ADAPTIVE MULTI-SENSOR INTEGRATION FOR MINE DETECTION*

James E. Baker
Oak Ridge National Laboratory
Robotics and Process Systems Division
P.O. Box 2008, Bldg. 7601, MS-6305
Oak Ridge, TN 37831-6305

RECEIVED
MAY 16 1997
OSTI

"The submitted manuscript has been authored by a contractor of the U.S. Government under contract No. DE-AC05-96OR22464. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes."

MASTER

To be published in the Proceedings of SPIE's AeroSense '97 Conference: Detection and Remediation Technologies for Mines and Minelike Targets II, Orlando, Florida, April 20-25, 1997

*Research sponsored by the U.S. Department of Energy, under contract DE-AC05-96OR22464 with Lockheed Martin Energy Research Corp.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

ln

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

Adaptive multi-sensor integration for mine detection*

James E. Baker

Oak Ridge National Laboratory, Robotics and Process Systems Division
P.O. Box 2008, Bldg. 7601, MS-6305, Oak Ridge, TN 37831-6305

ABSTRACT

State-of-the-art in multi-sensor integration (MSI) application involves extensive research and development time to understand and characterize the application domain; to determine and define the appropriate sensor suite; to analyze, characterize, and calibrate the individual sensor systems; to recognize and accommodate the various sensor interactions; and to develop and optimize robust merging code. Much of this process can benefit from adaptive learning, i.e., an output-based system can take raw sensor data and desired merged results as input and adaptively develop/determine an effective method of interpretation and merger. This approach significantly reduces the time required to apply MSI to a given application, while increasing the quality of the final result and provides a quantitative measure for comparing competing MSI techniques and sensor suites.

The ability to automatically develop and optimize MSI techniques for new sensor suites and operating environments makes this approach well suited to the detection of mines and mine-like targets. Perhaps more than any other, this application domain is characterized by diverse, innovative, and dynamic sensor suites, whose nature and interactions are not yet well established. This paper presents such an outcome-based multi-image analysis system. An empirical evaluation of its performance and its application, sensor and domain robustness is presented.

Keywords: multi-sensor integration, sensor fusion, adaptive learning, genetic algorithms, outcome-based processing

1. INTRODUCTION

1.1. Mine detection domain

The reliable detection of mines and mine-like targets has become of increasing research interest in recent years. This is partly due to the vast size of the problem, i.e., it is currently estimated that over 11,000,000 acres of land already contain unexploded ordnance and between 500,000 and 1,000,000 new mines are deployed around the world each year. Unfortunately, this application domain is characterized by well concealed targets (generally buried or submerged), prevalent interference (false targets, data-corrupting material, varying environmental conditions, etc.), and an understandable need for highly accurate detection (low false negative—due to the danger of overlooking explosives, and low false positive—due to high eradication costs). The very purpose of mine deployment requires them to be largely hidden from view, making many of the best understood sensor modalities less than optimally applicable (e.g., CCD, IR, SONAR), and has spawned many innovative sensor technologies to be investigated. Such sensors focus on detecting target features which are not fully occluded and integral properties with high spatial resolution (e.g., high sensitivity IR, explosive element enzyme dispersal and detection, focused chemical “sniffers”, ground conductivity sensing, ground penetrating radar). In each case, the sensor must be robust to the full range of anticipated targets, be highly accurate, and not induce detection/detonation by the ordnance. It seems unlikely that any single sensor will be able to achieve the desired detection proficiency, though an efficacious integration of a diverse and/or dynamic sensor suite may. Unfortunately, traditional Multi-Sensor Integration (MSI) application techniques require detailed knowledge of the sensors, the sensor interdependencies, the desired targets, and the target/application domain. The research presented in this paper is intended to help satisfy this requirement by presenting an MSI approach which requires minimal a priori knowledge of the sensors and the target domain.

1.2. Multi-sensor integration

Although MSI has a generally accepted overall purpose, i.e., to generate a more reliable and consistent representation, the *specific* goals of each MSI implementation vary from one application to another.^{1,2,3} MSI techniques combine multi-modal sensor images (e.g., sonar and visual,⁴ distance and reflectance^{5,6}), multiple single modal sensor images (e.g., multi-sampling, a single sensor's images displaced in time or in space), or combinations of these.

*Research sponsored by the U.S. Department of Energy, under contract DE-AC05-96OR22464 with Lockheed Martin Energy Research Corp.

MSI is necessary because of basic ambiguities inherent in our current sensor imaging technologies. Sensor ambiguities derive from two basic causes: limitations of the physical attribute being measured by the sensor (e.g., visible light's inability to permeate opaque surfaces, sonar edge effects, etc.) and the sensor's inaccuracies in making measurements (e.g., noise, resolution). Ambiguity exists as long as the mapping from reality to image is not 1-to-1. That is, if different "realities" lead to identical images, a single image cannot reveal the particular reality which was the truth, e.g., a 2-D visual image of an opaque object cannot reveal its interior or its hidden surfaces and hence, an infinite number of "realities" would result in the same sensor image.

Choosing the most advantageous MSI technique is dependent on the given application and available sensors. Multi-sampling is perhaps the simplest MSI technique and is primarily used to reduce noise. Merging multiple single modal images⁷ displaced in space can improve 3-dimensional information, reveal otherwise occluded areas, and reduce orientation induced artifacts (e.g., glare, sonar positional anomalies). Merging multiple single modal images displaced in time can be used to reduce temporal effects or to isolate them. Employing multi-modal sensors permits exploitation of each sensor's strengths without suffering their intrinsic weaknesses, e.g., using both sonar and vision can provide both accurate distance (sonar) and edge detection (vision).⁴ Clearly the choice of sensors and MSI techniques is critical to achieving performance gains from MSI; MSI provides nothing if the images to be merged are not complementary.

While the goal of MSI is to improve the representation's reliability and consistency, the means to achieve this goal can be divided into three categories⁶ based on the relative information content of the original images with that of the desired representation.

In the first case, "detail enhancement," the relative information content of the original images is less rich than the desired representation. This case occurs whenever the original images are integrations of the target (e.g., raw NMR, X-ray photography) or undergo a distorting function, e.g., unfocused lenses, interference patterns, etc. These images can be translated into a more accurate and detailed representation if the original integration or distortion function is sufficiently known to permit an inverse mapping function to be well approximated.⁸

In the second case, "data enhancement," the MSI techniques are concerned with improving the accuracy of the data rather than either increasing or decreasing the level of detail. Techniques within this category include noise reduction, resolution enhancement and reduction of sensor artifacts.

In the third case, "conceptual enhancement," the image contains more detail than is desired, making it difficult to easily recognize objects of interest. In these images one can group together pixels corresponding to the same conceptual object and thereby reduce the level of extraneous detail. This task may require significant amounts of global knowledge and processing time. For example, if one wishes to distinguish "walls" from "floors," one could separate objects on planar edges, however this would subdivide complex structures into their many facets. Collecting these facets into single multi-planar objects without combining separate, adjacent objects, requires detailed knowledge of each complex structure of interest (object definition) and recognition of that structure within the image (object recognition). This problem extends to all sensor modalities as well (e.g., color, texture, composition, shape). Either one accepts the shortcomings inherent in making clustering decisions independent of global knowledge, or one accepts the time and computational complexity associated with object definition and recognition code. This research concerns real-time applications and, hence, restricts its focus to optimizing the performance ability of rapidly implementable approaches.

Rapid implementation is not only useful but indispensable for many applications.^{9,10} Unfortunately, many image processing algorithms require significant processing time.^{1,2,11,13} This is especially true of feature extraction, object isolation, and object recognition algorithms due to their typical reliance on global or large neighborhood information. For example, many object isolation algorithms are based on expected object templates, feature extraction, and "hypothetical" feature extension, i.e., extension of edges to possible intersection points in order to form closed polygons from only partial edge line segments. These methods have the advantage of global knowledge, large neighborhood features, a priori expectations, etc., but are, for the same reasons, rather slow and domain or target specific.

State-of-the-art image digitizers¹² are capable of performing complex functions at the pixel level at full, image acquisition rates. This permits one to obtain pixel level features with little processing delay time and may permit MSI images to be produced at near original image acquisition rates, leading to virtual "multi-modal sensors." The methods described herein exploit this speed capability by developing MSI algorithms based on pixel-level features. Hence, the basic direction of this approach to conceptual enhancement is the potentially faster and more robust formation of clusters from pixels rather than the slower process of segmenting images into clusters. (Note that while this process and the resulting representation is generally called "segmentation" in the literature,^{8,13,14,15} we will use the term "clustering" to reinforce this distinction in the basic direction of the approaches.)

This pixel level approach to conceptual enhancement assumes that the sensors have sufficient resolution, relative to the objects, to ensure that adjacent pixels corresponding to the same object are similar, i.e., in zero, first, or second order

derivatives, in at least one of the sensor modalities and that the objects of interest span several pixels. If these conditions are not met, either the objects are out of range or the sensors' modalities are inappropriate for the task.

It is a great temptation to develop and evaluate image processing (IP) algorithms on simulated data with the rationalization that the results can be easily adapted to real images. While this may well be true theoretically, often the problems and difficulties inherent in processing real sensor data are such that simulation based techniques resolve few of the real problems (e.g., the distortions, sensor anomalies, noise) inherent in real images. These problems can so dominate/modify the processing problem that the simulation based technique is either solving such a trivial component or the anticipated mapping function from real images to simulation quality images is at least as complex as the original IP problem. It is our view, therefore, that real sensor data must be the focus from day one for all but function encoding verification, graphical explanation, or preliminary concept generation. The techniques presented in this paper are developed and evaluated on actual multi-modal sensor data consistent with real-world applications, i.e., a laser range camera used for robotic navigation and ground conductivity data used for buried waste assessment.

The following section describes the traditional approach to MSI and our adaptive learning approach is detailed in Sect. 3. Section 4 describes the sensory target domains, including a description of the physical nature of each sensor system (laser range camera, LRC, and ground conductivity, GC, sensor) and their targeting environments. The fifth section presents an empirical evaluation of the performance of our automated MSI system applied to LRC images. Section 6 empirically validates these results in the learned (LRC) and unlearned (GC) domains. A summary section follows.

2. TRADITIONAL MSI

Current state-of-the-art in MSI application^{3,11} involves a series of sequential, time-consuming steps which often only highly qualified and experienced researchers in sensors and image processing can conduct. As the knowledge and experience gained in one step is often critical to successfully performing the next, a single researcher or team must generally perform most of the work. This substantially constrains the critical path of the schedule, i.e., in general, since the steps cannot be performed in parallel, the length of the overall schedule will not be shortened by the addition of researchers to the development team, but may be seriously lengthened by their departure. These basic steps are:

1. *Understand and characterize the application domain.* The conditions within the application domain which must be considered include ambient conditions (e.g., temperature, background radiation, lighting, humidity, air quality), targeting conditions (e.g., specific characteristics of the target, false targets, obstructions, background objects) and anomalous effects (e.g., glare, shadows, changing air conditions). Any conditions in the operating environment which will affect the choice of sensors, their performance, and the image interpretation techniques, must be sufficiently understood to direct selection and exploitation of the sensor suite. It is important to note that specific conditions affecting sensor operations is a function of the sensors themselves. Therefore, this step cannot be executed without a detailed familiarity with the candidate sensors, their physical operation, and corresponding interpretation techniques. This step interdependency and required sensor and image processing expertise necessitates the high qualifications required of the MSI researchers.

2. *Determine and define the appropriate sensor suite.* This step requires an in-depth knowledge of the various available sensors (i.e., their advantages, disadvantages, and complementary behaviors with respect to the application domain and task requirements) and modifications which potential manufacturers could implement. Sensor capabilities must be weighed against both the operating domain and available potential image processing techniques. Furthermore, secondary considerations (e.g., weight, kinetic shock survivability, radiation hardening, signal-to-noise ratios, long-term product support, power requirements, field-of-view) must also be understood and considered when determining the best sensor suite for a given application.

3. *Analyze, characterize, and calibrate the sensor suite.* An empirical evaluation of the sensor(s) must be performed within the designated operating domain, and must include realistic targets, obstacles, and all expected ambient conditions. The purposes of this step are to validate the sensor selection, drive image interpretation software development, and understand any sensor anomalous behavior. Such behavior includes: sensor accuracy as a function of measurement discontinuity (e.g., color quality at boundaries); impact of theoretically unrelated modalities (e.g., range accuracy as a function of surface reflectance); scene dynamics (e.g., moving targets and/or moving obstacles); active sensor "cross talk" (e.g., acceptable proximity and frame rate of SONAR ranging sensors); ambient noise (e.g., presence of ambient light source which corresponds to a laser camera's operating frequency). If a sensor proves to be unacceptable, one may need to return to step one for both the single erroneous sensor and potentially the entire sensor suite. The entire suite may be impacted, since in some cases, individual sensors are selected for their behavioral compatibility and/or complementary modalities. Hence, the unexpected replacement of an unacceptable sensor may impact the utility of the others.

4. *Recognize and accommodate sensor interactions.* This step involves two distinct types of sensor interactions. In the first type, the individual accuracy of the data is affected, e.g., "crosstalk" interactions affecting the interpretation of data correspondences. As mentioned in step 3, "crosstalk" can occur whenever two or more active sensors confuse each others

signal for their own. This is a common potential problem for SONAR range sensors and is usually redressed by controlled firing. Clearly this solution changes the effective data stream rate. Other interactions occur because of incompatibilities between the physical measurement properties of the sensors. For example, the performance of a magnetic transducer based positional system or magnetic conductivity sensor can be adversely impacted by the near proximity of a highly conductive object (e.g., another sensor system¹⁶). Further, since some sensors are based on physically unreliable phenomenon, multiple, physical characteristics of each point must be determined in order to accurately represent the scene. For example, in some range cameras, the accuracy of the range data is a function of the reflective quality of the surface being measured. Hence, by coupling a range sensor with a reflectance sensor, the information is available to correct the range image.

The second type of sensor interaction which must be understood and addressed is the accurate interpretation of data correspondences. Determining which pixels in multiple distinct images of the same or similar scene correspond to each other is a very difficult problem and has yet to be fully solved.^{17,18,19} The problem, called image or data registration, is complicated by the possibilities of occlusions, dynamic changes in the scene, variations in sensor quality as a function of position or time, etc., which may prevent *any* correctly corresponding pixel to exist. The image registration problem is not directly a subject of this research, but must be addressed if spatially or temporally (for dynamic scenes) disjoint images are to be merged. Generally, one concentrates on registering images exhibiting only trivial changes or one determines the precise relative spatial positions of the sensors so that simple geometric mapping can be used. However, even if these conditions are met, the actual integration can be extremely complicated by the association and resolution problems.^{17,18,19} That is, given two distinctly different pixels that are geometrically determined to lie in very close 4-dimensional proximity (space and time): (1) should they be merged (and by what function); (2) is one anomalous or less reliable (e.g., due to differences in sensor quality, lighting, target sheen, timing); or (3) do they represent distinct, albeit spatially proximate, objects? In most cases, the problem is even worse since 4-dimensional placement is unknown, e.g., CCD images effectively yield vector data not point data.

5. *Develop and optimize robust merging code.* Based on the results from the above four steps and the MSI goal (e.g., improved human display, rapid evaluation of vast sensor data, target detection, object characterization, target tracking, self location) specific merging software must be designed, implemented, tested, and empirically tuned. Because of the variety of potential MSI goals, the resulting merging code varies tremendously from one application to another. Developing this code involves issues similar to those for image registration described above. If two perfectly registered and corresponding pixels have different values for the same physical characteristic: (1) is one erroneous and to be ignored; (2) are both erroneous and can any information be garnered from them; (3) if they are to be merged, by what function—simple average, weighted average, the minimum, the maximum, or some complex weighted function of secondary variables (e.g., temperature, air quality, timing, lens focal length). The number of merging possibilities grows substantially if the two pixels represent different physical properties, e.g., distance and reflectance.

Once completed, such carefully orchestrated and expertly crafted MSI systems are expensive, fragile, and sub-optimal. Any change (e.g., to the sensor suite, environmental domain, task specification, target description) may require a complete reanalysis. This problem is exacerbated by the typically multiple man-years necessary to complete the cycle for real world applications, since replacing a researcher lost to natural attrition can often necessitate a significantly time-consuming learning curve. Similarly, since most state-of-the-art MSI implementations are performed on proof-of-principle demonstration testbeds, the entire development cycle will probably need to be reperformed, even if the initial development was successful (i.e., since prototyping the hardware system will likely result in sensor system changes due to obsoleted models, new weight or power constraints).

Another critical deficiency in this approach is that due to the development costs of even a single MSI "solution," the only empirical quality comparison generally available is against the original, non-MSI system. Hence, traditionally developed MSI systems can undergo little more than a minor tuning from their original design parameters. Since it is financially and time-wise infeasible to fully develop even two MSI systems to compare competing sensor suites, initial design decisions must be accepted by faith unless they prove unworkable; optimizing for them is cost prohibitive.

3. ADAPTIVE LEARNING APPROACH TO MSI

Much of the MSI development process described in the previous subsection could benefit from adaptive learning.⁵ That is, an Automated MSI Solution Generator (AMSG) could be designed which would take actual sensor data and the desired merged results as input during a training stage in order to adaptively develop/determine an effective method of merging the sensor data. An AMSG would significantly reduce the time required to apply MSI to a given problem, while increasing the quality of the final result and provide an objective analysis/comparison of competing MSI techniques and sensor suites.

While an AMSG could be developed for a variety of MSI tasks, we will restrict our consideration to conceptual enhancement, wherein the original image contains more detail than is desired, making it difficult to easily recognize objects of interest. This research focuses on merging pixels into "homogeneous" clusters and leaving the unification of conceptually related clusters for some more abstract, post-processing phase.

Our approach to conceptual enhancement is based on the premise that given appropriate sensors (i.e., those capable of distinguishing the significant conceptual regions, or *facets*) and sufficient resolution (i.e., providing multiple pixels of each such facet) clusters of pixels can be formed which correspond to the facets by accurately answering the fundamental question for each pair of adjacent pixels: "Do these pixels belong to the same facet?" Given an accurate response for each pair of adjacent pixels, forming accurate clusters is trivial. This question applies to cluster formation regardless of the sensor or application domain. The answer, in the form of a confidence value, provides a universal and uniform interface—applicable to any sensor system and MSI technique to be integrated.

This fundamental question suggests simply comparing the adjacent values across each sensor modality. But when posing the same question in the negative "Is there a surface edge lying between these two pixels?" one is drawn to compare the pixel neighborhoods on each side of the questioned interface (e.g., comparing the linear extrapolation of pixels on the first side of the interface with the second pixel's value, comparing the average adjacent pixel value variation on each side of the interface). The AMSG is a system which searches for a function which best answers this question for each sensor in the suite simultaneously, and optimizes those functions in accordance with the user-defined, desired results.

During a learning stage, the AMSG is provided with a set of sensor data and the corresponding, desired MSI results. The system then progresses through a basic feedback learning cycle, see Fig. 1: (1) apply candidate MSI algorithm to sensor data, (2) compare results with desired results and quantify performance quality, and (3) use quality measure to determine next MSI algorithm to evaluate. This cycle proceeds until some user-specified stopping criterion is met, e.g., minimum performance quality, given number of iterations, etc. When the user is satisfied with the performance quality of the MSI algorithm generated, that MSI algorithm can be directly applied to incoming sensor data.

The AMSG consists of four primary components, see Fig. 1: the Search Space, the Implementation Code, the Learning Strategy, and the Evaluation Function. The following subsections will detail each of these four components.

Fig. 1. Basic schematic of AMSG data flow.

3.1. Search space

While this is not a "component" in the strictest sense, the specification of the search space, more than anything else, defines the potential quality of the AMSG. The search space must be composed of algorithms sufficiently specialized to provide high quality MSI performance, while sufficiently robust to permit working latitude in the fielded system. Furthermore, the search space must be sufficiently broad to include solutions applicable to a wide variety of domains, while sufficiently small to permit reasonable search, i.e., preferably no more than perhaps 10^{20} candidate algorithms. (Note that each set of parameter values defines a distinct MSI candidate solution.)

There are competing interests driving the search space design. First, the space must be searchable, i.e., there must be sufficient regularity for an automated global search technique to be applicable. Second, grossly dissimilar sensors must be accommodated, while third, all sensor inputs must be transformed into a single format for merging. Fourth, the MSI algorithm must be very robust and general purpose to permit maximum domain application, while fifth, being highly specialized to ensure quality solutions. And sixth, the search space design must be extendible to yet unforeseen sensor types and maximally diverse sensor suites.

To these ends, we have chosen a hierarchical organization of highly parameterized functions. Each sensor's input is transformed into an edge confidence map for subsequent confidence combination merger. The edge confidence map was chosen as a universal/fundamental interface which is based on perhaps the most reliable and basic property of sensors—the ability to detect discontinuities in some physical phenomenon. (That is, while absolute measurements may be erroneous and the detectors may not even be measuring the physical property intended by their designers, significant changes in output can be reasonably inferred to correspond to some significant change in some physical property in the scene. Hence, at the very least, an edge confidence map can be generated.) Edge confidence maps have the added advantage of being sufficiently universal that almost any source (e.g., CAD model, intelligence reports) can be translated into one, and thereby merged by this system.

In our approach, each sensor to be merged is analyzed separately, according to its “affinity function.” This real-valued, parameterized function provides a “degree of match” for each pair of adjacent pixels within that sensor's image. This operation translates the image into an edge confidence map. These maps are then merged across the sensor suite by combining the confidences for each pixel interface. This merging operation results in a single edge confidence map for the entire sensor suite, see Fig. 2. A fully segmented image can then be produced by threshold-based region growing. Hence, a single **MSI algorithm** is defined by: (1) the pixel interface statistics used by each sensor's affinity function; (2) the affinity function and corresponding parameter values for each sensor in the suite; (3) the merge function and corresponding parameter values which define how to combine confidences across the sensor suite; and (4) the region growing algorithm and corresponding parameter values which indicate cohesion or disjunction for a given pixel interface.

Fig. 2. Schematic of sensory data flow through execution of an AMSG's MSI algorithm.

Hence, the **search space** consists of:

- 1) *statistical definitions* – A set of low-level pixel-interface statistics (e.g., change in pixel value, change in slope of adjacent pixel values).
- 2) *affinity function definitions* – A set of affinity functions, each of which is a parameterized function of various pixel-interface statistics.
- 3) *merge function definitions* – A set of merge functions, each of which is a parameterized function of affinity values.
- 4) *region growing algorithm definitions* – A set of region growing algorithms, each of which is a parameterized function of merge function values.
- 5) *parameter/threshold values* – The possible ranges and precisions of parameter values for each of the above functions/algorithms.

To be included in the above statistic and function sets, a definition need only be considered potentially useful in distinguishing facet interface boundaries in at least *some* sensor modality. Specific definitions can be found in my related publication²⁰.

This search space design has several significant advantages:

- 1) It is applicable to most sensor systems, sensor modalities, and application domains, by simply augmenting/modifying the definition sets or parameter spaces.
- 2) Its design is applicable to spatial or temporal images; a pixel interface need not be considered a uniform spatial quantity.
- 3) Its hierarchy permits arbitrarily complex definitions at each stage, e.g., an image edge finding algorithm or a misregistered, inaccurate world map could supply affinity values. The AMSG's architecture permits the results from diverse systems to be readily merged using standard confidence combination approaches.
- 4) Its reliance on bottom-up image processing, i.e., forming clusters from pixels rather than segmenting images into clusters, permits extremely robust statistic and function definitions, applicable to a wide variety of sensor systems and modalities.
- 5) Changing the sensor suite has little impact on the AMSG, requiring at most augmenting, not changing, the definition sets.
- 6) State-of-the-art image digitizers¹² are capable of performing complex functions at the pixel level at full, image acquisition rates. This may permit one to execute a given MSI algorithm with little processing delay time, leading to virtual "multi-modal sensors."
- 7) It transforms the MSI development problem into a global search optimization problem. This permits the system to leverage existing technology in the area of global optimization.
- 8) By enabling an automated global search mechanism to "find" an appropriate MSI algorithm, new application and sensor domains can be targeted with minimal development costs and increased speed and quality. This, in turn, should permit a better and more objective analysis of the synergistic capabilities of various sensor suites while reducing the need for the developer to recognize and understand the "best" sensor interactions to be exploited for a given task. This advantage makes the AMSG approach ideal for innovative sensor modalities whose detailed character and multi-sensor synergies have yet to be established, e.g., as are beginning to emerge within the remote mine detection field.

3.2. Implementation code

The Implementation Code is responsible for applying an MSI algorithm to the incoming sensor data at near real-time rates. Once an appropriate MSI algorithm is generated for a given application domain, this component will constitute the entire fielded system.

Application of a given MSI solution is performed in four separate stages, see Fig. 2:

- 1) *collect statistics* – Each sensor image to be merged is processed separately and meaningful low-level statistics (e.g., change in pixel value, change in slope of adjacent pixel values) are collected for each pixel interface in accordance with required affinity functions.
- 2) *process each sensor* – For each sensor image, its "affinity function" is applied to each pixel interface, e.g., for a two dimensional image grid, both "horizontal" and "vertical" interfaces would be processed. The resulting numerical edge confidence for each pixel interface is stored for the subsequent stage. This value indicates the "degree of confidence" that the two pixels sharing the interface belong to the same facet according to this sensor image.
- 3) *merge sensors* – A single edge confidence value is determined for each pixel interface by applying a merge function to the corresponding affinity values across the entire sensor suite. The process is repeated for each pixel interface and results in a single (merged) edge confidence map.
- 4) *form clustered image* – A region growing algorithm is employed to combine pixels whose interface "edge confidence value" is below some parameterized threshold.

3.3. Learning strategy

The learning strategy must select an MSI algorithm from the candidate solutions within the search space according to the quantified performances of prior selections. This learning strategy must efficiently exploit the feedback so that a sufficiently high performing solution can be found in a reasonable amount of time.

Genetic Algorithms, GAs, are well suited to this global optimization task.^{21,22} GAs are extremely powerful adaptive global search techniques derived from natural population genetics,^{23,24,25,26} and have been shown to perform well for many types of functions including those exhibiting very difficult characteristics, e.g., discontinuities, non-differentiability, multi-modality,

high dimensionality, huge search spaces and noise. GAs require no specific, a priori function information; only the form of a candidate solution, i.e., the number of parameters to be optimized and the desired level of precision (number of bits), and a comparative performance measure for candidate solutions is needed.

GAs are a simplified simulation of the natural genetic model. As such, GAs simulate a population of individuals evolving over multiple generations: individuals are specified by a series of genes (bits) that can be independently inherited; reproduction is accomplished by a crossover operation which forms offspring from the genetic material of their parents; and, an individual's reproduction frequency is based on his performance in the environment (evaluation function). Thus, trait encoding, sexual reproduction and "survival of the fittest" propagation are all simulated; and the average performance of the population tends to improve over successive generations. When some user-specified stopping criterion is met, the best individual produced is taken as the GA's solution for function optimization.

Clearly, GAs can not guarantee discovery of the optimal solution, but they have proven themselves to be powerful global search techniques capable of simultaneously searching extensive regions of the parameter space. The basic GA framework used in this research was provided by the GENESIS^{27,28} GA package with modifications suggested by Baker²⁹.

3.4. Evaluation function

The evaluation function compares the MSI results of the candidate algorithm with the user-provided, desired results and quantifies the algorithm's quality. This quantification can be used as an objective comparison of MSI algorithms. To efficiently guide the AMSG, this function should have a high resolution and monotonically encourage MSI quality, i.e., it should be such that even small improvements in the MSI algorithm will be reflected in the evaluation measure and have a positive impact on the learning process. More detail on the evaluation function itself can be found in related publications^{6,20}.

4. TARGET SENSOR DOMAINS

Two imaging domains are used to test and validate the AMSG. The first is a laser range camera targeting an indoor "warehouse" environment, typical of many robotic systems' operating environments. This domain will be used for training the AMSG and validating its application robustness. The second imaging domain is a ground conductivity sensor targeting a buried waste field. This domain was chosen specifically because of its highly dissimilar modality and data format. That is, while the laser range camera provides two 2-dimensional images of the surface characteristics targeted, the ground conductivity sensor provides two 0-dimensional data points of the integrated character of the targeted region and is highly applicable to the mine detection domain.

This section presents a description of the physical nature of each of these sensor systems and their targeting environment.

4.1. Ground conductivity sensor

The available subsurface data consists of a set of ground conductivity (GC) readings³⁰ of a cold test pit at the Idaho National Engineering and Environmental Laboratory (INEEL) waste storage site.³¹ This set was originally to include ground penetrating radar data, but the ground water and clay soil of that region rendered that sensor modality impotent. Despite that loss, MSI can still be applied by merging the multiple GC data sets and the two measured signals inherent in GC data: quadrature and in-phase strength.

The GC sensor used consists of a magnetic transmitting coil and a receiving coil placed 3.66 meters apart.³⁰ The transmitting coil sets up a magnetic field in the ground which induces eddy currents and in turn a secondary magnetic field, 90° out of phase. The receiving coil measures both the primary magnetic field (in-phase component) and the secondary field (quadrature component). Thus, a GC sensor measures the integrated dielectric constant of the ground in the 3-dimensional proximity of the two coils. By its very nature, data from a GC sensor is very ambiguous and unfocused since a buried conductive object will disturb all readings taken anywhere within its vicinity and the amount of disturbance is a function of: the object's size, dielectric constant, orientation with respect to the coils, distance from each magnetic pole, uniformity and distribution of ambient material, the relative position of other magnetic conductive objects, etc.. Hence, from a set of GC data, it is impossible to determine which of the infinite number of perfectly data-consistent possibilities actually corresponds to "reality."

Given GC data's ambiguous and unfocused nature, it would appear a perfect candidate for inverse mapping using MSI "detail enhancement" techniques. However, the inverse mapping function is unresolvable from the GC data alone, due to its inherent ambiguities. Ground penetrating radar was an excellent complementary sensor and if available may have permitted the inverse mapping function to be roughly approximated.

The GC data's representation may be improved by using MSI "clustering" techniques. However, since precise truth for the GC data is not available, these data sets cannot be used for designing or training MSI "clustering" algorithms. This

necessitated the use of a second sensor domain, the laser range camera, for the development of general-purpose, MSI clustering algorithms which could eventually be applied to the GC sensor domain. Even so, the GC data is still well suited for the validation of the AMSG's domain robustness; sensor-type independence; and applicability to integrated and unstructured data points—like many consistent with mine detection (e.g., sniffers, Geiger counters, ground penetrating radar).

4.2. Laser range camera

The Odetics laser range camera (LRC) used for this research³² produces images of 128x128 pixels where each pixel's value is determined by the reflection properties of a directed laser. This LRC measures two values: the reflected light's phase shift, indicating the distance to the target; and its intensity, indicating the target surface's sheen or degree of reflectivity at the camera's operating wavelength of 820 nm. This LRC provides a 60°x60° field of view with an unambiguous measurement range for targets lying between 3 and 10 meters distant. The target chosen for analysis was a portion of our laboratory,³³ with distances ranging from 3 to 15 meters and scattered, miscellaneous objects (e.g., furniture, boxes, 55 gal. drums). For this domain, MSI could be performed using either multiple LRC images or the multi-modal, distance and reflectance images from a single view.

The LRC was chosen for MSI algorithm development for four reasons: (1) it is consistent with the requirements of many autonomous and teleoperated robotics' environments, for which we have various sponsors; (2) since its two, multi-modal images (distance and reflectance) are obtained from the same reflected laser light, they are perfectly registered images, thus eliminating the dependency on separate image registration algorithms; (3) a calibrated LRC was readily available in our laboratory and fully integrated with our computer network; and (4) precise truth measurements could be taken and used for AMSG design, development, training, and evaluation.

5. EMPIRICAL EVALUATION

The evaluation function, see Sect. 3.4, provides a numerical measure of the quality of a conceptually enhanced image, i.e., how well it compares with the desired result. This measure is useful not only for feedback to guide the GA's global search, but also for the objective comparison of competing MSI techniques. However, it should be noted that while the GA is a powerful search technique, it does not guarantee any particular proficiency relative to that technique's global optimum, e.g., a guarantee of achieving X% of the technique's optimum performance. For example, in the empirical evaluation presented in this paper, each search was permitted the same amount of computing resources yet the search space sizes ranged from 2^{10} to 2^{47} . Hence, the competing techniques had significantly different proportions of their search spaces investigated. In each experiment, the GA was permitted to examine 5000 MSI algorithms. For the smaller search spaces, this was sufficient to cover the entire space, but for the largest spaces, the GA could examine less than 4 trillionth of the search space. Although the evaluation measure can be used to compare the quality of the resulting images, it can only imply the corresponding technique's relative potential.

For each experiment, the GA was given an initially random population of 50 candidate solutions (MSI algorithms). The GA was permitted to search until it had evaluated 5000 MSI algorithms or its population had stagnated (failed to produce a distinct individual for two successive generations), whichever came first. As described in Sect. 3, the MSI algorithm to be evaluated was applied to the sensor data and the resulting segmented image compared with the user-defined desired results. An evaluation of the relative quality and effectiveness of the various MSI approaches can be found in related publications^{6,20}.

The raw data, 3-D plotted using the range values and colored according to their reflectance values is shown in Fig. 3a. The user-supplied, truth image is shown in Fig. 3b. Its features include: (1) floor; (2) a box with two facets facing the camera; (3) a movable wall partition; (4) two facets of the back wall, divided by a supporting pillar; (5) a tool chest; (6) a segment of the left wall; (7) a small crate; (8) a supporting pillar in the middle of the back wall; and (9) a triple faceted, corner support structure.

For human display, adjacent pixels corresponding to the same facet are graphically connected. This display mode makes it easier to discern the facets and their edges. The color of the facets is based on their relative size, in pixels. The black space between facets is the result of: (1) missing data due to occlusion, e.g., behind the box in the middle of the room; (2) the display mode not graphically connecting adjacent pixels which correspond to different facets, e.g., separating the two faces of the box and separating the wall supports; or (3) erroneous data labeled as "don't care pixels," e.g., the random clutter next to Facet #7. (Note that the "don't care pixels" were not excluded from processing; the MSI techniques had to deal with them. However, the correct clustering of these pixels was not evaluated since we wished to optimize the MSI techniques for their ability to ignore erroneous data, not for their ability to correctly cluster both valid and erroneous data.) In the following figures showing MSI results, erroneous data is shown with its corresponding clusters, however, to "declutter" the display, only the larger clusters are shown.

Fig. 3a and b. LRC data, 3-D plotted according to range data: a) raw data points colored by reflectance data; b) truth image (user-defined, desired MSI results) colored by size of facet and numbered for reference.

Figure 4 shows the MSI results of the best performing MSI algorithm found during the learning stage. This image reveals the major facets (#1 through #6, above) and some of the minor ones, including two of the corner supports (#9). The top surface of the tool chest and a region of clutter in the upper right corner of the image were also correctly differentiated. These facets were considered too small and their data too erroneous to be included in the truth image.

Fig. 4. Clustering results of automatically generated MSI algorithm applied to data of Fig. 3a and trained by corresponding "Truth Image" (Fig. 3b), 3-D plot colored by size of cluster.

Although this algorithm failed to differentiate all of the known detail in the image, it was very successful. To get an impression of the quality of these results, one need only review the raw data shown in Fig. 3a. In that image, most of the facets, e.g., #3+, are difficult to visually discern despite our detailed knowledge of them.

It should be noted that the uniform emphasis on pixels rather than any special emphasis on finding each of the primary 14 truth facets explains this technique's "failure" to isolate the smaller facets. However, to do otherwise would require global and/or a priori knowledge which would demean the general applicability and robustness of the AMSG. The promise of the AMSG is the quality of its results despite its total lack of global, conceptual shape, or a priori knowledge—a common requirement in many real-world applications. These MSI results are of sufficient quality to enhance an automated world map, a teleoperated display, or a higher-level object recognition system.

6. VALIDATION

The AMSG is highly domain and application robust. Domain robustness refers to its suitability for a wide variety of environments and sensor types, once it is tuned to the level of image detail desired by the user. Application robustness refers to the suitability of a given optimized MSI algorithm to a wide variety of situations within a tuned or learned domain. This section validates the AMSG's application robustness by evaluating the performance of its optimized MSI algorithm in the learned domain (i.e., LRCs operating in an indoor environment) and validates its domain robustness by demonstrating its suitability to another, highly dissimilar domain (i.e., GC images of a waste test site).

6.1. Learned domain, LRC

An MSI algorithm and its control parameters were chosen from an extensive search space, (see Sect. 3) based on a single LRC image set (Fig. 3a) and user defined truth image (Fig. 3b). The following validation test is to demonstrate the suitability of that MSI algorithm and its optimized parameters to LRC images of other scenes taken within the same domain.

Figure 5a shows the LRC's reflectance image of another corner of our laboratory. For clarity of description, the primary facets of interest are numbered: (1) floor; (2) a cylindrical barrel; (3) a moveable wall partition; (4) two facets of the back wall, separated by a supporting pillar; (5) a 55-gallon drum; (6) a door set in the right wall; (7) a box with two facets in view; (8) a supporting pillar in the middle of the back wall; (9) and (10) two halves of a windowed door; (11) a suspended hoist and its draped extension cords; (12) three more facets of the back wall separated by #11; (13) a box with one facet in view; (14) a fire extinguisher; and (15) an electrical cable lying on the floor. This image partially overlaps the previous image; facets numbered 1, 4, and 8 match the correspondingly numbered facets of Fig. 3b.

The image in Fig. 5a is more difficult than the learned image of Fig. 3a. This image contains cylindrical objects (Facets #2 and #5), very narrow objects (Facet #15 and parts of Facet #11), and significantly more detail. The best performing MSI technique from Sect. 5 and its tuned parameter set were applied to the distance and reflectance image corresponding to Fig. 5a.

Figure 5b shows the MSI results of this MSI technique using the GA's optimized parameter settings. For image clarity, only the larger clusters are displayed. This technique resolved most of the facets enumerated in Fig. 5a, including those corresponding to both the cylindrical and narrow objects. (Note that cylinders are difficult because the various affinity functions²⁰ defined in the search space (see Sect 3.1) are based on *linear* extrapolation and the parameters were tuned on an image consisting entirely of planes.) Furthermore, three facets in the far right corner were unexpectedly differentiated: the door facing, parallel to Facet #6; the doorway offset, parallel to Facet #12; and a narrow strip of the right wall, again parallel to Facet #6. One can also discern the sign and doorknob of Facet #9, the wire connecting the two primary components of Facet #11, and the speaker and electrical outlet near the top of Facet #8. Unfortunately, Facet #8 itself was not distinguished from Facet #4 and the extreme detail of the hoist and its cables (Facet #11) prevented it from being amalgamated into a single, cohesive object, though this technique did correctly isolate it from the other objects. (Note that the support pillar, Facet #8, was not distinguished in the training image either, see Fig. 4, and hence was not expected to be resolved here.)

Figure 5b clearly demonstrates the application robustness of the MSI algorithm tuned for this domain. After training the technique to just one application image, it was able to isolate almost all of the significant facets in an a priori unknown image. These facets included surface types hitherto never seen: curved surfaces. The success of this MSI technique on curved surfaces, despite its purely linear extrapolation approach and complete lack of experience with this surface type, strongly supports the utility of using local pixel properties to form clusters.

6.2. Unlearned domain, GC

The MSI algorithms described in Sect. 3 are designed to be independent of the sensor modalities being merged and the application being addressed. These are general purpose MSI techniques for determining edge confidences from multiple sensor data sets and merging the results into a single, conceptually enhanced representation. The desired level of image detail, and hence, the intrinsic definition of facets, must be defined by the user according to his application and purpose. This definition must be communicated to the system in order for it to tune the parameters for a given application domain. For the LRC, the truth image implicitly embodied the user's definition of facets and of extraneous detail. From this image, the system could infer the relative information content of the sensor modalities being merged, and hence, optimize the various tuning parameters. If a "truth" definition is not supplied, the user must either rely on general purpose parameter settings (tuned to a diverse suite of domain types), or "manually" adjust the parameters to achieve the desired level of detail.

Fig. 5a and b. New scene from Fig. 3's domain, i.e., LRC data of another corner of our laboratory, with new objects interspersed: a) 2-D plot of raw data colored by reflectance; b) clustering results of MSI algorithm using parameters tuned to Fig. 3, 3-D plot colored by size of cluster and numbered for reference.

The unlearned domain chosen to demonstrate the domain robustness of the AMSG is Ground Conductivity, GC, images (see Sect. 4.1) of a buried waste test site. This sensor is an excellent test case because of its extreme dissimilarity with the LRC: it measures internal rather than surface properties; it obtains an integration measure rather than a point measure; it has a wide signal dispersion, leading to sampling overlap; its truth is almost indeterminable, due in part to the effects of varying environmental conditions (e.g., ground water, soil type) across the region being mapped; it is a single rather than an array sensor, and hence its sample positioning must be determined externally; etc. Since accurate truth for this GC data is unknown, the system can not be automatically trained as it was for the LRC images. Furthermore, developing general purpose parameter settings will require images and truth from many diverse sensor modalities and was not the intention of this research. Hence, we must demonstrate the robustness of this MSI approach applied to GC data by examining the results of manually selected MSI algorithms and parameter settings. The purpose here is not to provide a definitive MSI solution for the GC domain, but rather to demonstrate that the AMSG is applicable and that given a "truth image," the system could be trained on this domain as it was on the LRC domain in Sect. 5.

The raw GC data consists of hundreds of samples taken over a three day period. These samples consist of an X-Y position (based on dead-reckoning from a given starting point) and two sensor readings: the quadrature component and the in-phase component (see Sect. 4.1). The data is plotted in Fig. 6. Because of the unequal spacing of the data points, the MSI techniques described in Sect. 3 can not be *directly* applied. The data must first be transformed into a grid representation. A grid resolution and registration is chosen such that all of the internal grid elements are represented in the original data set. The value of each grid element is then defined as the average value of the raw sample points lying within its corresponding area, see Fig. 7. (Note that the white elements on the lower edges have no corresponding data samples and that the extreme edges of the original data set were discarded, see Fig. 6.)

Figure 8a presents the MSI results of one technique from the search space described in Section 3.1 with a manually selected parameter set. With this MSI algorithm, the main cluster (near the middle of the image) is isolated, as are smaller clusters just below it and in the upper right corner of the image. This result is intuitively consistent with the data in Fig. 7, based on casual, visual inspection. However, without defined truth (e.g., the relative importance of the two sensor images, an indication of the desired level of detail) other solutions are equally viable. Figure 8b presents the MSI results of another MSI technique with a manually selected parameter set. In this image, the smaller cluster in the upper right corner of the image is considered insignificant, while the main cluster is examined more critically. Again, neither image of Fig. 8 is offered as a *solution*. Rather, without "ground truth," they represent different, yet equally valid interpretations of the data.

Fig. 6a and b. Raw GC data of INEEL waste storage site, superimposition of three data sets in a 2-D plot colored by: a) quadrature; b) in-phase measure.

Fig. 7a and b. Raw GC data of Fig. 6, represented in a 2-D uniform grid colored by: a) quadrature; b) in-phase measure.

Fig. 8a and b. MSI clustering results applied to data of Fig. 7 using manually selected parameter values, 2-D plot colored by size of cluster: a) "4-pt, OR" MSI algorithm; b) "6-pt, AND."

Despite the lack of a truth image to which the approach could be trained and its results evaluated, the two results presented in Fig. 8 clearly indicate the ability to isolate significant facets and yield significantly different conceptual interpretations. As such, these figures demonstrate the suitability of the AMSG to the GC domain and provide an indication of its overall domain robustness.

7. SUMMARY

The Automated MSI Solution Generator (AMSG), presented in this paper, is designed to take actual sensor data and a user's definition of the desired merged results as input during a training phase and independently find an efficacious method of merging the sensor data. While the basic AMSG approach could be applied to a variety of MSI tasks, we chose to restrict our consideration to the conceptual enhancement problem. Conceptual enhancement is an MSI technique used when the original image contains more detail than is desired, making it difficult to easily recognize objects of interest. This problem can be resolved by segmenting the image into meaningful clusters of pixels which correspond to conceptually homogeneous regions. The AMSG's underlying MSI approach is based on the premise that given appropriate sensors, clusters of pixels can be quickly and accurately formed by evaluating the continuity between each pair of adjacent pixels based on the local pixel neighborhoods. Hence, the basic direction of this approach is the potentially faster and more robust formation of clusters from pixels rather than segmenting images into clusters. This approach is based on low level functions which capture fundamental characteristics of continuity and, as such, are applicable to cluster formation regardless of the sensor or application domain. This approach constitutes a robust, automated, MSI methodology.

The AMSG employs a hierarchy of potentially useful definitions for solving conceptual enhancement tasks, from the bottom up. These definitions form a pool of low-level MSI tools which can be automatically selected and optimized for a given sensor suite and application domain. In our implementation, this adaptive search for superior MSI solutions is performed by GA global optimization. Thus, the overall system is domain and application robust and provides a universal interface for diverse MSI tools and techniques to be merged, through standard confidence combination mechanisms.

The AMSG is validated on real distance and reflectance images from a laser range camera (LRC) by its ability to develop a high quality MSI solution for the merger of this sensor's images in the defined application domain. This MSI solution was able to isolate nearly all of the significant facets of the training image. The domain robustness of this optimized solution was validated by applying it to an unknown LRC image set within the same application domain, containing new types of objects. Nearly all of the significant facets of the new image were correctly clustered. The application robustness of the

AMSG was validated by applying the techniques to an extremely dissimilar sensor and application domain—ground conductivity data of buried waste.⁶

The quality of its results, its general applicability, and its automated nature make the AMSG approach of MSI a valuable resource for emerging applications characterized by divergent sensor suites including new and unrigorously characterized sensor modalities. Hence, mine detection appears to be a perfect application domain for this technology.

Future research includes: time optimizing the MSI solution on a state-of-the art digitizer¹² to produce a virtual “multi-modal sensor”; extending the sets of definitions being investigated; permitting heterogeneous resolution sensors to be easily merged; improving the conceptual quality of the evaluation function by the addition of facet-based terms; and extending the GA's search to include the actual formation of affinity functions for consideration.

REFERENCES

1. M. A. Abidi and R. C. Gonzalez, eds., *Data Fusion in Robotics and Machine Intelligence*, Academic Press, New York, 1992.
2. J. M. Brady, Foreword, Special Issue on Sensor Data Fusion, *Int. J. Robotics Research*, 7(5):2–4, 1988.
3. A. Kak and S. Chen, eds., “Spatial Reasoning and Multi-Sensor Fusion, *Proceedings of the 1987 Workshop*, Morgan Kaufmann Publishers, Inc., San Mateo, California, 1987.
4. M. Beckerman and D. L. Barnett, “Performance of Visual and Ultrasound Sensing by an Autonomous Robot,” ORNL/TM-11733, Martin Marietta Energy Systems, Inc., Oak Ridge Natl. Lab., January 1991.
5. J. E. Baker, “Adaptive Learning of Multi-Sensor Integration Techniques with Genetic Algorithms,” *Evolutionary Algorithms in Robotics II, Proceedings of the 5th International Symposium on Robotics and Manufacturing*, 383–390, August 1994.
6. J. E. Baker, “Image Accuracy and Representational Enhancement Through Low-Level, Multi-Sensor Integration Techniques,” *Proceedings of Sensor Fusion and Aerospace Applications*, SPIE Proc. 1956, 170–186, April 1993.
7. J. P. Jones, “Real-Time construction of Three-Dimensional Occupancy Maps,” *Proceedings of the IEEE International Conference on Robotics and Automation*, 52–57, Atlanta, Georgia, May 1993.
8. R. C. Gonzaliz and P. Wintz, *Digital Image Processing*, Advanced Book Program/World Science Division, Addison-Wesley Publishing Company, Reading, Massachusetts, 1983.
9. M. Jamshidi, et. al., ed., *Proceedings of ISRAM '96, the Sixth International Symposium on Robotics and Manufacturing*, Montpellier, France, TSI Press, Albuquerque, New Mexico, May 1996.
10. M. Jamshidi, et. al., ed., *Proceedings of WAC '94, the First World Automation Congress*, Maui, Hawaii, TSI Press, Albuquerque, New Mexico (August 1994).
11. I. Kadar and V. Libby, eds., *Proceedings of Signal Processing Sensor Fusion, and Target Recognition V*, SPIE Proc. No. 2755, June 1996.
12. Datacube, Inc., 300 Rosewood Drive, Danvers, Massachusetts.
13. E. R. Dougherty, *An Introduction to Morphological Image Processing*, Vol. TT9, SPIE Optical Engineering Press, Bellingham, Washington, 1992.
14. C. R. Giardina and E. R. Dougherty, *Morphological Methods in Image and Signal Processing*, Prentice-Hall, Englewood Cliffs, New Jersey, 1988.
15. R. Kasturi and R. C. Jain, *Computer Vision: Principles*, IEEE Computer Society Press, Los Alamitos, California, 1991.
16. *Flock of Birds(TM): Six Degrees-of-Freedom Measurement Device*, technical description paper, Ascension Technology Corporation, P.O. Box 527, Burlington, Vermont, 1996.
17. *Proceedings of the 1996 Conference on Computer Vision and Pattern Recognition (CVPR '96)*, IEEE Cat. No. PRO7258, San Francisco, California, June 1996.
18. *Proceedings of the International Conference on Image Processing (ICIP '96)*, IEEE Cat. No. 96CH35919, Lausanne, Switzerland, September 1996.
19. *Proceedings of the 13th International Conference on Pattern Recognition (ICPR '96)*, Vienna, Austria (August 1996).
20. J. E. Baker, “Automated Generation of Multi-Sensor Integration Solutions for Real-Time Applications,” invited chapter for *Event-Based Sensing, Planning and Control of a Robotic System: An Integrated Approach*, co-ed. T. J. Tarn, Bijoy K. Ghosh and Ning Xi, CRC, submitted February 1997.
21. L. J. Eshelman, ed., *Proceedings of the Sixth International Conference on Genetic Algorithms*, Morgan Kaufmann Publishers, Inc., San Francisco, California, 1995.
22. S. Forrest, ed., *Proceedings of the Fifth International Conference on Genetic Algorithms*, Morgan Kaufmann Publ., Inc., San Mateo, California, 1993.
23. J. E. Baker, “An Analysis of the Effects of Selection in Genetic Algorithms,” Ph.D. thesis, Vanderbilt University, Nashville, Tennessee, May 1989.

24. K. A. De Jong, Analysis of the Behavior of a Class of Genetic Adaptive Systems, Ph.D. dissertation., University of Michigan, Ann Arbor, Michigan, 1970.
25. D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, Massachusetts, 1989.
26. J. H. Holland, *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, Michigan, 1975.
27. J. J. Grefenstette, *A User's Guide to GENESIS*, Version 5.0, direct distribution: gref@aic.nrl.navy.mil, October 1990.
28. J. J. Grefenstette, *Optimization of Control Parameters for Genetic Algorithms*, IEEE Transactions of Systems, Man, and Cybernetics, SMC-16(1), 122-128, 1985.
29. J. E. Baker, "Correcting Stochastic Sampling in Genetic Algorithms," to appear in a special issue of *Intelligent Automation and Soft Computing*, 1997.
30. M. Beckerman, "Phase I Report on Sensor Technology, Data Fusion and Data Interpretation for Site Characterization," ORNL Report No. CESAR-91/39, Martin Marietta Energy Systems, Inc., Oak Ridge Natl. Lab., 1991.
31. J. E. Baker, "Image Accuracy and Representational Enhancement Through Low-Level, Multi-Sensor Integration Techniques," ORNL/TM-12218, Martin Marietta Energy Systems, Inc., Oak Ridge Natl. Lab., May 1993.
32. *3-D Laser Imaging System User's Guide*, Odetics Publ No. 8533031-PR, Odetics Inc., 1515 S. Manchester Avenue, Anaheim, CA, March 29, 1990.
33. Center for Engineering Systems Advanced Research, Oak Ridge National Laboratory, Lockheed Martin Energy Research Corp., Oak Ridge, Tennessee.