR sensor, Marcus specified a 15-cm-aperture, $f/2$ charge-coupled device (CCD) camera operating in the 3-to-5- μm waveband. The passive-IR sensor was designed to achieve an NE ΔT of 0.1 K $^\circ$ and to operate at a signal-to-noise ratio (SNR) of 32 dB at 100-km range. (Note: NE ΔT , the noise-equivalent delta temperature, is defined as the thermal contrast that produces an SNR of 1.)

Model for Laser Radar Range Profiles

The range profile model is derived from on-going modeling work for laser radar range signatures by J.H. Shapiro [3, 4] and T.J. Green [5]. The model extends earlier laser radar imaging theory for peak-pulse logic and first-pulse logic to account for range-spread targets. In a peak-pulse system, the sensor measures the time between the transmitted laser pulse and the peak of the pulse's reflected return. This time provides the range to the largest or most highly reflective feature within the sensor FOV. In a first-pulse system, the sensor measures the time delay between the transmitted pulse and the first pulse of the reflected return that exceeds a prescribed threshold. Neither type of system, though, uses the entire reflected laser radar pulse, which contains detailed information about the target from which the pulse was reflected, specifically, the target range extent within the sensor IFOV. It has been proposed that a unique target representation, suitable for identification, would be obtained if this range-extent signature were measured.

The sensor model formulated by Shapiro and Green [6, 7] describes the statistics of a coherent laser radar receiver detecting a single reflected laser pulse. The model assumes that the sensor IFOV resolves a range-spread speckle target that has uniform, diffuse reflectivity. This assumption is a reasonable estimate of the target characteristics in the absence of detailed reflectivity information. The model accounts for the effects of the transmitter pulse width and the matched intermediate frequency (IF) filter pulse width, as well as for laser radar speckle and local-oscillator shot-noise statistics. The coherent detection receiver is modeled as a continuous wave (CW) local oscillator, a bandpass IF filter, and a linear envelope video detector. A brief mathematical description of the model is contained in "Appendix 1: Synthetic Laser Radar Range Profiles."

Using the model, we can calculate a synthetic profile for each element of the 2×2 detector. Each element of the detector records an independent range profile for the portion of the target that occupies the element's IFOV. Each range profile is unique with respect to the aspect and range of a particular target. For the final target signature, we sum all

the range-registered IFOV profiles that the sensor has scanned. This sum will also be aspect dependent.

Model for Passive-IR Signatures

Passive-IR imagery is a measure of the thermal radiation from a target and background within a particular waveband of detector sensitivity. For the proposed passive-IR sensor, we selected the 3-to-5- μm waveband after considering the target characteristics, atmospheric transmission, and sensor technology [1, 2].

At the shallow depression angle of the proposed sensor, a naval vessel would be observed against the ocean, with no sky background in the image. From this study, we modeled the scenario as a warm isothermal target on an ambient isothermal background without using highly detailed thermal models. The synthetic passive-IR image was constructed with a simple image template that consisted of a target silhouette, at the desired observation aspect, against a featureless background plane, representing the ambient ocean. For each sensor pixel in the template, a pair of passive-IR intensities was calculated, one intensity value for the target and the other for the background. The appropriate target or background value was then inserted into the final synthetic passive-IR image by using the image template as a guide. We modeled the passive-IR signatures of the target and the background by using two overlapping Gaussian distributions. (For further details of the method used, see "Appendix 2: Synthetic Passive-Infrared Imagery" and Reference 8.) The mean and standard deviation of each distribution were selected from experimental data [2]. In the estimation of the thermal distributions, the sensor NE Δ T was taken into account.

Ship Models and Database of Synthetic Signatures

The four ship models used in this research were of actual naval vessels: the USS *Duncan* (FFG 10), the USS *Horne* (CG 30), the USS *Spruance* (DD 963), and the USS *Texas* (CGN 39). Based on Navy drawings, photographs, and physical models of the ships, the four synthetic models were constructed as a collection of elementary scattering shapes by Georgia Tech Research Institute (GTRI) using the MAX geometric database editor [9]. These solid representations

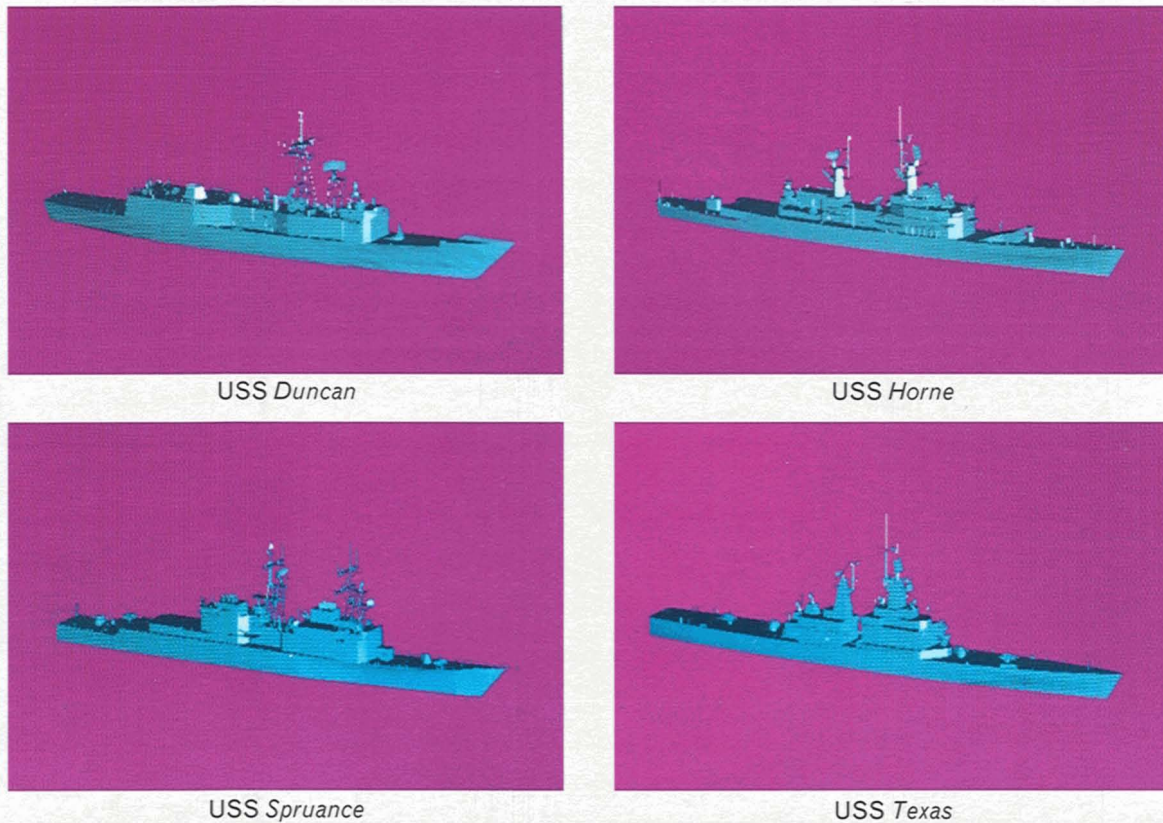


FIGURE 2. Georgia Tech Research Institute (GTRI) ship models visualized in the Army's Ballistic Research Laboratory computer-aided design (BRL-CAD) solid modeling environment: the USS *Duncan* (FFG 10), the USS *Horne* (CG 30), the USS *Spruance* (DD 963), and the USS *Texas* (CGN 39). These four models were used to generate the synthetic laser radar range signatures and passive-IR imagery used in our research.

were converted to facets for portability to other modeling environments.

The four models, each containing more than 15,000 facets, were then imported into the Army's Ballistic Research Laboratory (BRL) computer-aided design (CAD) solid modeling environment [10]. In BRL-CAD, the entire structure of each ship was represented as a single solid continuous skin. BRL-CAD rendered the solid models with high resolution at several viewing aspects and sensor depression angles. Figure 2 contains four examples of the renderings.

In the final step, a ray-tracing algorithm was used to convert each visual model to a range image. The ray-tracing algorithm sampled the model surfaces represented by each pixel in the output image, and assigned to a pixel the range nearest the observer—a

process not unlike the operation of a first-pulse laser radar ranging system [11]. Ship models were generated at 31 viewing aspects, from bow to port, in 3° increments. Each of the models was rendered and ray-traced with 0.25-m resolution on target to form the finely resolved synthetic imagery that would be used to calculate the active and passive-IR signatures. This process ensured that small-scale details of the models would be retained in the calculation of the target signatures.

Signatures were then generated for the ship-identification sensor suite in two domains: active and passive IR. The active signatures were generated for a nonimaging laser radar range sensor. The passive-IR signatures were generated in a coarse imaging mode. The following subsections describe the synthetic signatures that were calculated for the two sensor do-

mains, and the resulting databases.

Laser Radar Range Signatures

For the laser radar range sensor, we calculated a database of ship signatures based on a 100-km operational range and 32-dB CNR, yielding 1-m range resolution [1, 2]. The complete database consisted of synthetic signatures from the four ships at 31 different aspects, from bow to port, in 3° increments.

With the high-resolution-rendered BRL-CAD models, the range-signature database could be calculated with the equations in Appendix 1. The process called for subsampling a highly resolved model to obtain a range-extent histogram for the portion of the target resolved by the detector IFOV. This histogram was convolved with the transmitted laser pulse and the IF filter to form an IFOV profile. Range profiles were calculated for each instance of the detector's IFOV within the sensor's FOR and then accumulated to form a single target and aspect-dependent range signature, or *frame*. To understand the effects of laser speckle on signature characterization, we used Equation A of Appendix 1 to generate signatures subject to speckle statistics. We then used Equation C of Appendix 1 to compute the expected value of the signatures, averaged over speckle statistics. (Note: The expected value of a signature is the trace that would result if an infinite number of sample signatures, subject to the random effects of speckle, were averaged together.)

Realizing that the nonimaging sensor model measures range-extent information, we expected that the naval vessels measured near bow aspect (head on) would have significant range extent, while vessels measured near beam aspect (broadside) would have little range extent. Figure 3 demonstrates this expected aspect dependence for noiseless range signatures of the *Duncan*.

It was important to determine whether the signatures of the different ships were distinguishable. Figure 4 displays noiseless range signatures of the four ships at a 30° viewing aspect. Note the unique characteristics that should make classification possible. At first glance, these noiseless signatures do appear readily distinguishable.

The examples displayed in Figures 3 and 4 represent the expected values of the ship range signatures.




FIGURE 3. Expected-value range signatures of the *Duncan* at different aspect angles from bow to beam: (a) 0° , (b) 30° , (c) 60° , and (d) 90° . The signatures are for a 30-nsec pulse. (Note: A closed-form solution was used to calculate the signatures.)

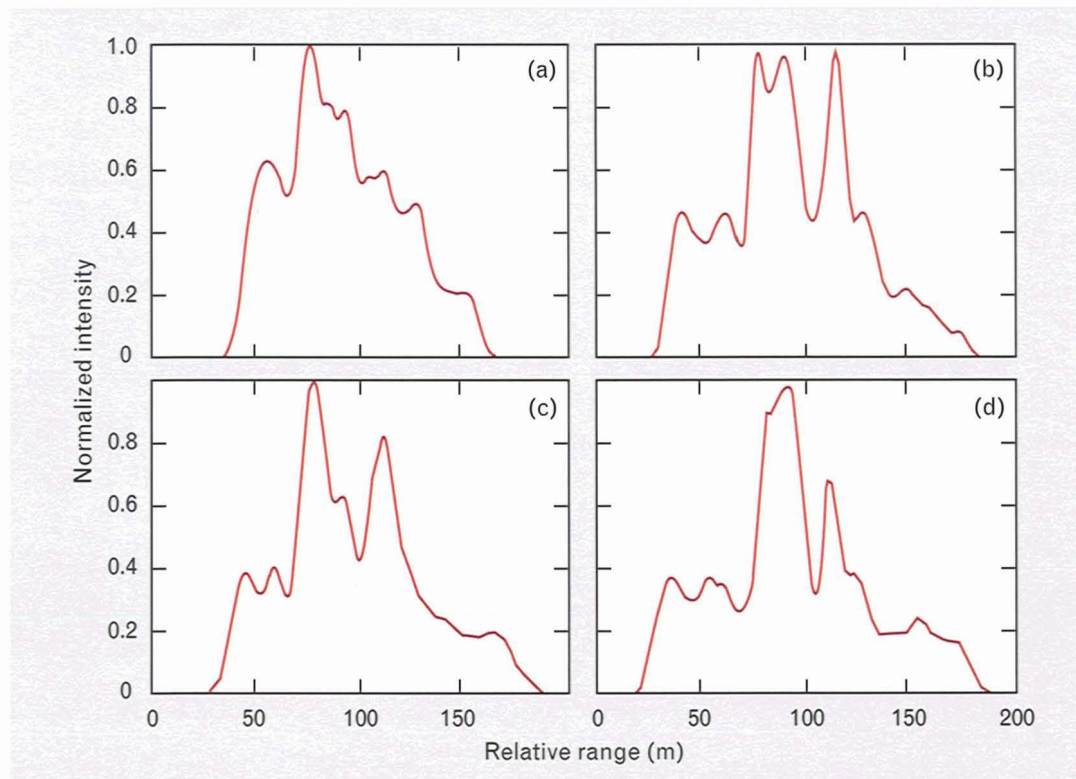


FIGURE 4. Expected-value range signatures of four ships at a 30° aspect angle: (a) the *Duncan*, (b) the *Horne*, (c) the *Spruance*, and (d) the *Texas*. The signatures are for a 30-nsec pulse. (Note: A closed-form solution was used to calculate the signatures.)

We were able to calculate these signatures directly because of the availability of a closed-form solution (Equation C in Appendix 1). In the proposed sensor, the measured signatures would be obtained by averaging numerous frames together to reduce the effects of laser speckle. This averaging is necessary when speckle is present because speckle severely limits the SNR of a laser radar system. The results of cumulative averaging, which increases the SNR of the signatures, are shown in Figure 5. Note that the effects of laser speckle dominate the range signatures when fewer than eight frames are averaged. In such cases, the intensity fluctuations caused by laser speckle are very significant. Indeed, such fluctuations will have a tremendous impact on the classification of naval vessels unless an adequate number of frames are averaged.

Last, we modeled the laser pulse at widths of 10, 30 and 50 nsec with the corresponding matched IF filter widths. Increasing the laser pulse width smoothed the contributions of the various components and superstructure of the ships; decreasing the width preserved

the characteristic details for identification of ship range signatures.

Passive-IR Signatures

Initially, the NCTI sensor suite was designed so that a target would be detected by a passive-IR sensor and identified by a laser radar. Early in the program, however, we noted the potential to develop ship silhouette recognition with the passive-IR imagery. To understand the difficulties associated with long-range passive-IR silhouette recognition, we synthesized passive-IR imagery at ranges of 25, 50, and 100 km. Each image had the appropriate sensor SNR based on the sensor design in References 1 and 2, and was calculated as a function of range and atmospheric extinction coefficient (a measure of how strongly the atmosphere will attenuate the signal). The conditions modeled were based on a summer scenario with the sensor at an altitude of 22,000 ft and a slant range of 100 km, yielding 18.1-dB atmospheric attenuation. Because a high frame rate was technically




FIGURE 5. The reduction of laser speckle by averaging (a) 1, (b) 4, (c) 8, and (d) 16 range signatures. For comparison, the expected value of the range signature is shown with a dashed line. The signatures are for the *Duncan* at a viewing aspect of 30°.

feasible with the sensor design, the recommendation was that all detection and recognition of ship silhouettes be studied with 16-frame-average imagery. The 16-frame average would improve the recognition performance because such averaging would enhance the effective sensor SNR by a factor of 4. Table 1 summarizes the sensor parameters. The complete database consisted of ship silhouettes generated at 31 different aspects, from bow to port, in 3° increments for each of the four ships.

The passive-IR imagery was synthesized from the high-resolution-rendered BRL-CAD ship models. Basically, each ship model was scaled to the appropriate range and converted to a binary silhouette template. Using the method described in Appendix 2, we synthesized the passive-IR imagery (Figure 6) with sensor- and range-dependent thermal distributions for the ship and background.

Synthesizing the passive-IR imagery with the proposed sensor angular resolution of 150 μ rad, we soon discovered that the identification of the ship silhou-

ettes would be difficult at long ranges because of the limited number of pixels on target. At a range of 100 km, for example, a ship with length 172 m and with 17 m of deck above the waterline would present only one row of 15 m \times 15 m square pixels to the sensor. Because of the limited number of pixels on target, the calculation of *super-resolution* (SR) sensor pixels was proposed to enhance the recognition process. In the SR technique [12], subsequent frames of imagery are shifted by fractional pixels, thus allowing the sen-

Table 1. Passive-IR Sensor Parameters

Range (km)	Attenuation (dB)	16-Frame SNR	$\Delta T_{\text{Contrast}}$ (K°)
25	4.5	164.0	1.28
50	9.05	58.0	0.76
100	18.1	7.2	0.27

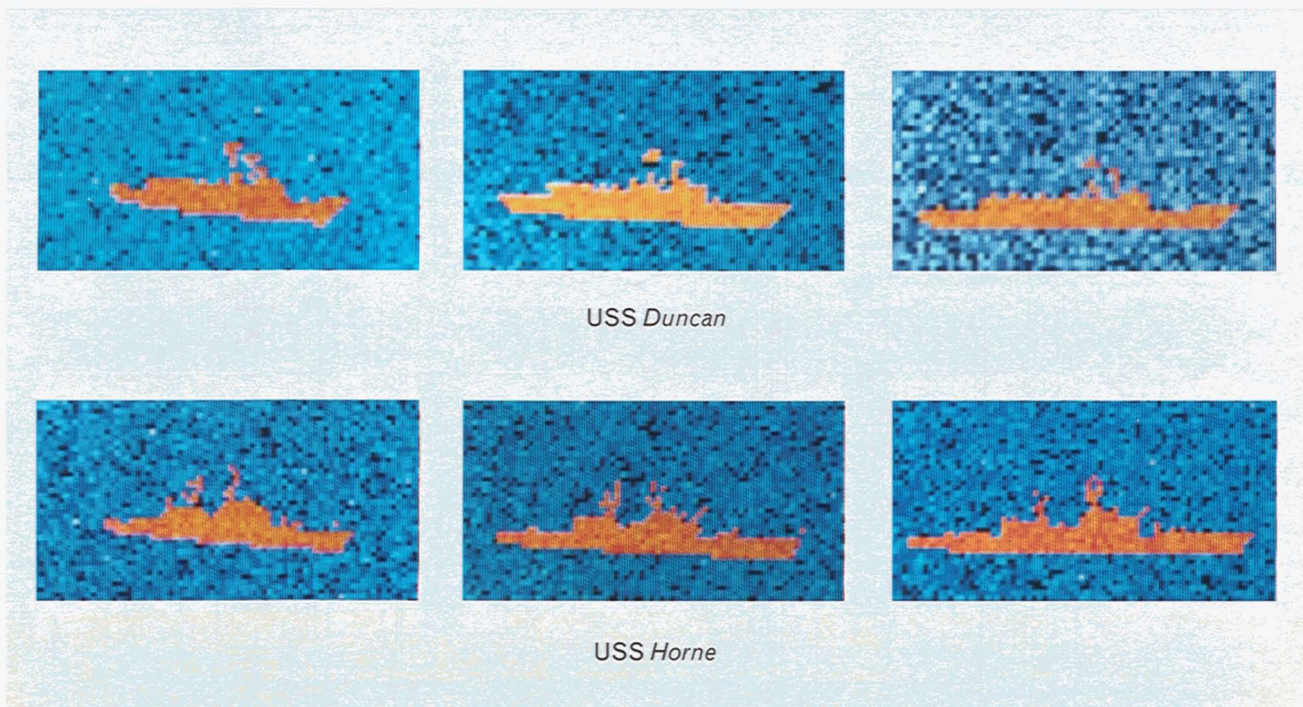


FIGURE 6. Synthetic passive-IR imagery of two ships at different aspect angles. The left images are for 50°, the middle images are for 70°, and the right images are for 90°. The ships are at a range of 25 km.

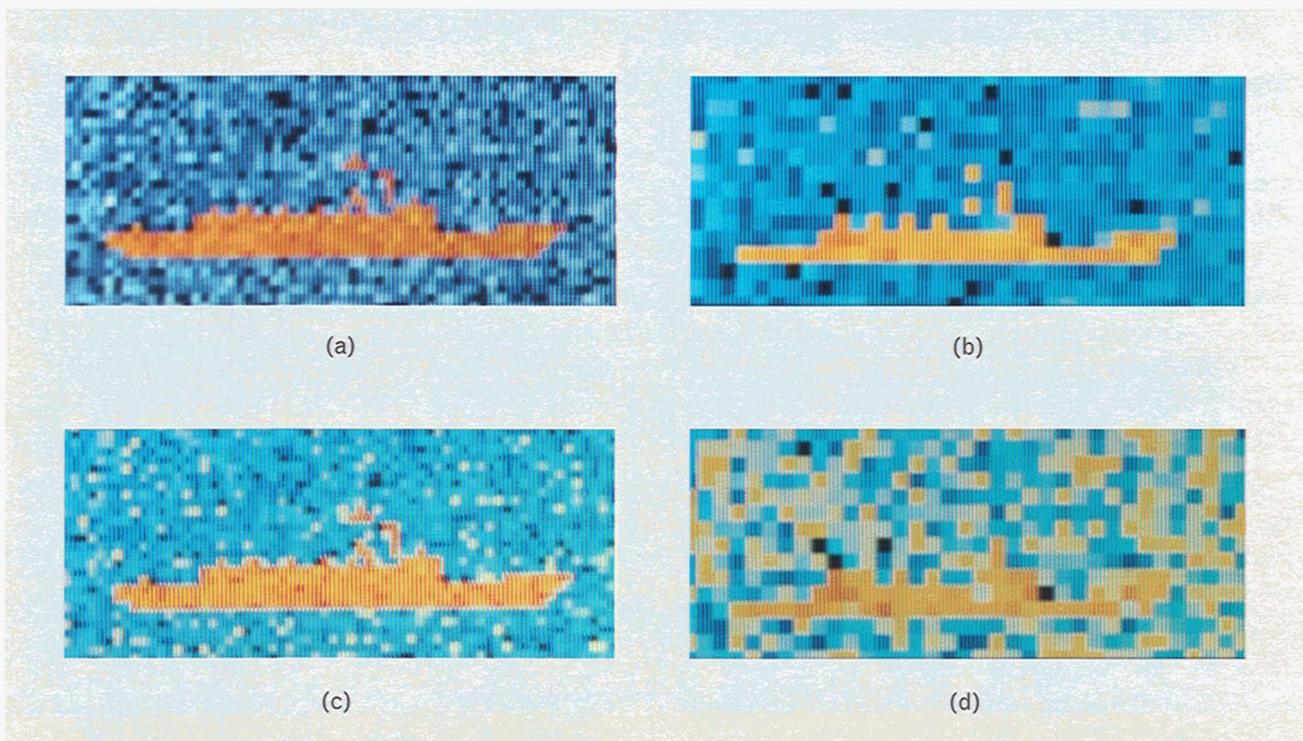


FIGURE 7. The effects of super-resolution (SR) on synthetic passive-IR imagery: (a) image of a ship at a range of 25 km, (b) image of the same ship at 50 km, (c) image of the ship at 50 km after the application of SR, and (d) image of the ship at 100 km after the application of SR. Note that the 50-km SR image has twice the detail of the standard 50-km image.

sor, with additional computer processing, to obtain more detailed information about the passive-IR target signature. This pixel shifting is possible because the pointing accuracy of the sensor is less than the IFOV. A single sensor pixel can be subdivided algebraically by isolating its sub-pixel components in each of the shifted overlapping images. For example, we can super-resolve a series of $150\text{-}\mu\text{rad}$ images at 100 km by shifting the images by $75\text{ }\mu\text{rad}$ in azimuth and elevation. The resulting image will have an angular resolution of $75\text{ }\mu\text{rad}$ at 100 km, and the image will look as though the $150\text{-}\mu\text{rad}$ sensor had been operated at a range of 50 km.

For the purposes of this initial study, we assumed perfect SR; that is, the 100-km $2\times$ SR case was modeled in the following way: sensor statistics that were appropriate for a 100-km image were applied to an image that had been modeled at 50 km. A sample of this imagery is shown in Figure 7. Although the assumption of perfect SR may be optimistic, it provided an initial data point to study classifier performance with realistic sensor parameters and target-to-background contrast ratios.

The database of passive-IR imagery was further processed for feature extraction to minimize the classifier memory storage requirements. By selecting features that were invariant under translation, in-plane rotation, and scale changes, we could use a single feature pattern stored in the classifier memory to recognize targets at a variety of ranges. The invariant representation we chose uses the log-radius polar-angle mapping (LPM) function on segmented targets. It has been shown [13] that LPM features are effective in clustering targets viewed from many different aspects into a few categories. Thus the LPM transform is an ideal choice for 3-D target recognition. Figure 8 shows the effects of target segmentation and invariance transformation for all four ships at two different views.

Data Collection

While the synthetic ship signatures were being generated, data were also being collected from two naval vessels deployed at sea. As discussed earlier, these measurement exercises had two purposes: (1) to determine whether range profiles of two different ships

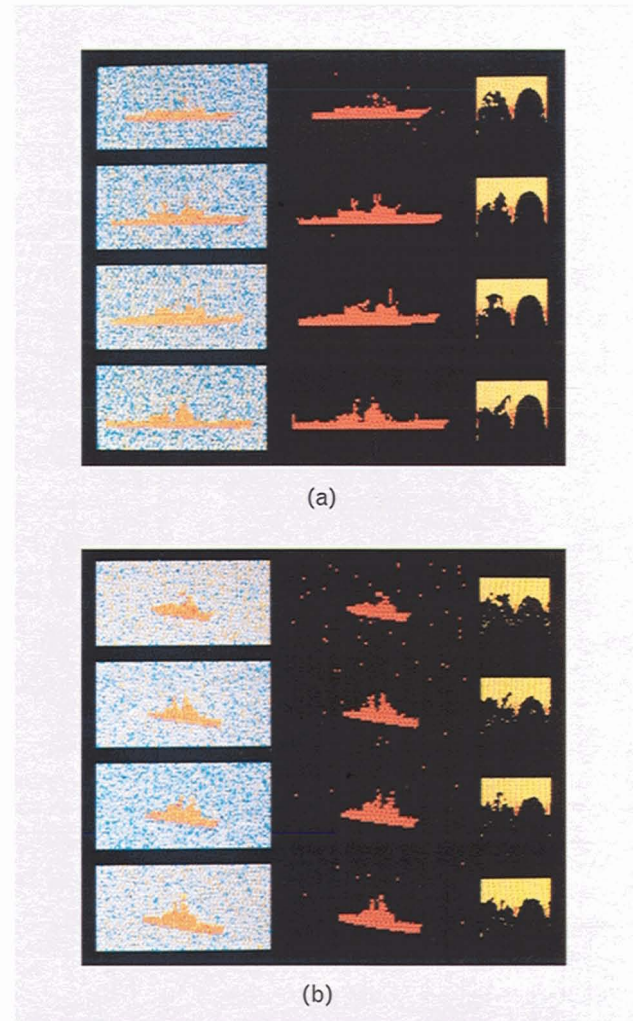


FIGURE 8. Target segmentation and invariance transformation performed on the four ships of Figure 2 at (a) 90° and (b) 30° viewing aspects. The left column of images shows the original synthetic passive-IR imagery, the middle column shows the silhouettes after the application of thresholds, and the right column shows the log-radius maps (LPM) of the silhouettes. The LPM provides a useful feature pattern that is invariant with respect to translation, in-plane rotation, and scale changes.

were distinct enough to be used for identification and (2) to validate the accuracy of the synthetic imagery.

The Infrared Airborne Radar (IRAR) System

To evaluate the active portion of the proposed NCTI sensor suite, we used the Infrared Airborne Radar (IRAR) [14] to collect data for the range profile studies. IRAR is a forward-looking multidimensional infrared sensor operated by the Opto-Radar Systems

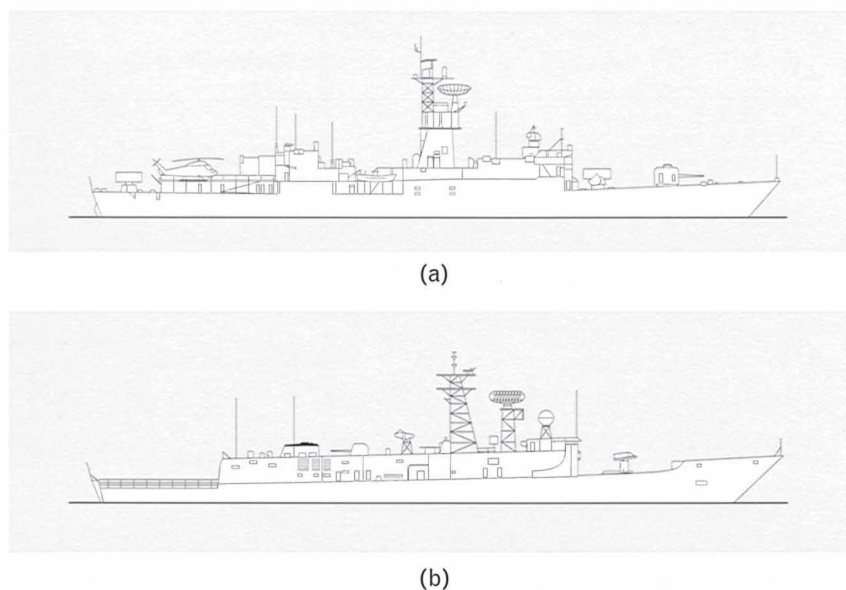


FIGURE 9. Ship targets used for the range-profile measurement exercises: (a) the USS *Connole* (FF 1056), a Knox-class frigate, and (b) the USS *Samuel B. Roberts* (FFG 58), a Perry-class guided-missile frigate. The drawings are from *Jane's Fighting Ships, 1991–1992*.

Group at Lincoln Laboratory. In its usual mode of operation, IRAR collects pixel-registered laser radar intensity, range, and passive-IR images. The active channel of the sensor uses a waveguide CO₂ laser operating at 10.6 μm to measure absolute range with 1-m precision. The passive-IR channel measures target and scene thermal characteristics in the 8-to-12- μm band. The system contains two linear 12-element arrays of HgCdTe photovoltaic detectors. One detector array is dedicated to the active channel and the other array to the passive channel. In typical operation, the arrays are oriented vertically and are scanned perpendicular to the airplane's motion.

The IRAR system was reconfigured to test the concept of range profiling for ship targets at distances up to about 5 km, which is farther than the system's usual operational range. Measurements at these longer ranges required an increase in the system's carrier-to-noise ratio (CNR), which we accomplished by removing the beam expander and focusing the laser so that only one detector would be illuminated. With only one detector being scanned instead of 12, the scan rate had to be increased to maintain contiguous coverage on the ground. At the higher scan rate, however, each scene element was oversampled. We compensat-

ed for this effect in software, as described in the subsection "Data Enhancement and Compensation."

The passive detector array was rotated 90° so that the array was scanned along its length rather than perpendicular to its length, as is the usual case. Thus each scene element was sampled 12 times during a scan. To take advantage of this oversampling, time-delay-and-integration (TDI) processing was applied to the passive-IR data in software. The TDI processing improved the effective SNR by a factor of approximately $\sqrt{12}$.

Measurement Exercises

The reconfigured IRAR sensor flew two missions on board a Gulfstream G-1 aircraft to collect data from two target ships. Figure 9 contains line drawings of the two ships: the USS *Connole*, a Knox-class frigate (FF 1056) that is 134 m long, 14.3 m wide, and 4.6 m high from the waterline to the main deck; and the USS *Samuel B. Roberts*, a Perry-class guided-missile frigate (FFG 58) that is 138 m long, 13.7 m wide, and 4.5 m high from the waterline to the main deck.

Data Processing

An example of the raw oversampled data is shown in

Figure 10 for a broadside view of the *Connole*. To process and analyze the data, we developed software for removing the overscanning effect, increasing the effective SNR of the remaining data, and correcting the range values and geometric distortions (created by the motion of the ship and aircraft). The corrected range and intensity data were then combined to produce a range profile: a plot of intensity as a function of range.

Data Enhancement and Compensation

With a single “look” at a diffuse target, the effective SNR cannot exceed unity for high CNR [15]. In these data-collection exercises, however, each scene element was overscanned; i.e., there was more than one measurement of intensity, range, and passive-IR data at each imaged pixel area. Consequently, to take advantage of the multiple measurements, we applied processing techniques to each data domain in order to increase the effective SNR beyond unity.

For the intensity data, we averaged the overlapping values at each pixel. Because each intensity value in an overscanned scene element represented an independent measurement, this frame-averaging process improved the SNR by reducing the speckle effect. When N scans of data are averaged, the improvement in SNR is proportional to \sqrt{N} . At the longest ranges, the number of overlaps in the data was between 8 and 12. Thus the frame averaging improved the SNR by a factor of approximately 3.

The range data represented instantaneous range values from the moving aircraft to the ship. To create an accurate range profile, we had to correct these values so that they would represent measurements from a single stationary point of view. Thus range differences from scanline to scanline were taken into account to produce a snapshot type of image. We then applied a histogram-cluster-average (HCA) filter to the motion-corrected (but still overscanned) range data to improve the range precision.

As with the intensity and range data, the passive-IR data also represented overscanning of each scene element. Thus we averaged the overlapped passive-IR values to improve the SNR, as we had done with the active intensity data. In addition, because the 12-element passive-IR detector array had been rotated so that it scanned along its length, the detectors imaged a

given spot 12 times with each scan in a time-delayed fashion. Readouts from all 12 detectors were recorded and a TDI performed on the data in software. With TDI processing and by averaging the typical 8 to 12 overscans, we improved the SNR of the processed passive-IR data by a factor of approximately 10.

Next, the target was segmented from the sea background in the active imagery by applying a threshold to the intensity data. Values below the threshold were set to zero in the intensity image, and the corresponding pixels in the registered range image were also set to zero. Figure 11 contains an example of the processed and segmented active imagery. We filled in the drop-outs in the imagery by using the techniques described earlier, and we removed most of the background noise by applying a threshold. The pseudo-color stripes in the range image indicate that the aircraft's forward motion has been removed. (The pseudo-color technique assigns a single color to a user-defined number of range counts on the displayed image.) Even after compensation for overscanning and for the aircraft's




FIGURE 10. Example of laser radar overscanned data: (a) intensity and (b) range imagery for a broadside view of the *Connole*.





FIGURE 11. Example of speckle reduction and segmentation of laser radar ship data: (a) intensity and (b) range imagery for a broadside view of the *Connole*. Note that the ship mast appears tilted and the bottom of the ship is not horizontal. In the range image, a pseudo-color technique was used to assign a single color to a user-defined number of range counts in the image.



forward motion, geometric distortions are still apparent in the images. These distortions include a general rotation of the entire ship and an arching of the central mast toward the stern. To correct for such distortions, we used the following processing steps.

During each approach of the aircraft to the ship, the sensor's line of sight was generally not aligned with the aircraft's centerline. This angular offset created an artificial cross-range motion between the ship and the sensor from one scan to the next. Consequently, tall structures, such as masts, appear tilted in the resulting imagery, as in Figure 11. In addition, the ship's own motion contributed to the cross-range distortion. To correct for these effects, we decomposed the ship and aircraft motions into components perpendicular to the sensor's line of sight, as shown in Figure 12. Using this vector analysis, and including the overscanning effect, we derived the following equation for ΔP , the azimuthal shift, of a pixel per scan:

$$\Delta P = \frac{N[V_a \sin \theta - V_s \cos(\beta_a + \theta + 270^\circ - \beta_s)]}{R\phi f},$$

where V_a = aircraft speed in m/sec, β_a = aircraft head-

FIGURE 12. Vector analysis of ship and aircraft motions for correcting geometric distortions. For more details of the analysis, see the main text.




FIGURE 13. Example of laser radar imagery corrected both for relative motions between the sensor and target ship and for sensor mirror rotation: (a) intensity and (b) range imagery for a broadside view of the *Connole*. Note that, compared with the imagery of Figure 11, the ship mast has been straightened, and the bottom of the ship is now horizontal rather than tilted.

ing in degrees (clockwise from magnetic north), V_s = ship speed in m/sec, β_s = ship heading in degrees (clockwise from magnetic north), θ = azimuthal offset of the sensor line of sight (LOS) from the aircraft heading in degrees, ϕ = IFOV in radians, R = ship range in meters, N = number of overscans, and f = pulse-repetition frequency of the sensor system. This correction was computed at each pixel location and propagated through every scanline to obtain geometrically correct images.

Finally, the mirror system of the IRAR sensor causes a rotation of the entire image. A complex 3-dimensional ray-tracing calculation by R.J. Hull [15] showed that θ_v , the amount of rotation for a vertical object, is given by

$$\tan \theta_v = -\sin \gamma \tan \psi,$$

where γ is the mirror depression angle and ψ is the mirror pointing angle. In general, the mirror pointing angle is 45° . When the depression angle is small, as is the usual case, the equation reduces to $\theta_v \approx -\gamma$. This result implies the image should be corrected by a clockwise rotation of the same magnitude as the mirror depression angle (confirming a rule of thumb suggested by D.G. Biron [16]).

Figure 13 shows an example of imagery that has been corrected both for motion distortions and mirror rotation. In particular, note that the ship mast has been straightened, and the bottom of the ship is now

horizontal rather than tilted (compare to Figure 11). These corrected data were then passed to the range profiling code.

Range Profiling

The final piece of software took the motion- and range-corrected data from the previous stage and created a range profile—a plot of returned intensity versus range.

In the proposed system, range profiling enables ship classification. To obtain a range profile, we would record a time history of amplitude returns for each pixel as the sensor scans a target. Signals would be received first from those parts of a ship which are closest to the sensor, while reflected signals from portions farther away would be sensed later. Thus the intensity of the return as a function of time (range) is recorded for each pixel, and the ship's signature may be obtained by combining all of the range profiles that are generated from scanning the sensor's FOV. With the design parameters provided by Marcus [1], a single pixel of the proposed detector would cover an area $18 \text{ m} \times 18 \text{ m}$ on the ship at a nominal operating range of 100 km.

The IRAR sensor, however, is a peak-detecting system. That is, for each pixel only the peak value of the returned signal intensity is recorded with the time-of-flight (range) measurement of that peak signal. Thus many scans of the reconfigured (one-detector) IRAR




FIGURE 14. Measured range profile for a near-stern view of the *Roberts*.

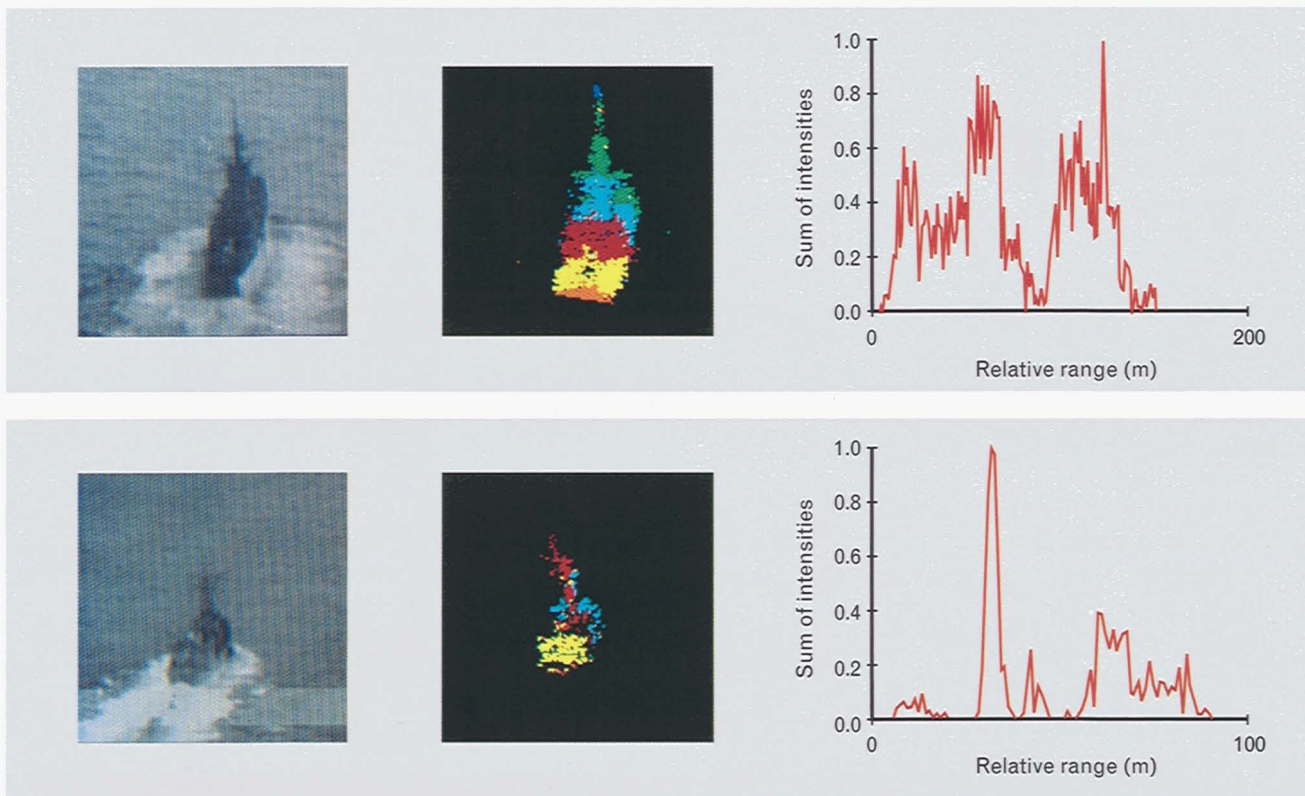


FIGURE 15. Comparison of near-stern views of (top) the *Connole* and (bottom) the *Roberts*. From left to right, the figure contains video images of the two ships, the measured range images, and the resulting range profiles. The *Connole* was at a range of 3 km, and the *Roberts* at 3.9 km.

sensor must be used to obtain spatially resolved images of a ship in both the range and intensity domains. Range profiles can then be created by summing the intensity values into their corresponding range bins, with each bin representing 1 m, the downrange resolution of the IRAR sensor.

At typical mission ranges, an IRAR pixel covers an area about $1\text{ m} \times 1\text{ m}$ on a target ship. This area is a sub-pixel, in a sense, of the proposed NCTI sensor. Considerations described by S.J. Rak [17] indicate that range profiles created from the IRAR resolved peak-detected data would be representative of profiles generated by the unresolved, range-intensity signals from the proposed NCTI sensor.

Initially, we tested the range profiling code by using measurements taken of an airplane parked on the ground. Gross features such as the airplane's length were discernible from the profile. Next, we created a profile of a synthetic range image of a ship. Detailed ship features could be identified in the range profile. This success provided confidence in using profiles for ship identification. Examples from the measurements of the *Connole* and the *Roberts* are shown and discussed in the following section.

Discussion of Range Profiles from Ships

A few examples from the measured database demonstrate some of the features and usefulness of range profiles. The examples include an interpretation of a single profile and a comparison of profiles from the *Roberts* and the *Connole*.

Figure 14 contains an example of a range profile of the *Roberts*, viewed from just off the ship's stern. An analysis of the profile and the original image determined that the higher peaks in the profile corresponded to regions of the ship with the most surface area. Thus the highest peak represented returns from the hangar door at the rear of the ship. This door, which was at a fixed range, presented the largest reflective surface to the sensor. The other peaks in the profile represented (from left to right) returns from the ship's smokestack and other features on top of the superstructure, the main mast, the secondary mast in front of the main mast, and the bridge area of the ship.

A major task of these missions was to determine

whether range profiles could be used to discriminate one ship from another. Figure 9 shows that the *Connole* and the *Roberts*, which belong to different ship classes, do appear different from each other. For example, the superstructure of the *Connole* has a relatively irregular shape, while the superstructure of the *Roberts* is more uniform and rectangular. In addition, the *Connole* has one massive main mast, while the *Roberts* has two slender masts. Such structural differences should be reflected in range profiles from the two ships, and the differences should be especially apparent at near-bow or near-stern viewing angles. Figure 15 shows profiles from the two ships at about the same range and same near-stern viewing aspect. Clearly, the profiles have different shapes. For example, the profile from the *Connole* shows a good deal of structure in the flight-deck region (behind the hangar-door spike), while the profile of the *Roberts* shows few features in that area. The reason for this difference can be seen in the ship outlines of Figure 9. The *Connole* has a stepped rear deck behind the hangar with a gun mounted on the lower level. Also, during the measurement exercise, a helicopter had been parked on the ship's flight deck. The stern region of the *Roberts*, on the other hand, has a uniform structure with a flat landing pad, hence the lack of features in that area of the profile.

Validation of Synthetic Range Profiles

The second major purpose of the flight missions was to collect data with which to validate the synthetic ship models described earlier. We chose the *Duncan* (FFG 10) as the model for comparison with the *Roberts* (FFG 58) because both ships are Perry-class guided missile frigates, although the *Duncan* was built earlier than the *Roberts*.

Stern views of the ships were desirable to yield range profiles with many features for comparison. From the limited database of measured data, we selected a view of the *Roberts* at 5° off the stern. Because the data were recorded at a range of about 4 km with a sensor depression angle of 6.7° , each pixel covered a square of about $80\text{ cm} \times 80\text{ cm}$ on the ship. For comparison with these measured data, we generated a view of the *Duncan* model at a similar aspect angle, depression angle, and pixel size, and used this noiseless syn-

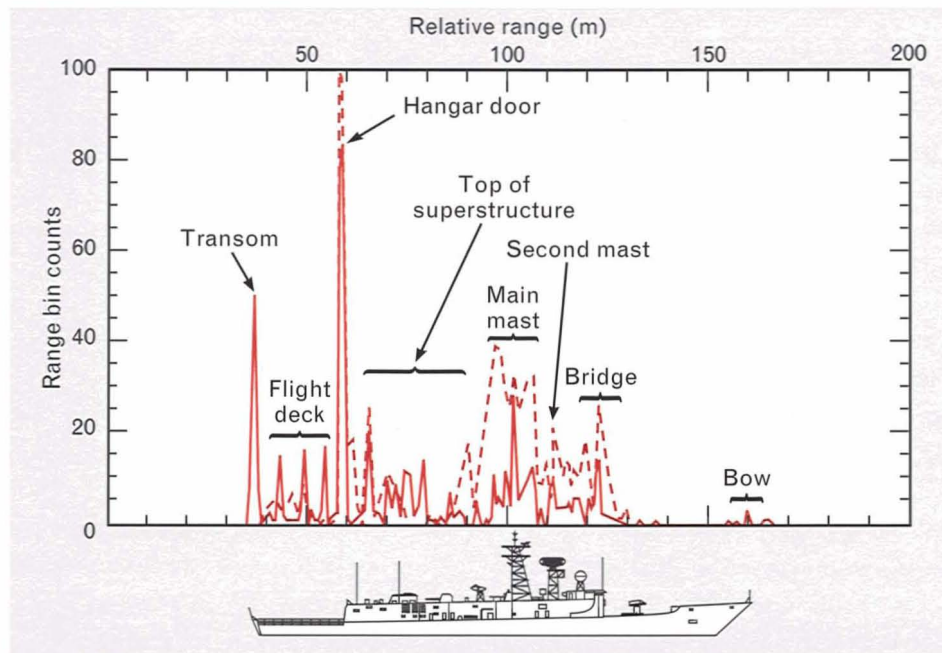


FIGURE 16. Comparison of synthetic and measured range profiles for two similar ships. The solid line represents the synthetic profile of the *Duncan* (FFG 10) and the dashed line represents the measured profile of the *Roberts* (FFG 58). Both profiles are for near-stern views at 4 km.

thetic image to obtain a range profile of the ship.

Figure 16 shows both the measured profile of the *Roberts* and the synthetic profile of the *Duncan*. The similarities of the two profiles indicate that the measured ship data validate the accuracy of the model.

There are, however, a few differences between the two profiles. At the stern end of the profiles (the left end of the range axis), the synthetic profile of the *Duncan* has a large spike, representing the ship's transom, while the measured profile of the *Roberts* contains no such spike. This difference exists because the ships were constructed about six years apart. The *Duncan* was built with its transom relatively vertical. By the time the *Roberts* was built, the flight deck of Perry-class frigates had been extended backwards, and this extension required the transom to be angled down at about 45° [18]. Because a surface at this angle reflects laser energy toward the water and not up toward the sensor's detector, the measured profile of the *Roberts* shows no returns at that location.

From left to right along the synthetic profile of the *Duncan*, three distinct spikes are apparent in the region representing the flight deck. Note that these

three spikes do not appear in the measured profile of the *Roberts*. The spacing between the spikes represents the range discretization of the software range sensor model that was used to calculate the range image. The software behaves, in some sense, like a first-return-detected sensor. The IRAR sensor, however, is a peak-detecting sensor. Thus the three spikes do not appear in the measured profile. Furthermore, because the IRAR sensor has the statistical uncertainty inherent in a real system, the measured range values from the flight deck appear as clusters rather than discrete spikes.

The peaks representing the main mast of the *Duncan* model are smaller (relative to the hangar-door spike) than the corresponding peaks in the measured data from the *Roberts*. The model of the *Duncan* assumes that all surfaces are diffuse reflectors, whereas the main mast of the *Roberts* is constructed of many narrow, interlacing metal struts that produce specular, as well as diffuse, returns. Because the gain of specular over diffuse cross section can be substantial (as shown by Hull [19]), we expected this disparity in the main-mast peaks of the two profiles.

Despite these differences, the two profiles are strikingly similar. In both profiles, the same length for the ship from stern to bridge is displayed, the spike representing the hangar door is clearly visible, small peaks representing identifiable features on top of the superstructure are present, the main mast appears at the same relative location, a spike corresponding to the mast just forward of the main mast can be identified, and, last, structures representing bridge features are shown in the same locations.

Because of these similarities and because of our ability to account for the major differences described above, we feel that the synthetic image of the *Duncan* accurately models the measured data from the *Roberts*. Thus the classification results reported below should apply to measured ship data in addition to synthetic imagery.

Neural Network Classifiers

Artificial neural networks are making a dramatic comeback this decade because of advances in computer hardware technology and the ease of implementation of these types of networks in massively parallel architectures. Because they are capable of learning from sample data and because they are tolerant of noise, artificial neural networks are frequently used for pattern-recognition applications. The utility of artificial neural networks for image enhancement and target classification of range imagery has been demonstrated by the Opto-Radar Systems Group at Lincoln Laboratory under the Autonomous Infrared Sensor Technology Program [13] of the Defense Advanced Research Projects Agency (DARPA).

A few of the issues investigated under the ship-identification program were the range-profiling resolution necessary for target classification, the sensitivity of the profiles to changes in viewing angle, and the angular resolution and SNR required by an imaging system to identify targets successfully. For target classification using either range profiles or features extracted from passive-IR imagery, we applied two artificial neural networks: a multilayered neural network (MNN) [20] and the Adaptive Clustering Network (ACN) classifier, which is similar to the adaptive resonance theory (ART) [21] neural network.

The MNN is a feedforward multilayered network

that constructs an *internal representation* by learning from training data. The internal representation is a set of linkweights that defines the connections between the nodes in either the input or the output layer and the nodes in one or more of the hidden layers. Once training is completed and the internal representation has been learned, the representation allows the MNN to map an unknown input to a particular output category that has been previously defined (in the training process) by similar inputs. The learned internal representation thus enables unknown input patterns to be classified.

The ACN is an unsupervised classifier that determines its own categories, rather than having the category separation determined *a priori*. The formation of categories is a function of an association threshold parameter (called the *vigilance* ρ) and the rate at which the memory is updated (called the *learning rate*).

The following subsections discuss the training algorithms for both types of neural network classifiers together with their algorithm design parameters.

Multilayered Neural Network

Feedforward MNNs have been attracting much interest in target-recognition applications ever since D.E. Rumelhart [20] developed a learning algorithm called backward error propagation (BEP) for this type of MNN. With the BEP algorithm, MNNs have been applied successfully in many experimental works. For the conventional BEP, however, the convergence speed had been slow. Consequently, M.M. Menon and E.J. Van Allen [22] developed a deterministic gain-annealing mechanism in BEP to speed up the convergence. We have employed this gain-annealing BEP training algorithm in our MNN classifier, using range profiles as input target features.

Our MNN classifier contains three layers of fully interconnected nodes in a feedforward network. Figure 17 shows the architecture of the network. A sigmoid function is used for the node function at each node output:

$$x_j^n = a \left(\frac{1}{1 + \exp \left[\beta \left(\sum_i w_{ji}^n x_i^{n-1} + \theta \right) \right]} + c \right),$$




FIGURE 17. The multilayered neural network (MNN) classifier: (a) diagram of three layers of interconnected nodes and (b) activity at the j th node in the n th layer with linkweights w .

where x_j^n is the output of the j th node in the n th layer. To minimize the error, the linkweights w_{ji}^n as well as the gain β and offset θ are allowed to vary from node to node. For the experiments described in this article, we kept the scaling parameter a and offset parameter c at constant values of 2.0 and -0.5 , respectively.

The gain-annealing algorithm is based on the hypothesis that the optimum representation at a low gain (small β) is a good approximation for the optimum at a high gain (large β). At low gain, the optimization problem is easy to solve. Hence, by starting at a low gain and by slowly raising the gain levels in the cycles of the gain-annealing process, we can make the iterative BEP converge much faster than the conventional algorithm.

For our MNN classifier, the input layer has approximately 250 nodes, corresponding to the number of range bins in the simulated range profiles. We have chosen to use 12 nodes in the top (output) layer to map each of the four ships into 3 separate aspect categories (0° – 27° , 30° – 57° , and 60° – 90°). For the hidden layer, the number of nodes is determined automatically by the training data. During training, a new node (up to 25 total) is added to the hidden layer until the error tolerance has been satisfied.

Adaptive Clustering Network

The ACN [23] is an unsupervised neural classifier

that clusters input vectors into a few stable categories. Figure 18 shows a schematic diagram of the ACN.

At the beginning of a training cycle, the long-term memory (LTM) traces are initialized to the first input vector. During training, each input vector \mathbf{X}_i is compared with all of the stored LTM traces \mathbf{W}_k to find the best match, which is given by

$$\rho_{\max} \equiv \max_k \frac{\mathbf{X}_i \cdot \mathbf{W}_k}{\|\mathbf{X}_i\| \cdot \|\mathbf{W}_k\|}.$$

If ρ_{\max} for a given input vector \mathbf{X}_i is greater than the preset vigilance ρ , the input vector \mathbf{X}_i is “learned” onto the LTM according to

$$\mathbf{W}_k = \mathbf{W}_k + (\mathbf{X}_i - \mathbf{W}_k)\Delta t,$$

where Δt is the learning rate. If ρ_{\max} does not exceed ρ , the ACN creates a new category for the input vector \mathbf{X}_i . Training data are presented repeatedly to the network until the categories formed no longer change. Once the LTM traces have stabilized, a test pattern can be presented to the ACN and classified by its correlation with the LTM traces.

Classification Results Using Synthetic Range Profiles

The synthetic range profile data for training consisted of 124 ideal profiles (four ships, at 31 different view-

ing aspects, from bow to port) for each transmitter pulse width (10, 30 and 50 nsec). We eliminated speckle from the profiles and normalized them with respect to their peak values. Figure 4 shows four examples.

The test data contained speckle-reduced averages of 32, 16, 8, and 4 profile frames to simulate various SNRs. There was a total of 12 sets of testing data: for each of the three transmitter pulse widths, there were four sets (one for each of the four different frame averages) of 124 profiles. Figure 5 contains examples of the test data.

Using the range profile data, we tested the effectiveness of both the MNN and the ACN. A strength of the ACN is its ability to separate the training data automatically into different categories without specific supervision. The resulting categories revealed the intrinsic classes in the data. MNN, on the other hand, forced the separation of the data into different categories by design. The training results from the MNN gave a measure of how successfully this could be done.

MNN Results

As mentioned previously, the MNN classifier was designed for a 12-category problem—three aspect groups (0°–27°, 30°–57°, and 60°–90°) for each of four ships—at a given pulse width. We used the ideal range profiles to train the MNN classifier and the frame-averaged noisy profiles to test the classifier. Table 2 contains the results for all four frame averages of testing data for a pulse width of 30 nsec. Classification results for the training data have also been included to show how well the network could be trained.

We designed these initial experiments for the case in which there was no *a priori* information about either the viewing aspect or the ship class. A classification was called correct only if both the ship ID and the viewing aspect were classified correctly.

Not surprisingly, the most difficult profiles to train on and to classify were those from the broadside or near-broadside aspects. Specifically, training of the MNN classifier started to break down somewhere in the 30°–57° viewing range. For 60°–90° aspects, the classifier could not be trained into the four ship classes

at all. Also, as expected, the classifier was more effective in separating profiles that had higher SNR (i.e., profiles that had been obtained through an averaging of a greater number of frames). Lastly, Table 2 shows that the MNN could classify nearly 80% of the target ships correctly for 0°–27° (near-bow) aspects with 16 frames averaged and a transmitter pulse width of 30 nsec. Similar results were found for pulse widths of 10 and 50 nsec.

ACN Results

Using the ACN, we examined the effectiveness of range profiling to determine both the amount of category compression that was possible during training, and the rate of correct classification during testing. We also studied category formation as a function of viewing angle to determine the angles at which the profiles were no longer distinguishable. This informa-

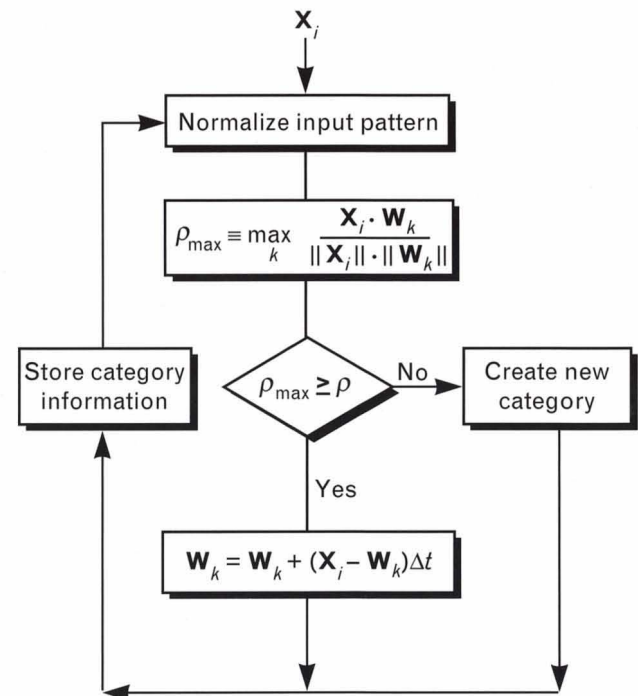


FIGURE 18. Algorithm for the Adaptive Clustering Network (ACN). During training, each input vector \mathbf{X}_i is compared with all of the stored long-term memory (LTM) traces \mathbf{W}_k to find the best match, which is called ρ_{\max} . If ρ_{\max} for a given \mathbf{X}_i is greater than a preset vigilance ρ , \mathbf{X}_i is “learned” onto the LTM at a learning rate of Δt . If ρ_{\max} does not exceed ρ , a new category is created for the input vector \mathbf{X}_i .

Table 2. MNN Classification Results¹ for Synthetic Range Profiles of Four Ships at Different Viewing Aspects

	Viewing Aspect			Overall
	0°–27°	30°–57°	60°–90°	
Training Data	100%	93%	43%	77%
Testing Data ²				
32	83	60	39	60
16	78	40	34	50
8	63	20	25	35
4	48	25	23	31

¹Results are given in percent of profiles correctly classified by the system for a 30-nsec pulse width.

²Four different datasets were used for testing. The datasets were created by averaging different numbers (32, 16, 8, and 4) of profile frames to investigate classifier performance at enhanced SNR levels.

tion would indicate whether, and at which aspects, another discriminator (e.g., passive-IR data) would be necessary.

We presented the ACN with a four-class problem that consisted of range profiles created from the four synthetic ship models at 31 viewing aspects each. The training results determined (1) the number of categories formed by the 124 range profiles and (2) the viewing aspects at which the categories became confused.

Before the ACN could be trained, we had to choose a value for the vigilance ρ , which controls the number of categories formed, given a particular

input set. In general, a high vigilance value should be used to separate input patterns having small feature differences, while a low vigilance value should be used for a rough separation of patterns having large differences. Thus the vigilance should be chosen based on the particular application at hand. For our analysis, we chose five vigilance values: 0.99, 0.95, 0.90, 0.85, and 0.80.

During the training process, the 124 noise-free expected-value range profiles (with a given pulse width) were presented repeatedly to the classifier until the LTM traces had stabilized. The procedure was repeated for each of the three pulse widths and for the five vigilance values.

First, we examined the training results to determine the amount of category compression achieved by the ACN. For example, with a vigilance of 0.95, the classifier formed 26 categories for the 10-nsec pulse, 18 categories for the 30-nsec pulse, and 16 categories for the 50-nsec pulse (Table 3). These results represent memory compression factors of 5, 7, and 8, respectively. As expected, the number of categories decreased with increasing pulse width because longer pulses do not resolve ship features as well, causing the profiles to look similar to each other. Shorter pulses yield better range resolution, thus emphasizing differences between the profiles. This trend was also appar-

Table 3. ACN Training Results¹ for Synthetic Range Profiles of Four Ships

Pulse Width (nsec)	Number of Categories Formed
10	26
30	18
50	16

¹Results are for 31 viewing aspects of each of the four ships with a vigilance of 0.95.

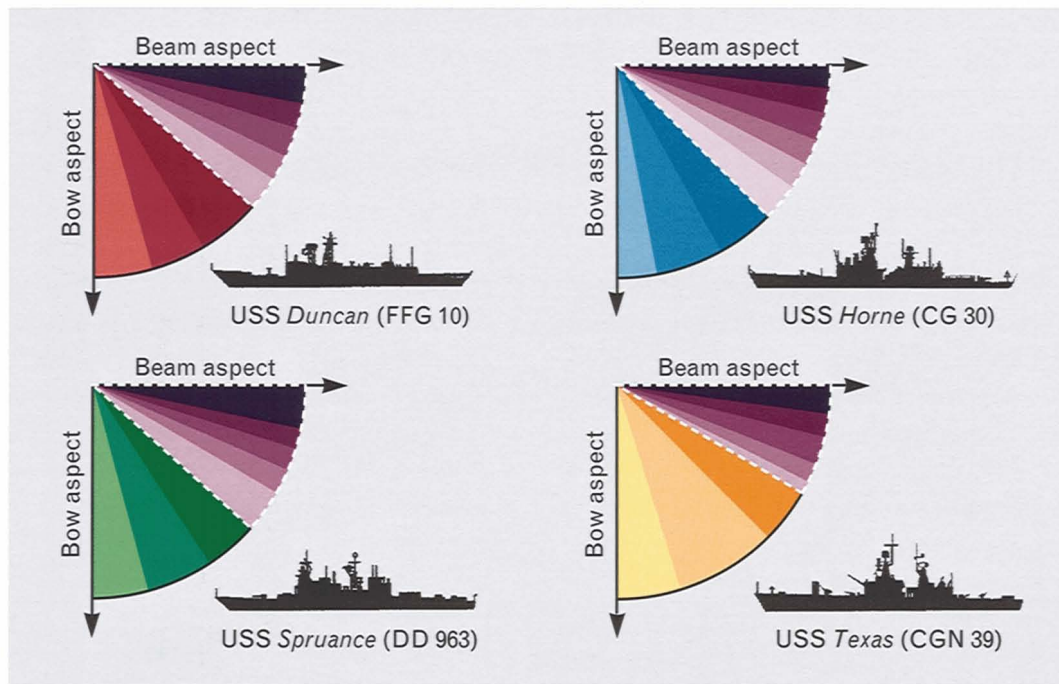


FIGURE 19. Typical ACN category formation for ship range profiles. The four different colors (red, blue, green, and yellow) in the lower parts of the pie charts represent the four ship classes. For a given ship, each shaded wedge represents those viewing angles which the classifier grouped into a single category. For example, the results for the *Duncan* show that all viewing aspects from 0° (bow) to about 50° form only three distinct categories (the three shades of red). The confused categories, represented by the different shades of purple bounded by the white dashed lines, begin to form at about 50°. These results are for a pulse width of 30 nsec, a vigilance of 0.95, and 16-frame averages.

ent for the other vigilance values.

Next, we studied the training output to determine the angles at which the range profiles became difficult to distinguish. As discussed earlier, range profiles contain the most information when a sensor views a ship at aspects providing the most range depth. When a ship is viewed from just off the bow or stern, almost the entire ship length contributes to the range profile. Thus the classifier should be able to discriminate easily between the profiles created at near-bow and near-stern aspects. When a ship is viewed broadside, however, it presents far less information to the range sensor because nearly the entire ship is at the same range. Thus the classifier should have more difficulty distinguishing the profiles at near-broadside aspects.

Figure 19 shows the category-formation results for range profiles at a pulse width of 30 nsec. The four different colors (red, blue, green, and yellow) in the

lower parts of the pie charts represent the four ship classes. For a given ship, each shaded wedge represents those viewing angles which the classifier grouped into a single category. For example, the results for the *Duncan* show that all viewing aspects from 0° (bow) to about 50° form only three distinct categories, as denoted by the three shades of red. The figure shows that confused categories, represented by the different shades of purple in the upper sections of the pie charts, begin to form around 50°. Similar behavior was found for the other pulse widths. For those viewing angles at which the profiles become indistinguishable, another discriminator is needed. This analysis implies that the critical viewing angle is about 50° off the bow (or stern).

After the LTM traces of the stable categories had been established with training data, we presented the test data to the ACN. For each input test pattern, the classifier computed the correlation between the pat-

tern and each of the existing LTM traces. If the highest correlation exceeded the set vigilance, the classifier assigned the input pattern to the category corresponding to the LTM trace that produced the highest correlation value. When the input pattern and the assigned category were of the same ship, we scored the match as a correct classification, regardless of the viewing aspect. This procedure was different from that used by the MNN classifier, which required both the ship class and viewing aspect to be correct.

As before, we created test data of 4-, 8-, 16-, and 32-frame averages to study the classification performance as a function of SNR. The test sets consisted of these frame-averaged profiles at all 31 aspects for the four ships and for each pulse width.

During the testing process, we kept the learning rate at the same value as that used during the training phase. We set the vigilance to a small, near-zero value to ensure that the test patterns would be classified into one of the established categories. (Alternatively, the vigilance could have been set to a higher value and the patterns that did not correlate well with any of the LTM traces could have been assigned to a category labeled "unknown.") All of the test results reported in this section are from training with a vigilance of 0.95.

We studied the classification performance of the

ACN in terms of the same 30° increments described earlier. Table 4 shows the results of this analysis for the 30-nsec pulse width. For a given number of frames averaged, performance is best for the 0°–27° (near-bow) viewing aspects because of the reasons described earlier. In fact, classification for this range of angles is quite good for the 32- and 16-frame averages. Classification begins to degrade for the 30°–57° aspects, as the profiles become narrower. Finally, the test results fall off dramatically for the 60°–90° aspects. The performance at these aspects is not much better than if the profiles had been assigned randomly to one of the four ship classes. These results reflect the confusion in category formation during the training phase.

Classification Results Using Synthetic Passive-IR Imagery

Using the ACN classifier, we tested the effectiveness of synthetic passive-IR imagery for ship identification. The classification results for the passive-IR ship silhouettes can be presented similarly to the range-signature results in that both have a strong dependence on ship aspect. With range profiles, the range depth of a ship was shown to be the critical factor. For aspects near broadside (representing the minimum target range depth), the recognition performance de-

Table 4. ACN Classification Results¹ for Synthetic Range Profiles of Four Ships at Different Viewing Aspects

	Viewing Aspect		
	0°–27°	30°–57°	60°–90°
Training Data	100%	93%	43%
Testing Data ²			
32	83	60	39
16	78	40	34
8	63	20	25
4	48	25	23

¹Results are given in percent of profiles correctly classified by the system for a 30-nsec pulse width and a training vigilance of 0.95.

²Four different datasets were used for testing. The datasets were created by averaging different numbers (32, 16, 8, and 4) of profile frames to investigate classifier performance at enhanced SNR levels.

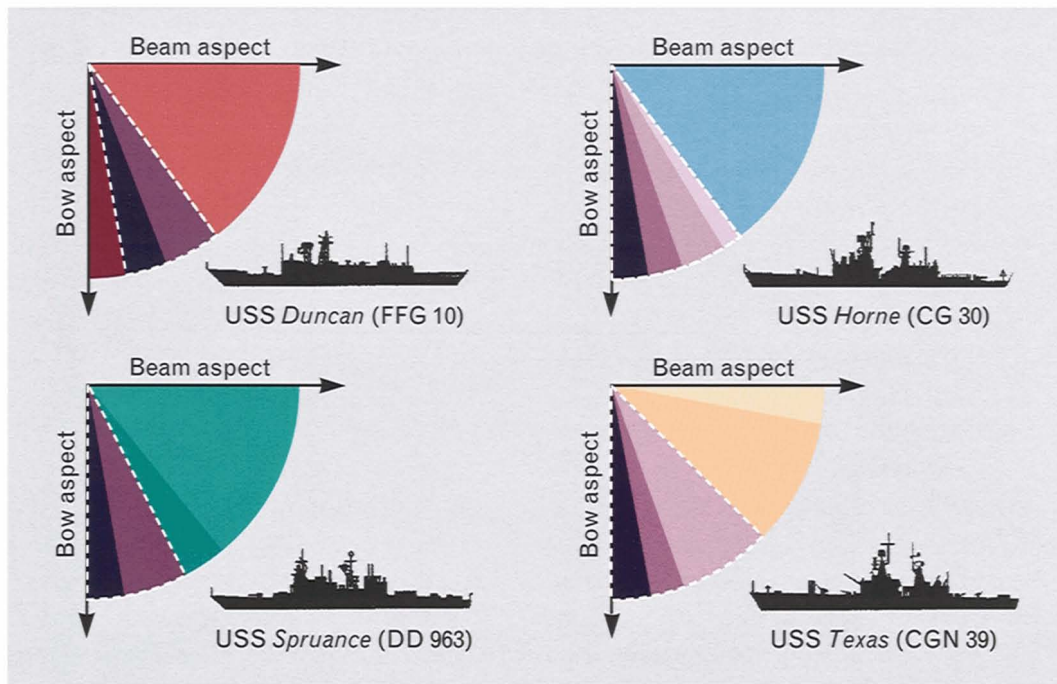


FIGURE 20. ACN category formation for passive-IR ship silhouettes. As in Figure 19, each red, blue, green, and yellow shaded wedge represents viewing aspects that have been mapped into a single category. The different shades of purple bounded by white dashed lines represent confused categories.

creased dramatically. We expected the opposite to be true for passive-IR ship silhouettes. At broadside aspects, a ship's silhouette is very distinct. At near-bow (or near-stern) aspects, on the other hand, most ship silhouettes look alike, which should result in a decrease in recognition performance.

As described earlier, the passive-IR inputs to the classifier were log-polar-mapped target silhouettes. The database consisted of four ships, each at 31 views from 0° (bow) aspect to 90° (beam) aspect in 3° increments, for a total of 124 inputs.

The goal of the training phase was to determine the classifier vigilance that would provide the greatest target separability while minimizing the overall number of categories. We initialized the ACN memory by training the classifier with noiseless ship silhouettes (i.e., ship silhouettes with infinite SNR) scaled to a range of 25 km. Using noiseless silhouettes ensured that a highly detailed silhouette representation was learned onto the ACN LTM traces. (Note: All recognition testing for ship silhouettes was performed against this benchmark.) Using a training vigilance

of 0.87, the ACN mapped the 124 inputs into just 11 categories. This level of memory compression was consistent with the level achieved by similar processing methods applied to other 3-D targets rotated out of plane [13].

Figure 20 shows the category distribution within the classifier memory as a function of angle, from bow to beam aspect. As before, the shaded wedges in the pie charts indicate the range of aspects that the classifier has mapped into a single category. The four different colors (red, blue, green, and yellow) indicate the unique classes for each ship. Note that more categories were formed at near-bow aspects because the silhouette changes dramatically with angle at these aspects. In addition, note that at 0° – 30° aspects the ships became indistinct to the classifier; that is, several ship inputs fell into a single category.

In the testing phase, we presented the noisy datasets to the classifier to determine the sensor operating characteristics that could be tolerated. The 25- and 50-km imagery achieved correct recognition levels of 77% and 68%, respectively, with 16-frame averages.

Table 5. ACN Classification Results¹ for Synthetic Passive-IR Imagery of Four Ships at Different Viewing Aspects

<i>Test Set</i>	<i>Apparent Range (km)</i>	<i>16-Frame SNR</i>	<i>Viewing Aspect</i>		
			<i>0°–27°</i>	<i>30°–57°</i>	<i>60°–90°</i>
25 km	25	164	57%	93%	93%
50 km with 2× super-resolution	25	58	48	93	93
50 km	50	58	41	89	79

¹Results are given in percent of profiles correctly classified by system. All datasets were tested against the 25-km noiseless training set.

This result was about the same as the performance of the ACN on the training data. Further experiments measured the effectiveness of using super-resolution (SR) to enhance the recognition performance. The 50-km, 2× SR imagery, effectively scaled to a range of 25 km, achieved a correct recognition performance of 76%, almost as good as the 25-km imagery itself. (Note: As mentioned earlier, all of the imagery was tested after the classifier had been trained with the 25-km noiseless silhouettes.)

We ran additional experiments to determine the target aspects at which the classifier was able to distinguish the ships. These results, summarized in Table 5, show that the classifier was clearly able to distinguish ships at aspects of 30°–57° and 60°–90°, achieving correct recognition levels generally greater than 90%. For aspects of 0°–27°, the correct recognition level was only 40% to 60%.

Summary and Conclusion

A multidimensional sensor suite consisting of a laser radar and a passive-IR sensor has been evaluated for detecting and identifying ships at long ranges from an airborne platform.

Active and passive-IR measurements of two Navy frigates belonging to different classes were taken with an existing airborne, multidimensional sensor system. The missions demonstrated that the range profiles were distinct enough to discriminate between the ships, despite their nearly identical sizes. The measurements also validated the synthetic range profiles, which had been created under a variety of controlled sensor operating character-

istics and target scenarios.

Using synthetic range profiles, the multilayered neural network (MNN) and the Adaptive Clustering Network (ACN) achieved correct recognition of up to 100% for the classification of ships at 0°–27° (near-bow) viewing aspects. Using the synthetic passive-IR imagery, the ACN achieved correct recognition levels generally greater than 90% for 30°–57° and 60°–90° (near-beam) aspects. Thus our results, from both synthetic and measured data, indicate that range-profile and passive-IR signatures complement each other in covering all viewing aspects for long-range ship classification.

Acknowledgments

The laser radar range and passive-IR sensor models described in this article are based on the work of Jeffrey Shapiro, Associate Head of the MIT Electrical Engineering and Computer Science Department, and Thomas J. Green, Jr., staff member at Lincoln Laboratory. The authors are indebted to them for their contributions to this modeling effort.

The authors also thank Murali M. Menon for the use of the ACN software that he developed and implemented, and for many helpful discussions on the subject.

The authors acknowledge Steve Marcus for his proposed design of the noncooperative target identification sensor. He, along with Ted Quist, Dan Corbosiero, and Bob Hull, redesigned and reconfigured the Infrared Airborne Radar (IRAR) sensors for this mission. The authors are grateful to the plane crew and, in particular, to Ted Quist for handling the IRAR

sensor in a novel environment. The authors also thank the Newport Naval Base for allowing access to the ship targets.

Finally, the authors thank Vivian Titus, who helped develop and code the data-processing software. She and Lee Labrecque edited the raw data and created the database from the measurement exercises.

This work was sponsored by the Department of the Navy.

REFERENCES

1. S. Marcus, private communication.
2. S. Marcus, private communication.
3. J.H. Shapiro and P.L. Mesite, private communication.
4. J.H. Shapiro, R.W. Reinhold, and D. Park, "Performance Analyses for Peak-Detecting Laser Radars," *SPIE* 663, 38 (1986).
5. T.J. Green, Jr., "Three-Dimensional Object Recognition Using Laser Radar," Ph.D. Thesis, MIT Department of Electrical Engineering and Computer Science (1992).
6. J.H. Shapiro, private communication.
7. T.J. Green, Jr., and J.H. Shapiro, private communication.
8. S. Hannon, "Detection Processing for Multidimensional Laser Radars," Ph.D. Thesis, MIT Department of Electrical Engineering and Computer Science (1990).
9. M.M. Horst et al., private communication.
10. P. Dykstra and M. Muuss, "The BRL CAD Package: An Overview," in *Proc. Fourth USENIX Computer Graphics Workshop*, Cambridge MA (1987), pp. 73–80.
11. P.N. DiCaprio, private communication.
12. P.J. Kolodzy, private communication.
13. S.J. Rak and P.J. Kolodzy, "Performance of a Neural Network Based 3-D Object Recognition System," *SPIE* 1471, 177 (1991).
14. P.J. Kolodzy, "Multidimensional Automatic Target Recognition System Evaluation," *Linc. Lab. J.* 6, 117 (1993).
15. R. J. Hull, private communication.
16. D. G. Biron, private communication.
17. S.J. Rak and J.E. Baum, private communication.
18. R. Sharpe, ed., *Jane's Fighting Ships, 1991–1992* (Jane's Information Group, Coulsdon, UK, 1991).
19. R. J. Hull, private communication.
20. D.E. Rumelhart and J.L. McClelland, *Parallel Distributed Processing 1* (MIT Press, Cambridge, MA, 1986).
21. G.A. Carpenter and S. Grossberg, "Neural Dynamics of Category Learning and Recognition: Attention, Memory Consolidation, and Amnesia," in *Brain Structure, Learning, and Memory*, eds. J.L. Davis, R.W. Newburg, and E.J. Wegman, pp. 233–290, AAA Symposium Series (Westview Press, Boulder, CO, 1985).
22. M.M. Menon and E.J. Van Allen, "Automatic Design of Signal Processors Using Neural Networks," *SPIE* 1469, 322, (1991).
23. M.M. Menon, E.R. Boudreau, and P.J. Kolodzy, "An Automatic Ship Classification System for ISAR Imagery," *Linc. Lab. J.* 6, 289 (1993).

APPENDIX 1: SYNTHETIC LASER RADAR RANGE PROFILES

THIS APPENDIX contains a brief mathematical description of the model used to obtain the synthetic laser radar range profiles. A full derivation of the model is contained in Reference 1 of this appendix.

The range profile measured at the detector output is the magnitude of a complex envelope:

$$\hat{l}(n_z) = |\hat{\mathbf{l}}(n_z)|$$

where $\hat{\mathbf{l}}(n_z)$ represents the video-detected, range-sampled output from the sensor at range increment n_z . (Note: The $\hat{}$ symbol denotes the estimated value of a quantity, and the boldface type indicates a complex quantity.) The detector output comprises contributions from the reflected target signature and local-oscillator shot noise:

$$\hat{\mathbf{l}}(n_z) = |\hat{\mathbf{l}}(n_z)^{\text{target}} + \hat{\mathbf{l}}(n_z)^{\text{noise}}|, \quad (\text{A})$$

which are defined below. The detector output at each sensor instantaneous field of view (IFOV) can be calculated by convolving the transmitted laser pulse with the target backscatter coefficient $\beta(n_R)$. Assuming that all transmitted laser power is distributed uniformly within the sensor IFOV and is incident on a range-resolved target with uniform diffuse reflectivity, we can estimate the backscatter coefficient, or normalized target-range-extent signature weighting function, to be

$$\hat{\beta}(n_R) = \frac{\beta(n_R \Delta R)}{\beta(R_0)}.$$

The variable $\hat{\beta}(n_R)$ represents the fraction of the target that is illuminated at range n_R , normalized to full illumination at reference range $\beta(R_0)$, for a reflectance model sampled at ΔR intervals in range.

Next, laser speckle statistics are introduced as a product of $\beta^{1/2}(n_R)$ with a sequence of independent, identically distributed (IID) complex Gaussian ran-

dom variables $\{\mathbf{w}^{\text{speckle}}(n_R): n_R = 0, 1, 2, \dots, +\infty\}$, for which

$$\langle \mathbf{w}^{\text{speckle}}(n) \rangle = 0,$$

and

$$\langle \mathbf{w}^{\text{speckle}}(m) \cdot \mathbf{w}^{*\text{speckle}}(n) \rangle = \delta_{mn},$$

where the brackets $\langle \dots \rangle$ denote the expectation of a quantity, the superscript * denotes the complex conjugate, and

$$\delta_{mn} = \begin{cases} 1, & m = n \\ 0, & m \neq n \end{cases}.$$

This result is convolved with discrete versions of the transmitter pulse $s(t)$ and the matched filter impulse response $h(t)$, both with duration $T = 1/\epsilon^2$ pulse-width. The two pulses $s(t)$ and $h(t)$ are given by the following:

$$h(t) = s(t) = \left(\frac{8}{\pi T^2} \right)^{1/4} e^{-4t^2/T^2}.$$

Thus the reflected target signature can be written as

$$\hat{\mathbf{l}}(n_z)^{\text{target}} = \sqrt{\text{CNR}(R_0)} \times \sum_{n_R=0}^{+\infty} \sqrt{\hat{\beta}(n_R)} \frac{e^{-4(n_z-n_R)^2/N_z^2}}{\sqrt{N_z}} \mathbf{w}^{\text{speckle}}(n_R), \quad (\text{B})$$

where the range sampling increment N_z is equal to $cT/(\sqrt{2}\Delta R)$ with c = speed of light. The quantity $\text{CNR}(R_0)$ represents the design carrier-to-noise ratio of the sensor at a reference range R_0 and is expressed in terms of real system design parameters.

Equation B is the target contribution to the signal at the detector output. The contribution due to local-oscillator shot noise is modeled as an IID complex Gaussian random variable $\{\mathbf{w}^{\text{l.o.}}(n_t): n_t = -\infty, \dots, -2, -1, 0, 1, 2, \dots, +\infty\}$, for which

$$\langle \mathbf{w}^{\text{l.o.}}(n) \rangle = 0,$$

and

$$\langle \mathbf{w}^{\text{l.o.}}(m) \cdot \mathbf{w}^{*\text{l.o.}}(n) \rangle = \delta_{mn},$$

convolved with the discrete version of $h(t)$:

$$\hat{\mathbf{l}}(n_z)^{\text{noise}} = \sum_{n_t=-\infty}^{+\infty} \frac{e^{-4(n_z-n_t)^2/N_t^2}}{\sqrt{N_t}} \mathbf{w}^{\text{l.o.}}(n_t),$$

where the time sampling increment N_t is equal to $T/\Delta t$, and the discrete-time sampling interval Δt is equal to $2\Delta R/c$.

Finally, the output at the video detector is calculated by summing the target and local-oscillator shot-noise contributions and calculating the magnitude of the complex envelope, as given by Equation A. The closed-form expected value of the range signature, $\langle |\mathbf{l}(n_z)| \rangle$, which predicts *a priori* the result of averaging an infinite number of measured profiles by using the equations described above, is

$$\langle |\mathbf{l}(n_z)| \rangle = \sqrt{A + B}, \quad (\text{C})$$

where

$$A = \frac{\pi}{2} \text{CNR}(R_0) \sum_{n_R=0}^{+\infty} \hat{\beta}(n_R) \frac{e^{-8(n_z-n_R)^2/N_z^2}}{N_z},$$

and

$$B = \sqrt{\frac{\pi^3}{2^5}}.$$

Equation C predicts the mean range signature the sensor would measure by averaging over the effects of laser speckle and shot noise. Because Equation C is a function of the sensor's design, via $\text{CNR}(R_0)$, and the target characteristics, via $\beta(R)$, the equation is useful for understanding the fundamental characteristics and differences of the target signatures obtained with the design under consideration.

Reference

1. T.J. Green, Jr., and J.H. Shapiro, private communication.

APPENDIX 2: SYNTHETIC PASSIVE- INFRARED IMAGERY

TO CREATE the synthetic passive-infrared (IR) imagery, we selected the Gaussian-distributed passive-IR intensities of both the target and the background by using the method described in Reference 1 of this appendix. First, a pair of independent random numbers U_1 and U_2 were sampled from a uniform distribution. Next we converted these numbers to a pair of independent zero-mean unity-variance Gaussian-distributed numbers G_1 and G_2 by using the standard Box-Mueller method:

$$G_1 = \sqrt{-2 \ln(U_1)} \cos(2\pi U_2),$$

and

$$G_2 = \sqrt{-2 \ln(U_1)} \sin(2\pi U_2).$$

The Gaussian values were then converted to the target or background Gaussian distribution according to the following:

$$\hat{P}_t = \sigma_t G_1 + \mu_t,$$


and

$$\hat{P}_b = \sigma_b G_2 + \mu_b,$$

where (μ_t, σ_t) and (μ_b, σ_b) represent the mean and standard deviation of the Gaussian thermal distributions for the target and background, respectively. We repeated this process of calculating target and background passive-IR pixel intensities for the entire image template. Further details of the processing are contained in Reference 1.

Reference


1. S. Hannon, "Detection Processing for Multidimensional Laser Radars," Ph.D. Thesis, MIT Department of Electrical Engineering and Computer Science (1990).



JERROLD BAUM is an associate staff member in the Opto-Radar Systems Group. His research has included battlefield modeling and simulation (studying how to reduce fratricide) and automatic target-detection algorithms that fuse laser radar and passive-infrared imagery. Jerry received an A.B. degree in physics, cum laude, from Brandeis University and an M.S. in physics from the University of Maryland. Before joining Lincoln Laboratory five years ago, he worked at Textron Defense Systems. Previously, he taught high school physics for eight years.



ELEANOR TUNG received a B.S. degree in chemistry from the National Taiwan University and a Ph.D. degree in theoretical chemistry from the Catholic University of America. She worked for Ford Aerospace and Communications (now the Loral Corp.) before joining Lincoln Laboratory 12 years ago. She is currently a staff member in the Opto-Radar Systems Group, where her focus of research has been on automatic target classification and identification with active and passive-infrared imaging systems. Eleanor is a member of Sigma Xi. Her hobbies include gardening and hiking.



STEVEN J. RAK is an associate staff member in the Opto-Radar Systems Group. He received both a B.S. and an M.S. degree in chemical engineering from Case Western Reserve University. In 1986 he joined Lincoln Laboratory, where his research has included the investigation of neural network techniques for biologically inspired image processing and for classifying laser radar imagery, the simulation of sensor performance for creating synthetic data and imagery, and the visualization of three-dimensional imagery for enhanced image understanding. Before joining Lincoln Laboratory, he was an engineering analyst at Westvaco's Laurel Research Center. Steve enjoys few things more than simply bicycling and running on the roads and trails of New England.