



3 4456 0374959 4

ornl

ORNL/TM-12218

**OAK RIDGE
NATIONAL
LABORATORY**

MARTIN MARIETTA

Image Accuracy and Representational Enhancement Through Low-Level, Multi-sensor Integration Techniques

OAK RIDGE NATIONAL LABORATORY
CENTRAL RESEARCH LIBRARY
CIRCULATION SECTION
1500N ROOM 175

LIBRARY LOAN COPY

DO NOT TRANSFER TO ANOTHER PERSON

If you wish someone else to see this
report, send in name with report and
the library will arrange a loan.

UCR-789 (8-77)

J. E. Baker

MANAGED BY
MARTIN MARIETTA ENERGY SYSTEMS, INC.
FOR THE UNITED STATES
DEPARTMENT OF ENERGY

This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from the Office of Scientific and Technical Information, P.O. Box 62, Oak Ridge, TN 37831; prices available from (615) 576-8401, FTS 626-8401.

Available to the public from the National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161.

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, makes 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.

ORNL/TM-12218

Engineering Physics and Mathematics Division

**IMAGE ACCURACY AND REPRESENTATIONAL
ENHANCEMENT THROUGH LOW-LEVEL,
MULTI-SENSOR INTEGRATION TECHNIQUES**

J. E. Baker

DATE PUBLISHED — May 1993

Research sponsored by the
Office of Technology Development
and by the
Engineering Research Program
Office of Basic Energy Sciences
U.S. Department of Energy

Prepared by the
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee 37831
managed by
MARTIN MARIETTA ENERGY SYSTEMS, INC.
for the
U.S. DEPARTMENT OF ENERGY
under contract DE-AC05-84OR21400



3 4456 0374959 4

CONTENTS

ABSTRACT	vii
1. INTRODUCTION	1
1.1 MSI OBJECTIVES	1
1.2 TARGET DOMAINS	2
1.3 OVERVIEW	3
2. LOCAL MSI METHODOLOGY	5
2.1 DATA ENHANCEMENT ALGORITHMS	5
2.1.1 Resolution Enhancement	5
2.1.2 Sensor Anomaly Correction	6
2.2 CONCEPTUAL ENHANCEMENT ALGORITHMS	9
3. MSI OPTIMIZATION	13
3.1 LEARNING STRATEGY	13
3.2 EVALUATION FUNCTION	13
3.3 EMPIRICAL RESULTS	15
4. VALIDATION	23
4.1 LEARNED DOMAIN, LASER RANGE CAMERA	23
4.2 UNLEARNED DOMAIN, GROUND CONDUCTIVITY	26
5. SUMMARY AND CONCLUSIONS	31
ACKNOWLEDGMENTS	33
REFERENCES	35

LIST OF FIGURES

<u>Fig.</u>		<u>Page</u>
1A	Unenhanced LRC data of a box on the floor, with corner facing camera, 3-D plot colored by reflectance	7
1B	Resolution enhanced LRC data of a box on the floor, with corner facing camera, 3-D plot colored by reflectance	7
2A	Raw LRC data of a corner of the laboratory, with various objects interspersed, 2-D plot colored by distance	8
2B	Raw LRC data of a corner of the laboratory, with various objects interspersed, 2-D plot colored by reflectance	8
3	LRC data of Fig. 2, 3-D plot colored by reflectance	10
4	Wrap-around corrected LRC data of Fig. 3, 3-D plot colored by reflectance	10
5	Cluster matching against truth cluster T . Cluster B associated with T , penalty is a function of A' and B''	15
6	Comparison of penalty factors for the best solutions found for each method, both for individual sensors and with MSI	16
7	Comparison of eight MSI techniques, showing combination function trade-off (AND, OR) with respect to match function strictness (2-pt and 4-pt G versus 4-pt P and 6-pt)	18
8	Truth image of Fig. 4, 3-D plot arbitrarily colored by cluster and numbered for reference	18
9	“Dynamic 6-pt OR” clustering of Fig. 4, 3-D plot arbitrarily colored by cluster	20
10	“6-pt OR” clustering of Fig. 4, 3-D plot arbitrarily colored by cluster	20
11A	Raw LRC data of another corner of the laboratory, with new objects interspersed, 2-D plot colored by distance	24
11B	Raw LRC data of another corner of the laboratory, with new objects interspersed, 2-D plot colored by distance	24
12	“Dynamic 6-pt OR” clustering of image from Fig. 11 using parameters tuned to Fig. 4, 3-D plot arbitrarily colored by cluster and numbered for reference	25

LIST OF FIGURES (cont'd)

<u>Fig.</u>		<u>Page</u>
13A	Raw GC data of INEL test pit, superimposition of three data sets in a 2-D plot colored by quadrature	27
13B	Raw GC data of INEL test pit, superimposition of three data sets in a 2-D plot colored by in-phase measure	27
14A	Raw GC data of Fig. 13A, represented in a 2-D uniform grid colored by quadrature	28
14B	Raw GC data of Fig. 13B, represented in a 2-D uniform grid colored by in-phase measure	28
15	“4-pt OR” clustering of Fig. 14 using hand-tuned parameters, 2-D grid plot arbitrarily colored by cluster	29
16	“6-pt AND” clustering of Fig. 14 using hand-tuned parameters, 2-D grid plot arbitrarily colored by cluster	29

ABSTRACT

Multi-Sensor Integration (MSI) is the combining of data and information from more than one source in order to generate a more reliable and consistent representation of the environment. The need for MSI derives largely from basic ambiguities inherent in our current sensor imaging technologies. These ambiguities exist 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. MSI techniques attempt to resolve some of these ambiguities by appropriately coupling complementary images to eliminate possible inverse mappings. What constitutes the best MSI technique is dependent on the given application domain, available sensors, and task requirements.

MSI techniques can be divided into three categories based on the relative information content of the original images with that of the desired representation: 1) "detail enhancement," wherein the relative information content of the original images is less rich than the desired representation; 2) "data enhancement," wherein the MSI techniques are concerned with improving the accuracy of the data rather than either increasing or decreasing the level of detail; and 3) "conceptual enhancement," wherein the image contains more detail than is desired, making it difficult to easily recognize objects of interest. In conceptual enhancement one must group pixels corresponding to the same conceptual object and thereby reduce the level of extraneous detail.

This research focuses on data and conceptual enhancement algorithms. To be useful in many real-world applications, e.g., autonomous or teleoperated robotics, real-time feedback is critical. Unfortunately, many MSI/image processing algorithms require significant processing time. This is especially true of feature extraction, object isolation, and object recognition algorithms due to their typical reliance on global or large neighborhood information. This research attempts to exploit the speed currently available in state-of-the-art digitizers and highly parallel processing systems by developing MSI algorithms based on pixel rather than global-level features. Hence, the basic direction of this research is the potentially faster and more robust formation of "clusters from pixels" rather than the slower process of extracting "clusters from images."

The "data enhancement" algorithms developed include general-purpose, sensitivity enhancement techniques and specific sensor anomaly detection and remediation techniques. These techniques are evaluated on actual multi-modal sensor data obtained from a laser range camera, i.e., range and reflectance images. A suite of over thirty conceptual enhancement techniques are developed, evaluated and compared on this sensor domain. For objective analysis, each algorithm's performance is optimized by Genetic Algorithms, a machine learning technique, for a predefined level of detail. However, the general applicability of these techniques is empirically verified by applying the laser range camera optimized results to different sensor domains, e.g., ground conductivity data including quadrature and in-phase strength images. The overall result is a general-purpose, MSI conceptual enhancement approach which can be efficiently implemented and used to supply input to a variety of high-level processes, including: object recognition, path planning, and object avoidance systems.

1. INTRODUCTION

Multi-Sensor Integration(MSI) has been defined as combining data and information from more than one source to generate a more reliable and consistent representation of the environment.¹ MSI techniques combine: multi-modal sensor images (e.g., sonar and visual, range and reflectance, etc.), multiple single modal sensor images (e.g., multi-sampling, images displaced in time, or images displaced in space), or various combination of these.²

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, etc.). 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 a 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 images displaced in space can improve 3-dimensional information, reveal otherwise occluded areas, and reduce relative position induced artifacts, e.g., glare, sonar anomalies, etc. 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 while avoiding their intrinsic weaknesses, e.g., using both sonar and vision can provide both accurate distance (sonar) and edge detection (vision).³ 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.

1.1 MSI OBJECTIVES

While the goal of MSI, by definition, is "to generate a more reliable and consistent representation," the means to achieve this goal can be divided into three categories 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, etc., 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 understood to permit an inverse mapping function to be well approximated.

2 INTRODUCTION

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/regions of interest. In these images one must group together pixels corresponding to the same conceptual entity 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 clusters on planar edges, however this would subdivide complex structures into their many facets. Collecting these facets into single multi-planar clusters without combining separate, adjacent clusters, 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, etc. Either one accepts the shortcomings inherent in making segmentation/clustering decisions independent of global knowledge, or one accepts the time and computational complexity associated with object definition and recognition.

1.2 TARGET DOMAINS

Two imaging domains are investigated in this paper: subsurface data of a buried waste site and sensor data typical of underground storage tanks or warehouse environments.

The available subsurface data consists of a set of ground conductivity (GC) readings of a cold test pit at the INEL waste storage facility. 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 provide some benefit by merging the multiple GC data sets and by merging the two measured signals inherent in GC data: quadrature and in-phase strength.

The GC sensor used¹ consists of a magnetic transmitting coil and a receiving coil placed 3.66 meters apart. 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 "clustering" algorithms. This forced us to investigate a second sensor domain, the laser range camera (LRC), for the development of general-purpose, clustering algorithms which could eventually be applied to the GC sensor domain, see Section 4.2.

Laser imaging is performed by scanning a laser beam in a two-dimensional pattern and measuring the returned (reflected) signal to determine the distance to the target at various points and its relative reflectance ability, or gloss. LRCs determine the target distance by measuring the time delay between broadcasting a signal and receiving its reflection. The LRC used in this research [4], produces images of 128×128 pixels corresponding to its $60^\circ \times 60^\circ$ field of view and has a measurement range of 3 to 10 meters. The target chosen for analysis was a portion of ORNL's Center for Engineering Systems Advanced Research (CESAR) laboratory, with distances ranging from 3 to 15 meters and scattered, miscellaneous objects, e.g., furniture, boxes, 55 gallon drums, etc. 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, including underground storage tank characterization and automated warehouse tasks; 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 need for and effects of 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 of the target scene could be taken and used for MSI design, development, tuning, and evaluation.

1.3 OVERVIEW

The second section of this paper outlines our basic MSI approach and its application to Data Enhancement and Clustering techniques. Section 3 outlines a general purpose parameter optimization method and its objective comparison of the competing MSI algorithms. The fourth section presents an empirical validation of these techniques on new images from the learned domain (LRC) and from a completely unfamiliar sensor domain (GC).

2. LOCAL MSI METHODOLOGY

This research focuses on MSI “data enhancement” and “conceptual enhancement” algorithms. To be useful in many applications, e.g., autonomous or teleoperated robotics, real-time feedback is critical. Unfortunately, many image processing algorithms require significant processing time. 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 specialized.

State-of-the-art image digitizers⁵ 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 rate, leading to virtual “multi-modal sensors.” This research attempts to exploit this speed capability by developing MSI algorithms based on pixel-level features. Hence, the basic direction of this research is the potentially faster and more robust formation of “clusters from pixels” rather than the slower process of extracting “clusters from images.”

This pixel level approach to clustering assumes that the sensors are well chosen and have sufficient resolution, relative to the conceptual entities, to ensure that adjacent pixels corresponding to the same entity are similar (e.g., in zero, first, or second order derivatives, in at least one of the sensor modalities) and that the entities of interest span several pixels. If these conditions are not met, either the entities are out of range or the sensors’ modalities are inappropriate for the task.

2.1 DATA ENHANCEMENT ALGORITHMS

2.1.1 Resolution Enhancement

Optimally, a sensor’s resolution is sufficiently fine to prevent its intrinsically discrete nature from impacting its use. This was not the case for the LRC. The LRC has a distance resolution of 1.44 inches⁴ and a vertical and horizontal displacement of approximately $\text{range} \cdot \sin(60^\circ/128^\circ) = 1.5$ inches at about its mid-range of 15 feet. This resolution level creates discretization artifacts, making constant slopes appear as step functions. Use of a simple smoothing function would reduce this effect but would also decrease significant detail and sharpness.

The solution chosen was to develop a discriminating smoothing operator which would be applied only if adjacent pixels indicated either the pixel in question was erroneous (noisy) or suffering from the coarse resolution constraint. The four pairs of pixels adjacent to and on opposite sides of the target pixel are chosen. The linear interpolation of each pair at the target pixel’s position is calculated and the

6 LOCAL MSI METHODOLOGY

common consensus values averaged. An erroneous value was defined as one without a neighboring pixel-pair interpolation having similar value. (Clearly, this assumes all features of interest are greater than one pixel in size.) If noise was present, the target pixel's value would be changed to the average linear interpolation defined above. If not, the target pixel's value would be changed in the direction of the average linear interpolation, but by no more than one half the sensor's resolution capability. This restraint is based on the assumption that if the target pixel is not noisy, then its value is the most accurate possible for the sensor's resolution. To change it by more than half that resolution would contradict this assumption.

Figure 1 shows a close-up of a carton pointed toward the camera and colored by distance. Figure 1A is the raw sensor data and Fig. 1B is the resolution enhanced data. Since the target consists of two perpendicular planes at approximately the same angle relative to the camera's line of sight, the readings should be smooth and nearly symmetric. This process improved the resolution and corrected much of the resolution (discretization) distortions. It is important to note that by resolution enhancement we do not mean an increase in the number of sample points, but rather an increase in the existing samples' measurement accuracy and precision.

This resolution enhancement technique is applicable to any type of image and is used as a preprocess to reduce spot noise and resolution-induced sensor anomalies. More sophisticated slope determinations could be performed on larger neighborhoods or subsequently resolved entities, however, such methods would be significantly slower and/or post processing methods.

2.1.2 Sensor Anomaly Correction

The LRC determines distance by comparing the relative phase shift between the original light transmitted and its reflection. This method is unambiguous only for targets within one phase length, since it is impossible to determine the integer number of phase lengths traversed. Thus, targets at 0.1, 1.1 and 2.1 phase length would all be detected as 0.1 out of phase. This problem is known as "wrap-around error"; the LRC assumes all reflections occur within one phase length of the transmitter. Figure 2 displays the raw data from the LRC. (The random, dark spotting on this and subsequent figures is a result of the photographic process and is not present in the sensor data itself.) Figure 2a presents the distance image shown as a 2-dimensional grid with darker color indicating shorter distance from the sensor. Figure 2b presents the reflectance image with darker color indicating greater reflectivity (sheen). In Fig. 2a, the dark region in the upper left corner actually corresponds to the most distant targets and is suffering from wrap-around errors. The fact that wrap-around errors do not effect the reflectance image, Fig. 2b, permits MSI techniques to resolve this anomaly.

Figure 3 presents the reflectance data plotted in 3-dimensions by using the distance data of Fig. 2a. Note that Fig. 3's point of view is offset from the original LRC's view point to amplify the 3-dimensional effect. Areas apparently missing data are a result of target occlusion. The wrap-around error in the distance readings results in a smearing of data points toward the camera's view position.

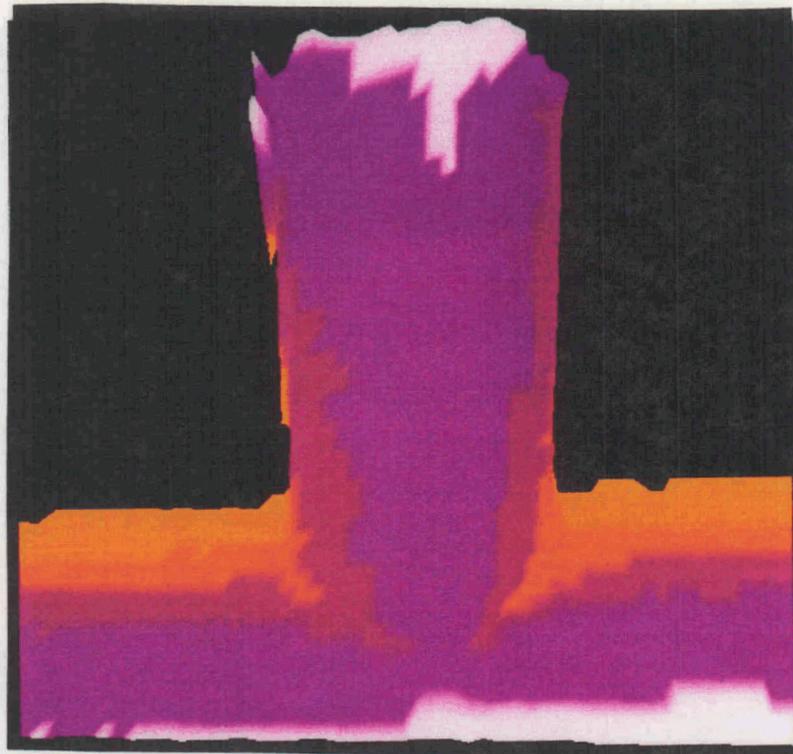


Fig. 1A. Unenhanced LRC data of a box on the floor, with corner facing camera, 3-D plot colored by reflectance.

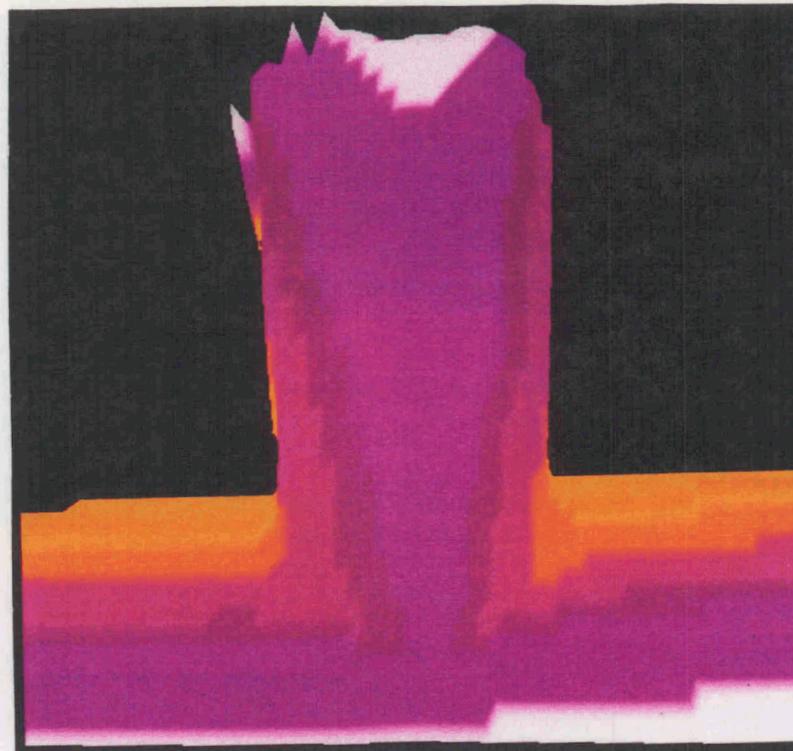


Fig. 1B. Resolution enhanced LRC data of a box on the floor, with corner facing camera, 3-D plot colored by reflectance.

8 LOCAL MSI METHODOLOGY

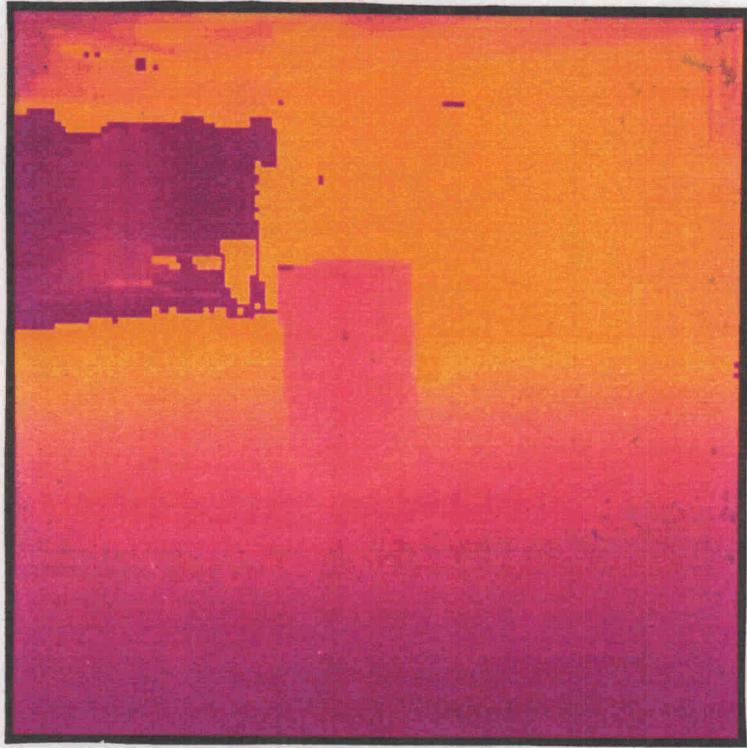


Fig. 2A. Raw LRC data of a corner of the laboratory, with various objects interspersed, 2-D plot colored by distance.

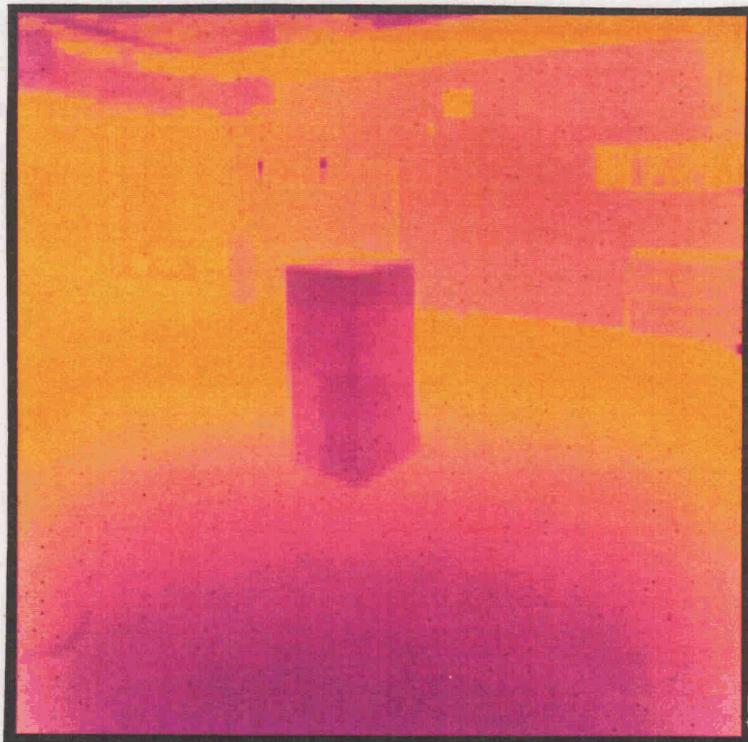


Fig. 2B. Raw LRC data of a corner of the laboratory, with various objects interspersed, 2-D plot colored by reflectance.

The wrap-around error was addressed by comparing adjacent pixels in both the distance and reflectance images. A wrap-around error is indicated if the difference in their distance values is very large (approaching the full phase length) while the difference in their reflectance is small (indicating little change). Whenever a wrap around error is indicated, the closer pixel's distance is incriminated by one phase length. The results of this MSI technique are shown in Fig. 4. Note that this solution cannot guarantee accurately correcting all wrap-around errors. For example, if the error is coincident with a sudden change in reflectance values, one must conclude that the pixel's short distance may in fact be accurate. While this occurred for several pixels in the given image, most of the wrap-around error was resolved correctly.

2.2 CONCEPTUAL ENHANCEMENT ALGORITHMS

The MSI conceptual enhancement algorithms presented in this paper are based on the premise that given appropriate sensors (i.e., those capable of "seeing" the targets) and sufficient resolution (i.e., providing multiple pixels of each target) clusters of pixels can be formed which correspond to the targets by accurately answering the fundamental question for each pair of adjacent pixels: "Do these pixels belong to the same target?" Given an accurate response for each pair of pixels, forming accurate clusters is trivial. This question applies for cluster formation regardless of the sensor domain, whether one is isolating shapes in CCD images, forms in LRCs, or conductivity anomalies in GC data.

This fundamental question suggests simply comparing the adjacent values across each sensor modality. But by 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, or comparing the average adjacent pixel value variation on each side of the interface. Such functions should be compared across all modalities and appropriately merged.

In this research, each modality is analyzed separately, and a real-valued, "degree of match" for each pair of adjacent pixels is determined based on the user specified "match function." These "degrees of match" are then merged across both modalities by a separate "combination function." Four different match functions are compared. These differ in the size and use of pixel neighborhoods: 1) "2-pt," which uses only the two pixels adjacent to the interface; 2) "4-pt General," which uses the best linear fit to any four adjacent pixels spanning the interface; 3) "4-pt Predictive," uses the three adjacent pixels on each side of the interface to predict the pixel on the other; and 4) "6-pt," which uses the three adjacent pixels on each side of the interface to measure the linear discontinuity at the interface. The combination function produces an "overall degree of match" for each adjacent pixel pair, based on the degrees of match returned by each sensor modality. Two different combination functions are compared: 1) "OR," which requires a high degree of match on only one of the sensor modalities; and 2) "AND," which requires a high degree of match

10 LOCAL MSI METHODOLOGY

on both. A total of eight distinct MSI techniques results from the different pairings of these match and combination functions.

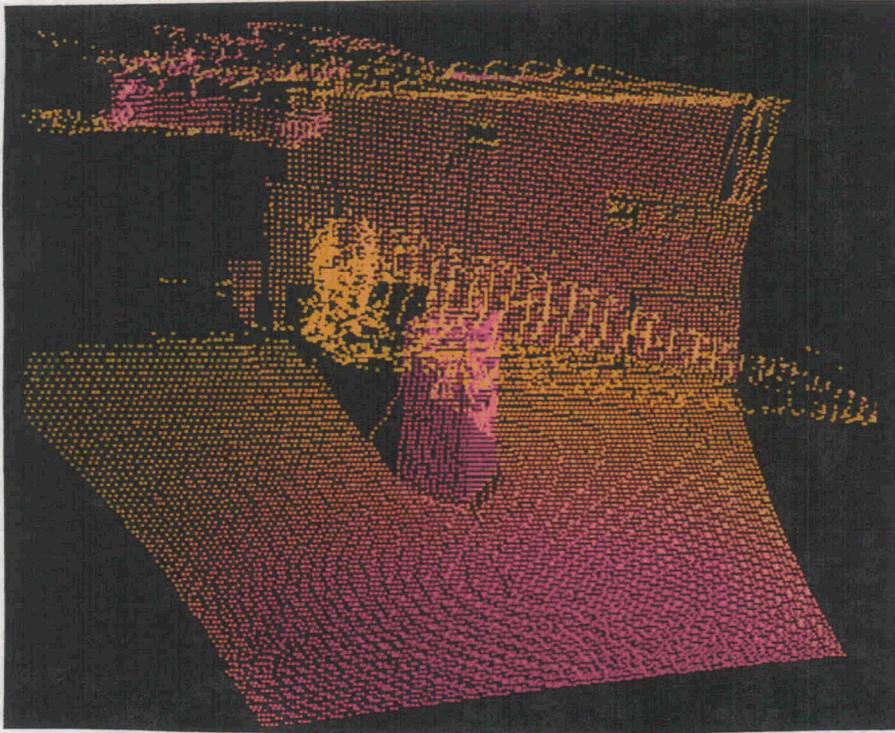


Fig. 3. LRC data of Fig. 2, 3-D plot colored by reflectance.

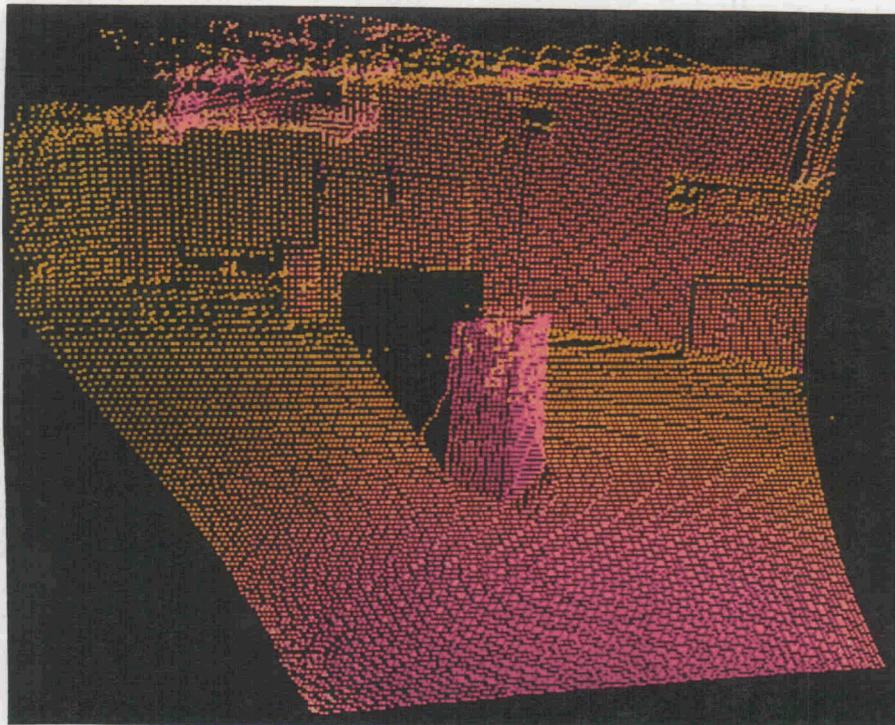


Fig. 4. Wrap-around corrected LRC data of Fig. 3, 3-D plot colored by reflectance.

The final "overall degree of match" resulting from the combination function could, in turn, be used as a meta-level description of the pixel interface. Hence, a ninth technique was devised which compared these overall measures in the neighborhood of the interface. This ninth technique provides a locally adaptive measure of local pixel predictability. By thresholding the relative degree of match with respect to neighboring interfaces, this measure becomes more strict in smooth regions, e.g., corresponding to an artificial surface, and more forgiving in rough regions, e.g., corresponding to natural surfaces like grass or brush. In this way, the ninth technique varies its clustering requirements over the image based on local indications of continuity. This ninth technique is actually eight separate techniques, depending on which of the original eight techniques is used to determine the underlying pixel interfaces' "overall degree of match."

With each of these sixteen techniques is associated various tuning parameters and/or thresholds. To quickly and objectively evaluate these competing techniques, an automatic parameter tuning and image evaluation system was developed. This system reduces the arbitrary effects of blind luck and personal bias and provides a more accurate indication of the various techniques' true performances. The following section details this system and its empirical results.

3. MSI OPTIMIZATION

The automatic parameter tuning system consists of three basic components: a learning strategy, an evaluation function, and a “truth” image. The learning strategy selects parameter settings for analysis, based largely on prior parameter performance, and assesses them by using the evaluation function to compare the resulting clustered images with the predefined optimal results of the “truth” image. Following is a detailed discussion of the learning strategy, evaluation function and the empirical results.

3.1 LEARNING STRATEGY

Genetic Algorithms, GAs, are extremely powerful adaptive global search techniques derived from natural population genetics.⁶ GAs 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 genetic material of two 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 as powerful global search techniques⁷ capable of simultaneously searching extensive regions of the parameter space. The basic GA framework used in this research was provided by the GENESIS⁸ GA package. However, the quality and efficiency of the search is almost completely dependant on the user defined evaluation function.

3.2 EVALUATION FUNCTION

The evaluation function guides and directs the GA’s search by providing performance feedback for candidate “individuals.” To efficiently guide, the function should have a high resolution and monotonically encourage clustering quality so that even small improvements will be reflected in the evaluation measure and have a positive impact on the search.

For this application, an “individual” is a set of parameter values for the conceptual enhancement technique being optimized, and the evaluation function

must return a measure of the quality of each candidate set. Since the goal is to form pixel clusters which most nearly approximate reality, a “truth image” was defined which had all of the pixels in a typical LRC image correctly grouped into their separate conceptual regions. Each of the erroneous pixels in the LRC image was associated with a single group understood to be “don’t care” pixels. That is, the cluster association of those pixels had no effect on the evaluation.

We chose to run the GA as a minimization technique, wherein the GA searches for the parameter set yielding the lowest evaluation measure. This requires the evaluation function to be written as a “penalty function”; the function increases its evaluation measure with each undesirable behavior that it detects.

For the evaluation function to measure how closely the image resulting from a given parameter set matches the truth image, each pixel is assigned a “cluster number” which is constant for all pixels within a single cluster and distinct from all other pixels. The particular cluster number associated with any given cluster is arbitrary. Hence, the evaluation function cannot simply count the number of pixels whose cluster numbers match the truth image’s conceptual region numbers. Rather it must consider each set of pixels which (according to truth) should be united as a single region, and determine how well those pixels are grouped and differentiated from “outside” pixels. Thus, the evaluation function must penalize both the fragmentation and amalgamation of true regions.

Evaluation was performed by considering each true region in order of decreasing size, and associating with each, the candidate cluster which has not already been associated with any other cluster and has the greatest representation within that true region. For example,

let $A_i \in \{\text{Candidate Clusters}\}$
 $T \in \{\text{Truth Regions}\}$
 $D = \{\text{Previously Associated Candidate Clusters}\}$

$$\text{and } T = \bigcup_{i=1}^n A_i$$

then A_k is associated with T if and only if

$$A_k \notin D \wedge |A_k \cap T| \geq |A_i \cap T| \quad \forall i \ni (A_i \notin D) \wedge (1 \leq i \leq n)$$

For each associated cluster, the evaluation function increments the penalty based on the amount by which that cluster failed to cover its true region and by the amount by which it lies outside its true region, see Fig. 5. For example, let $A, B \in \{\text{Candidate Clusters}\}$

$$A = A' \cup A''$$

$$B = B' \cup B''$$

$$T \in \{\text{Truth Regions}\}$$

$$T = A' \cup B'$$

and $|A'| < |B'|$

then the penalty resulting from T is $f(|A'|, |B''|)$, for some function f . This penalty combination encourages clusters to “grow” within the true regions’ bounds and “shrink” outside those bounds. Note also that this function makes no distinction between pixels of different clusters, i.e., every valid, mis-associated pixel causes the same amount of penalty. Hence, an “N” pixel improvement in a small region has the same impact as an “N” pixel improvement in a large region, though the two resulting images may intuitively appear to have very different levels of “clustering quality,” depending on the application. An “optimal” evaluation function is very subjective and application specific and not in the scope of this research. Although this evaluation function may not provide sufficient feedback for the GA to find the globally optimum setting for each competing MSI technique, it does provide sufficient feedback for the GA to objectively optimize, compare, and evaluate them.

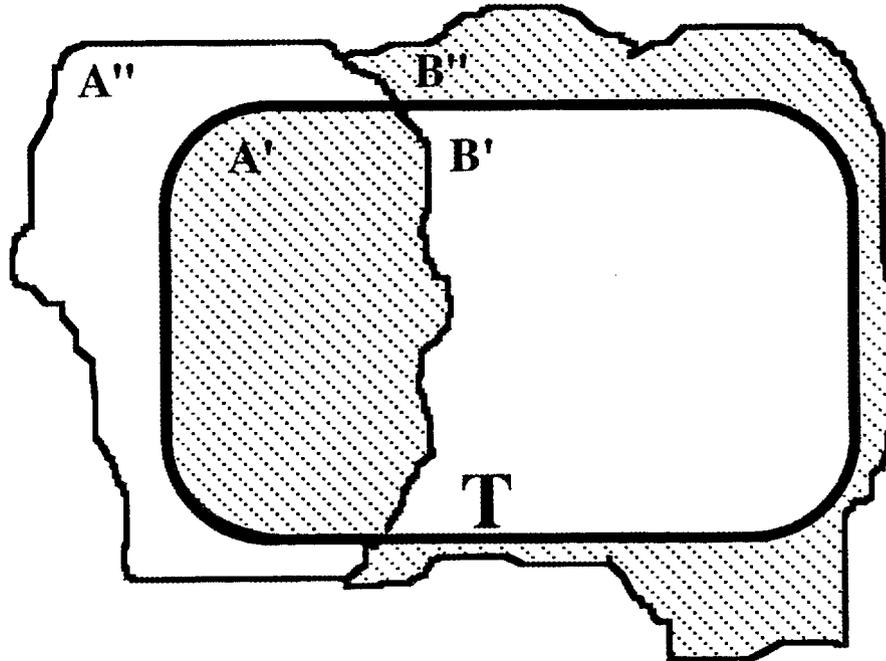


Fig. 5. Cluster matching against truth region T . Cluster B associated with T , penalty is a function of A' and B'' .

3.3 EMPIRICAL RESULTS

The evaluation function provides a numerical measure of the quality of a conceptually enhanced image. This measure is useful not only for the GA’s parameter optimization, 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. Since each search was permitted the same amount of resources, yet the search space sizes ranged from 2^{10} to 2^{47} , the competing techniques had significantly different proportions of their search spaces investigated.

The GA was permitted to examine up to 5000 parameter sets. 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. Hence, 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 of the four match functions, four modes of clustering were GA optimized: each of the two sensor modalities (distance and reflectance) and the two MSI combination functions ("AND" and "OR"). The best solution found after up to 5000 evaluations was taken as that modes optimized solution, see Fig. 6. (Note that the overall degree of match is a penalty function, and therefore, lower values are better.)

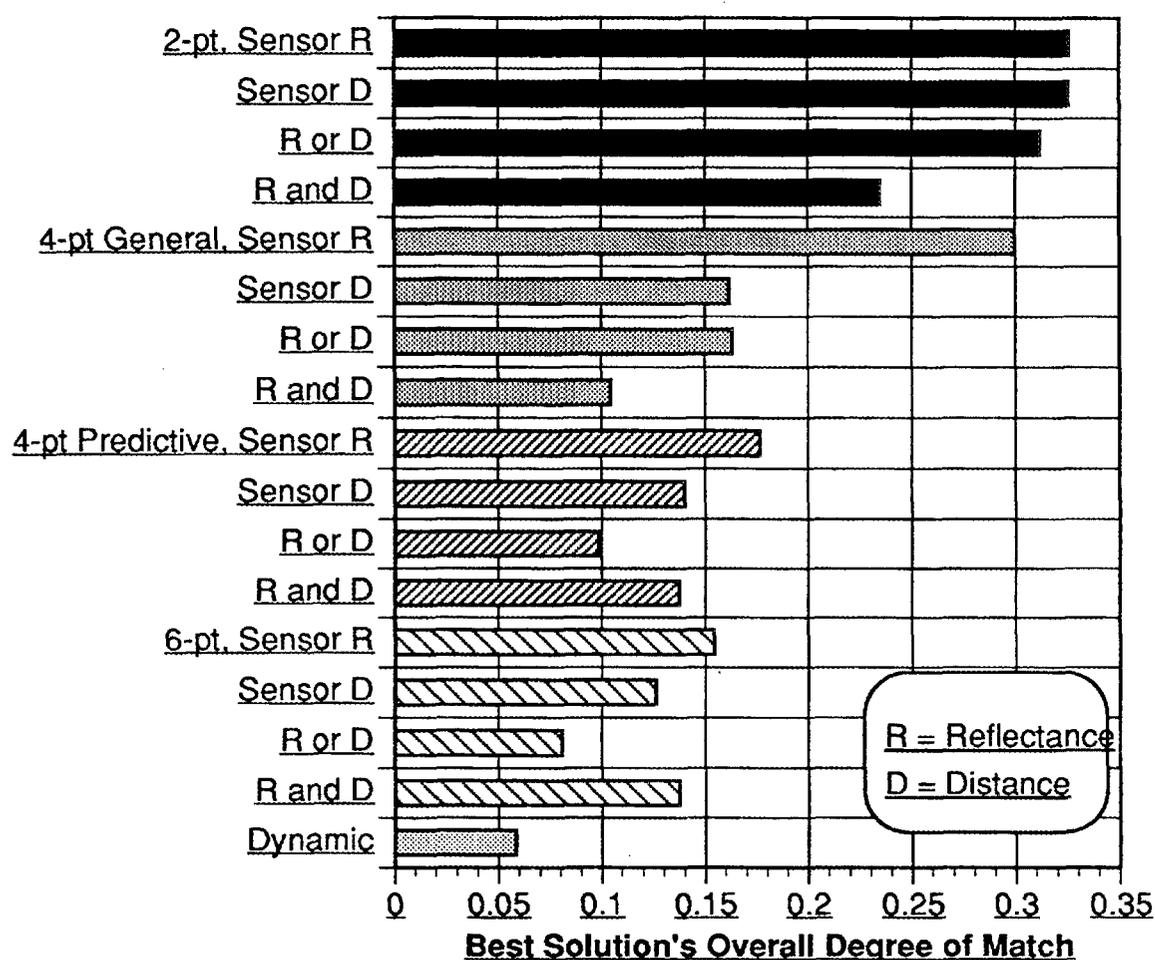


Fig. 6. Comparison of penalty factors for the best solutions found for each method, both for individual sensors and with MSI.

Several interesting tendencies are worthy of note. First, and most importantly, the basic goal of MSI was achieved; MSI improved clustering and therein led to better representations of the real environment. An MSI combination function outperformed both of the single sensors for every match function, see Fig. 6. (By

the theoretical nature of the combination functions, neither could perform worse than the single sensor modes if the optimum parameter values over the entire space $[0...∞]$ was discovered and employed.) Second, for all match functions, the distance image led to better clustering than the reflectance image, see Fig. 6. This may be a result of primarily planar surfaces and linear based matching functions; the distance image embodies the linearity of the surfaces while the reflectance image is subject to non-linear reduction of the intensity signal with distance. Third, there is an interplay between the strictness of the match function (in its size requirements of the true regions, see Table 1) and the strictness of the MSI combination function, see Fig. 7. For less strict match functions, e.g., 2-pt and 4-pt G, the stricter "AND" combination leads to better clustering by preventing over amalgamation. For more restrictive match functions, e.g., 4-pt P and 6-pt, the less restrictive "OR" combination leads to better clustering by preventing over fragmentation. This result implies that the methods we investigated span the proper range of true region sizes within our LRC application domain.

Table 1. Comparison of match function strictness by the size requirements placed on the interfaced regions

Method	Min. Size of Interfaced Clusters	Therefore . . .
2-pt	1 & 1	one cluster 1 pixels
4-pt General	2 & 2 or 3 & 1	one cluster 2 pixels
4-pt Predictive	3 & 1	one cluster 3 pixels
6-pt	3 & 3	both clusters 3 pixels

The dynamic method examines the local uniformity of the overall degree of match "image" for a given type of operation. Hence, there are sixteen different dynamic methods, one for each technique (match function and mode) which can be used to produce the overall degrees of match. Rather than execute the GA an additional sixteen times, the choice of technique was defined as an additional parameter to be simultaneously optimized by the GA. The GA determined that the best technique for dynamic operation is "6-pt OR," which is also the technique found to perform best non-dynamically. Figure 6 shows the superiority of dynamic operation, as it outperforms all other techniques by a significant margin (>27%).

The truth image is shown in Fig. 8. This image is based on the data shown in Fig. 4. The significant true regions include: 1) floor; 2) a box with two facets facing the camera; 3) a movable wall partition; 4) two regions 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.

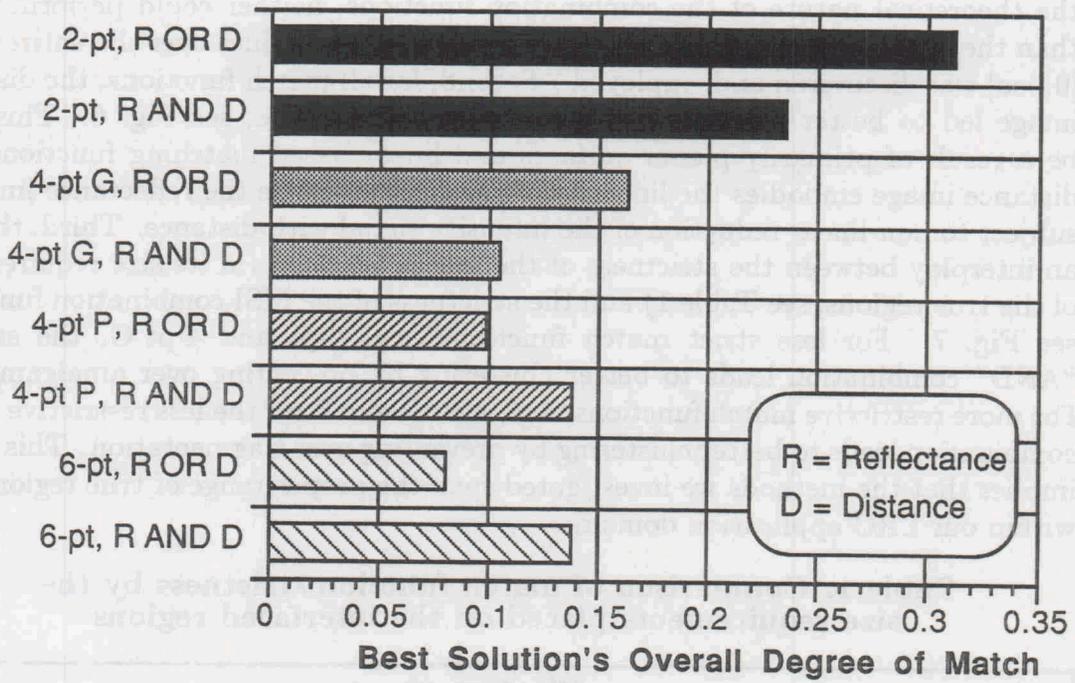


Fig. 7. Comparison of eight MSI techniques, showing combination function trade-off (AND, OR) with respect to match function strictness (2-pt and 4-pt G versus 4-pt P and 6-pt).

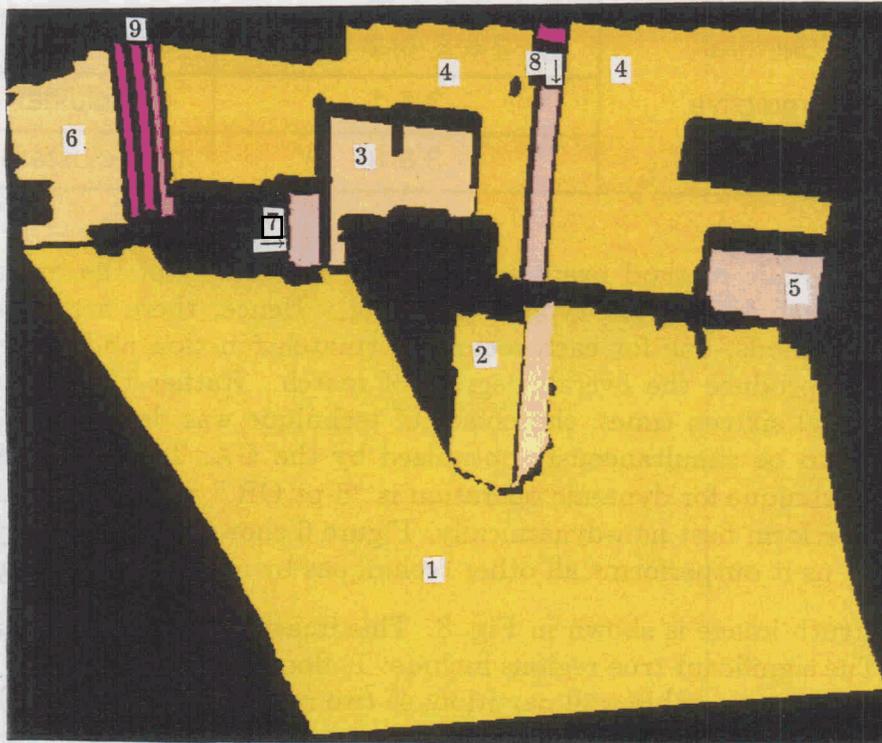


Fig. 8. Truth image of Fig. 4, 3-D plot arbitrarily colored by region size and numbered for reference.

For human display, adjacent points corresponding to the same region are graphically connected. This display mode makes it easier to discern the true regions and their edges. The color of each region is based on its relative size, in pixels. Black space between regions 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 regions, 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 region #7. (Note that the “don’t care pixels” were not excluded from processing; the MSI techniques had to deal with them. But the correct clustering of these pixels was not evaluated since the system should 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, erroneous data is shown with its corresponding cluster, however, to “declutter” the display, only the larger clusters are shown.

Figure 9 presents the results of dynamic operation with “6-pt OR” using the GA’s optimized parameter settings. This image reveals the major regions (#1–#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 considered clutter in the upper right corner of the image were also correctly differentiated. These regions were considered too small and their data too erratic to be included in the truth image. Furthermore, despite wrap-around effecting regions numbered 1, 4, 6, and 8, this MSI technique was able to smoothly form clusters across the wrap-around boundaries and lying within the wrapped area. These results verify the accuracy and continuity of the MSI wrap-around technique described in Section 2.1.

Although this MSI technique failed to differentiate all of the significant true regions in the image, it was very successful. To get an impression of the quality of these results, one need only review the original images which were merged, see Figs. 2–4. In these images, most of the regions, e.g., #3+, are difficult to visually isolate despite our detailed knowledge of them.

Figure 10 presents the results of the “6-pt OR” MSI technique using the GA’s optimized parameter settings. This technique also differentiated most of the major regions, failing only to amalgamate region #6. (Because of region #6’s greater distance, its spacial sampling resolution is lower. This led to greater variation in the data and for a non-dynamic technique to correctly amalgamate those pixels, it would fail to disassociate separate other, more consistent clusters. To correctly handle this problem one must base matching on the local dynamics of the data, e.g., like the dynamic MSI techniques.) All three of the corner supports, #9, the top of the tool box, #5, and a region in the upper right corner were found.

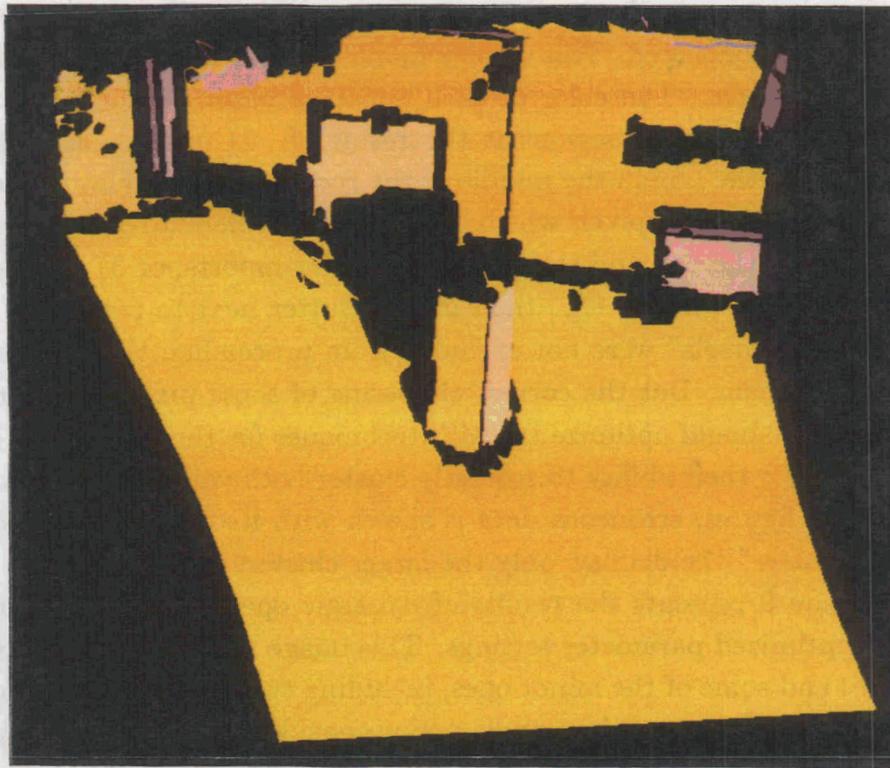


Fig. 9. "Dynamic 6-pt OR" conceptual enhancement of Fig. 4, 3-D plot arbitrarily colored by cluster size.

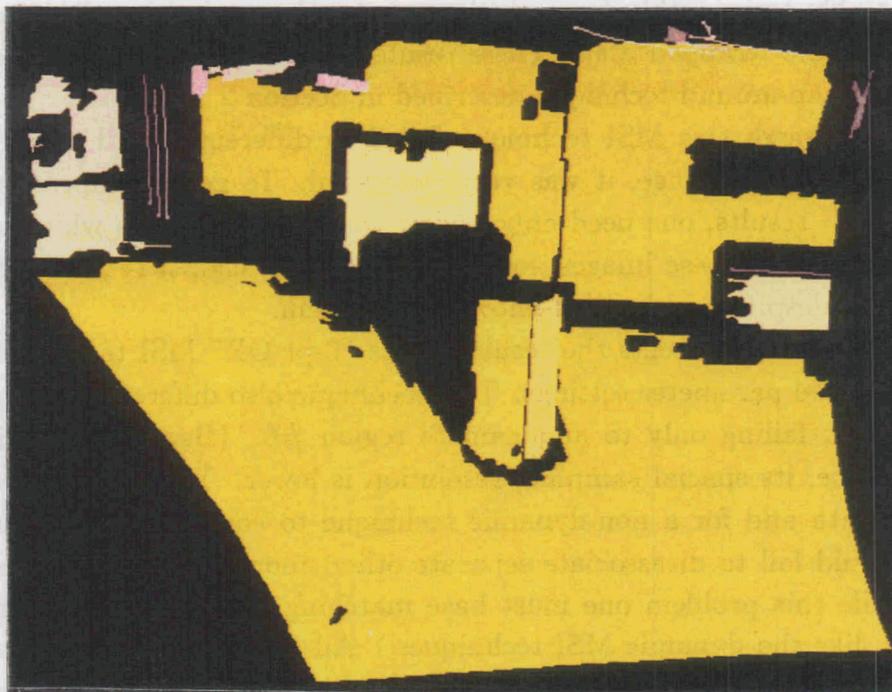


Fig. 10. "6-pt OR" conceptual enhancement of Fig. 4, 3-D plot arbitrarily colored by cluster size.

It should be recalled that the uniform emphasis on pixels rather than any special emphasis on finding each of the primary 14 truth regions helps explain the techniques' "failure" to isolate the smaller regions. However, to do otherwise might demean the general applicability and robustness of the techniques. The promise of these techniques is their ability to produce quality results despite their total lack of global, conceptual shape, or a priori knowledge. This ability is a common requirement in many robotic applications. These MSI results are of sufficient quality to enhance an autonomous system's world map, a teleoperated display, or a higher-level object recognition system.

4. VALIDATION

The MSI conceptual enhancement approach outlined in Section 2.2 is highly domain and application robust. By domain robustness we refer to its suitability for a wide variety of environments and sensor types, though it may first require tuning to the level of image detail desired by the user for that domain, via, e.g., the use of truth images as outlined in Section 3. By application robustness we refer to the suitability of a given parameter set to a wide variety of situations within a tuned or learned domain. This section validates our approach's application robustness by evaluating the performance of its tuned parameter set in the learned domain, i.e., LRC images of an indoor living or warehouse environment, and validates its domain robustness by demonstrating its suitability to another, highly dissimilar domain, i.e., GC images of buried waste.

4.1 LEARNED DOMAIN, LASER RANGE CAMERA

The control parameters of the MSI techniques described in Section 2.2 were tuned to a single LRC image (Fig. 4) and user defined truth image (Fig. 8.) The following validation test is to demonstrate the suitability of these tuned parameters to LRC images of other, "similar" scenes.

Figure 11 shows the LRC's images of another corner of the CESAR laboratory. Figure 11A presents the distance image shown as a 2-dimensional grid with darker color indicating shorter distance from the sensor. Figure 11B presents the reflectance image with darker color indicating greater reflectivity (sheen). For clarity of description, the significant true regions are numbered in Fig. 12: 1) floor; 2) a cylindrical barrel; 3) a moveable wall partition; 4) two regions 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 regions 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; regions numbered 1, 4, and 8 match the correspondingly numbered regions of Fig. 8.

The scene represented in Fig. 11 is more difficult than the learned scene of Fig. 8. This scene contains cylindrical objects (#2 and #5), very narrow objects (#15 and parts of #11), and significantly more detail. The best performing MSI technique from Section 3, i.e., "Dynamic 6-pt OR" and its tuned parameter set was applied to the distance and reflectance images of Fig. 11.

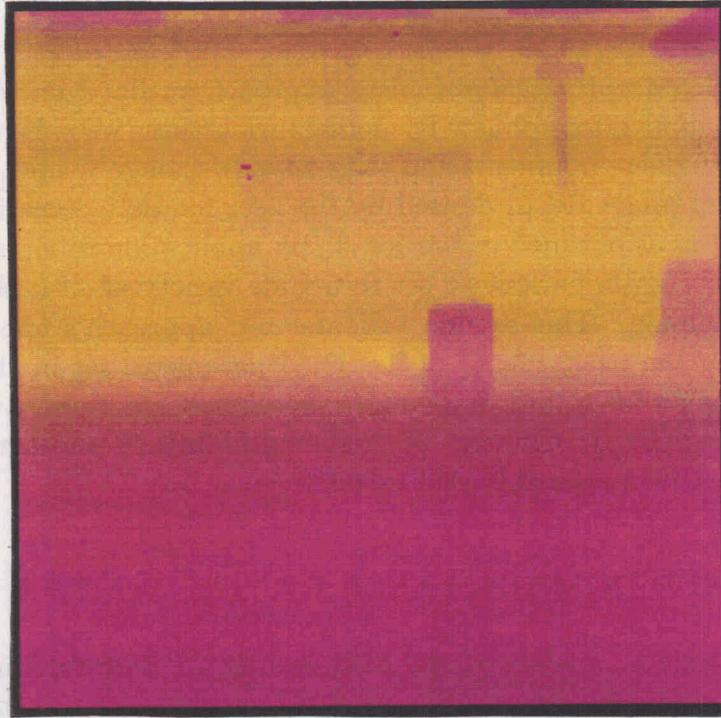


Fig. 11A. Raw LRC data of another corner of the laboratory, with new types of objects interspersed, 2-D plot colored by distance.

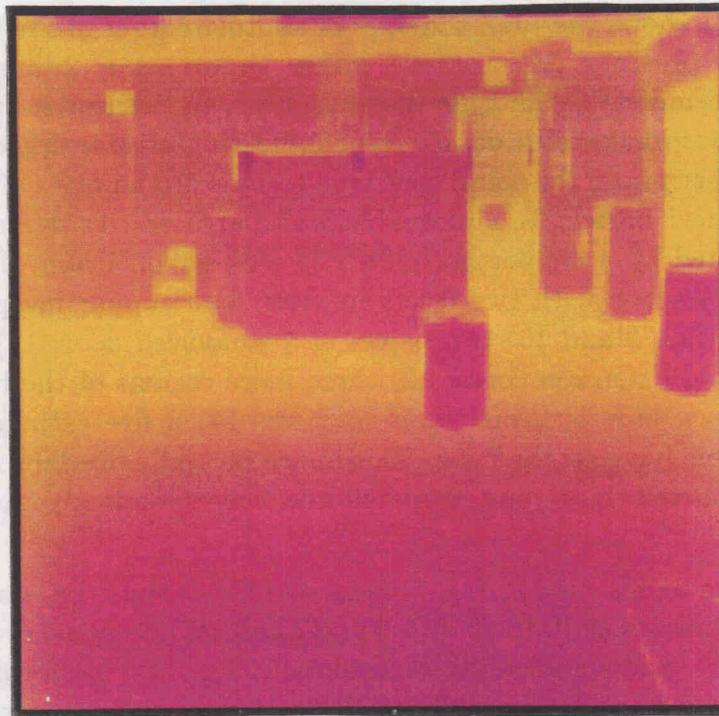


Fig. 11B. Raw LRC data of another corner of the laboratory, with new types of objects interspersed, 2-D plot colored by reflectance.

Figure 12 shows the results of the “Dynamic 6-pt OR” MSI technique using the GA’s optimized parameter settings. For image clarity, only the larger clusters are displayed. This technique resolved most of the regions enumerated in Fig. 11, including both the cylindrical and narrow regions. (Note that cylinders are difficult because the various MSI techniques are based on linear extrapolation and the parameters were tuned on an image consisting entirely of planes.) Furthermore, three regions in the far right corner were correctly, yet unexpectedly, differentiated: the door facing, the doorway offset, and a narrow strip of the right wall. One can also discern the sign and doorknob correctly excluded from region #9, the wire connecting the two primary components of region #11, and the speaker and electrical outlet near the top of region #8. Unfortunately, region #8 itself was not distinguished from region #4 and the extreme detail of the hoist and its cables (region #11) prevented it from being amalgamated into a single, cohesive cluster, though this technique did correctly isolate it from the others. (Note that the support pillar, region #8, was not distinguished in the tuned image either (see Fig. 9) and hence was not expected to be resolved here.)

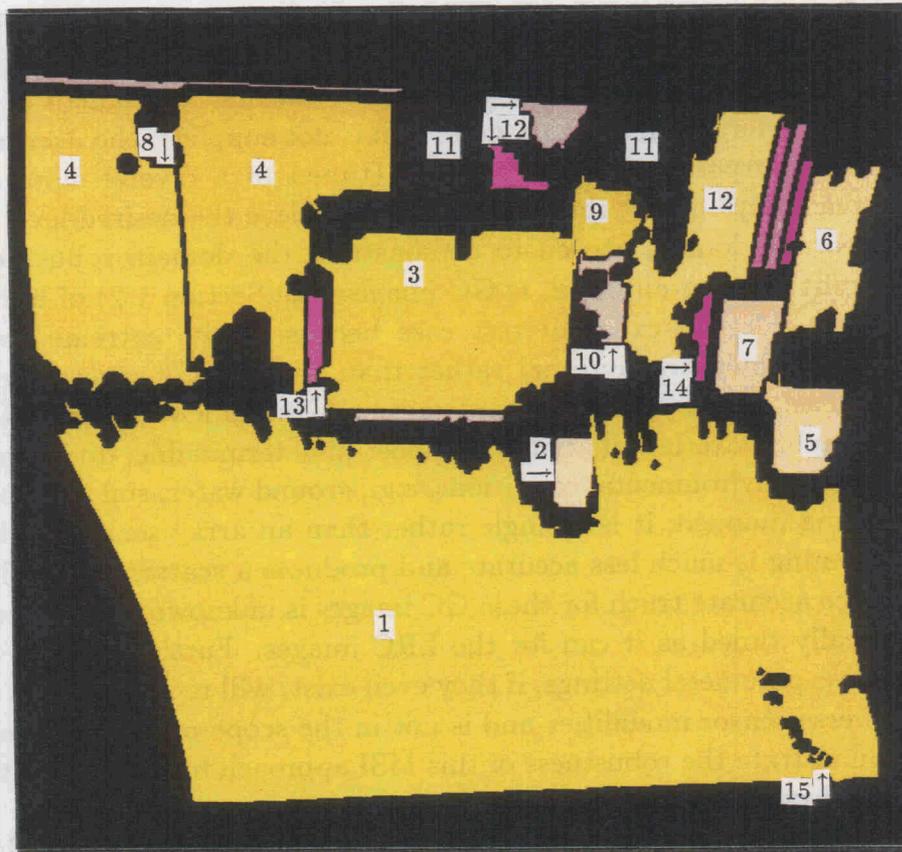


Fig. 12. “Dynamic 6-pt OR” conceptual enhancement of image from Fig. 11 using parameters tuned to Fig. 4, 3-D plot arbitrarily colored by cluster size and numbered for reference.

Figure 12 clearly demonstrates the application robustness of the tuned “Dynamic 6-pt OR” conceptual enhancement technique. After tuning the technique to just one application image, it was able to isolate almost all of the significant regions in an a priori unknown image. Furthermore, these regions included surface types hitherto never seen: curved surfaces. The success of this clustering technique on curved surfaces, despite its linear extrapolation approach and lack of experience with this surface type, strongly supports the utility and robustness of this MSI approach’s concentration on “pixel-level affinity.”

4.2 UNLEARNED DOMAIN, GROUND CONDUCTIVITY

The conceptual enhancement techniques described in Section 2.2 are designed to be completely independent of the sensor modalities being merged and the application being addressed. These are general purpose MSI techniques for merging multiple images into a single, enhanced representation of conceptually disjoint regions. The desired level of image detail, and hence, the intrinsic definition of regions, must be defined by the user according to his application and objectives. 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 meaningful regions and of extraneous detail. From this image, the system can determine the relative information content of the sensor modalities being merged. If a region 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.

The unlearned domain chosen to demonstrate the domain robustness of this conceptual enhancement approach is GC images (see Section 1.2) of buried waste. This sensor provides 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 undeterminable, due in part to the effects of varying environmental conditions, e.g., ground water, soil type, etc., across the region being mapped; it is a single rather than an array sensor, and hence its sample positioning is much less accurate and produces a scatter plot rather than a grid; etc. Since accurate truth for these GC images is unknown, the system can not be automatically tuned as it can for the LRC images. Furthermore, determining general purpose parameter settings, if they even exist, will require images and truth for many diverse sensor modalities and is not in the scope of this research. Hence, we must demonstrate the robustness of this MSI approach by examining the results of manually selected parameter settings.

The raw GC data consists of hundreds of samples taken over a three day period. These samples consist of an X - Y position and two sensor readings: the quadrature component and the in-phase component (see Section 1.2). The data is plotted in Fig. 13. Because of the inconsistent spacing of the data points, the MSI techniques

described above cannot be directly applied. The data must first be transformed into a grid representation. A grid resolution and registration is chosen so that each of the internal grid elements contains samples from 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. 14. (Note that the extreme edges of the original data set were discarded, see Fig. 13, and that the white elements on the lower edges have no corresponding data samples.)

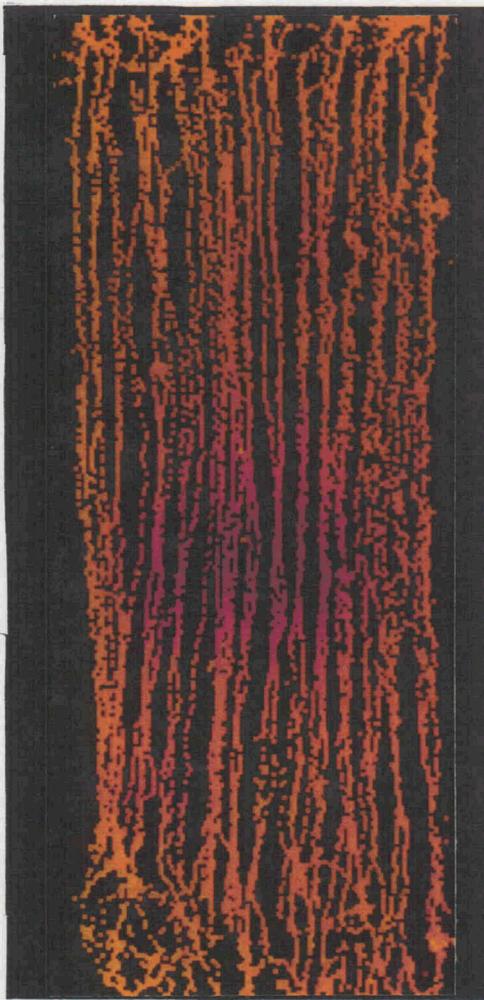


Fig. 13A. Raw GC data of buried waste field, superimposition of three data sets in a 2-D plot colored by quadrature.

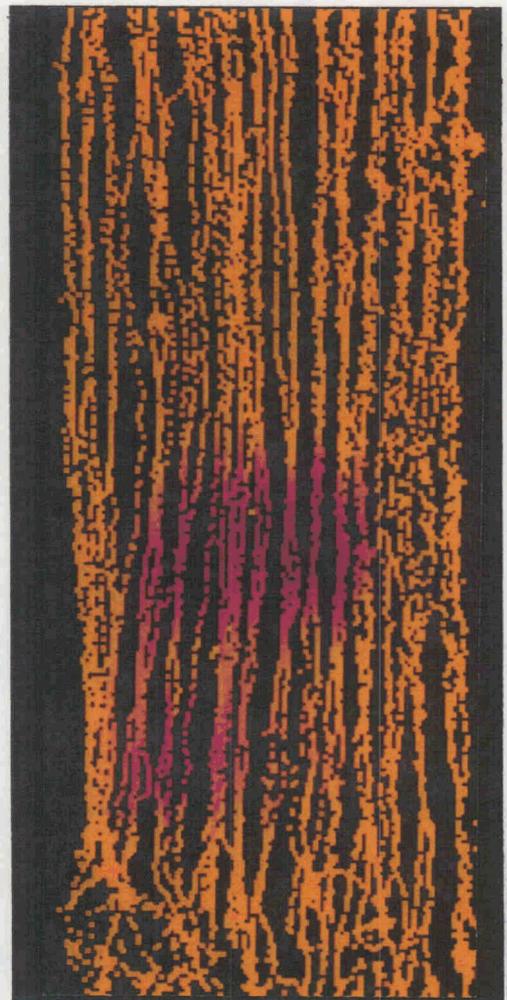


Fig. 13B. Raw GC data of buried waste field, superimposition of three data sets in a 2-D plot colored by in-phase measure.

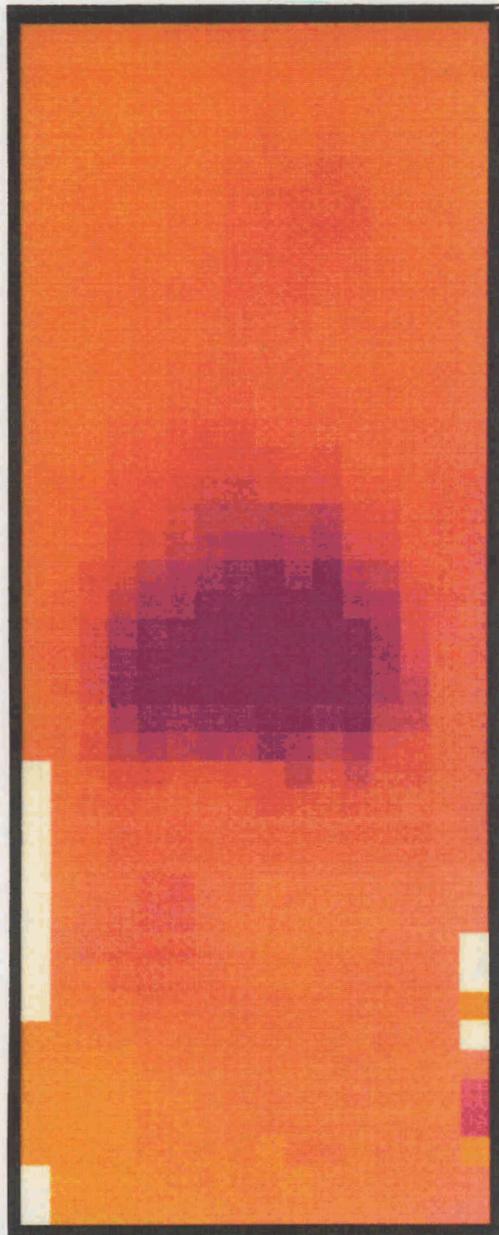


Fig. 14A. Raw GC data of Fig. 13A, represented in a 2-D uniform grid colored by quadrature.

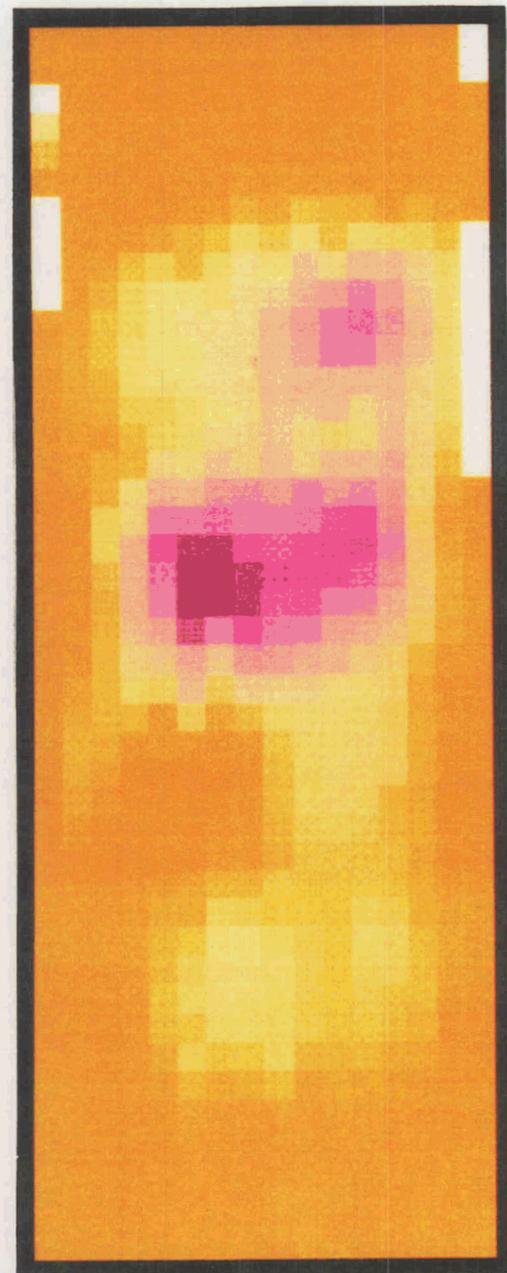


Fig. 14B. Raw GC data of Fig. 13B, represented in a 2-D uniform grid colored by in-phase measure.

Figure 15 presents the results of the "4-pt G-OR" technique with a manually chosen parameter set. With this particular parameter set, the main cluster, near the middle of the image, is isolated, as are smaller clusters just below it and in the upper right-hand of the image. This result is intuitively consistent with the data in Fig. 14. However, without defined truth, e.g., the relative importance of the two sensor images, an indication of the desired level of detail, etc., other solutions are equally viable. Figure 16 presents the results of the "6-pt AND" technique with a

manually chosen parameter set. In this case, the smaller cluster in the upper right of the image is considered insignificant, while the main cluster is examined more critically. Again neither Fig. 15 nor Fig. 16 is offered as a solution. Rather, they represent different, yet equally valid interpretations of the data.

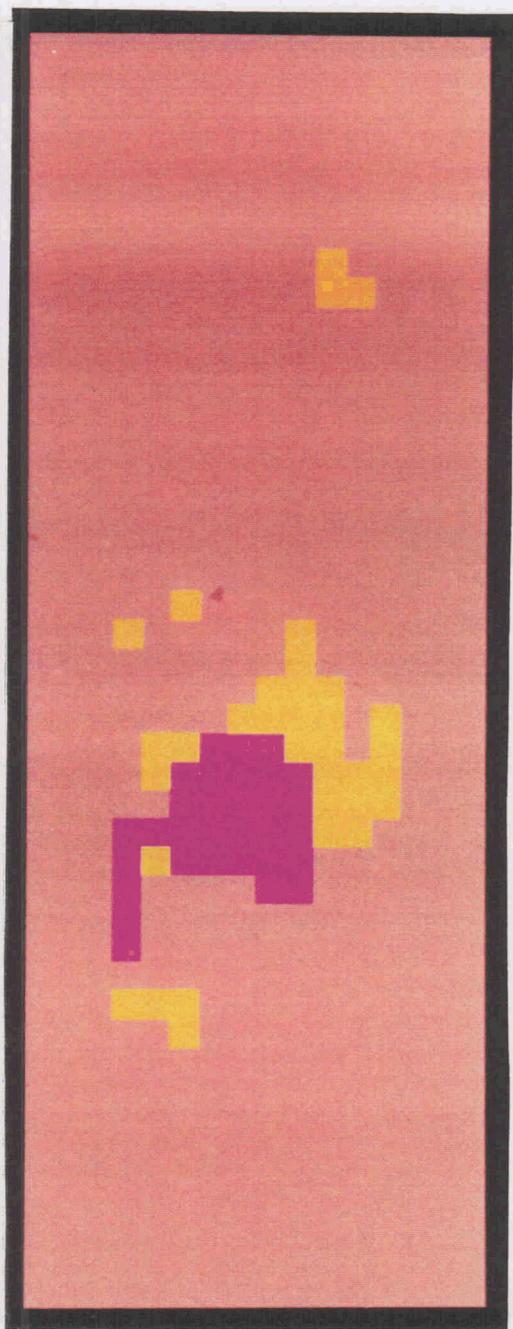


Fig. 15. "4-pt OR" clustering of Fig. 14 using hand-tuned parameters, 2-D grid plot arbitrarily colored by region.

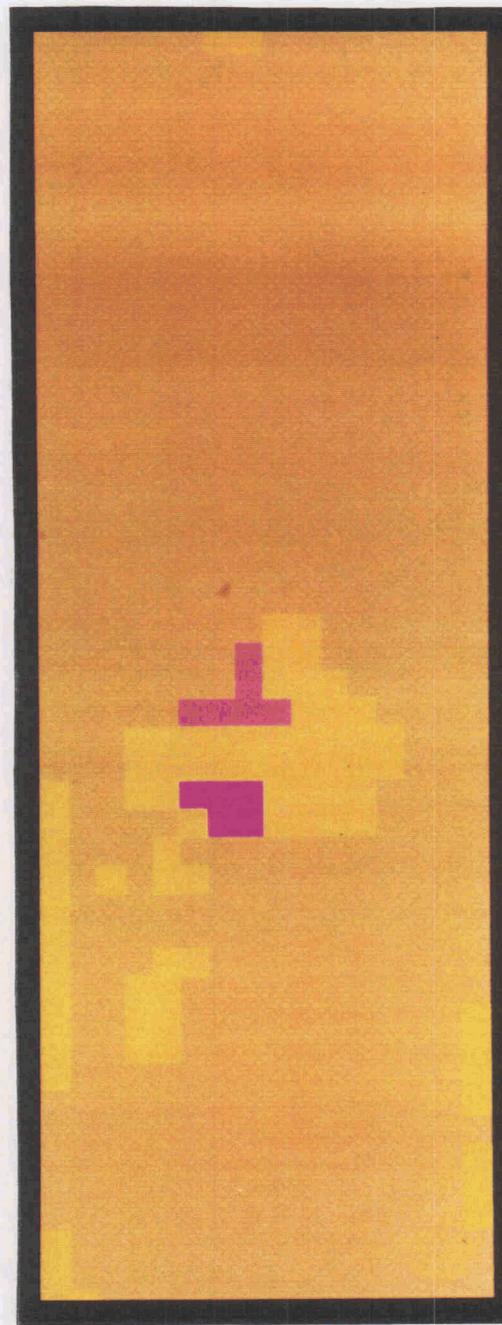


Fig. 16. "6-pt AND" clustering of Fig. 14 using hand-tuned parameters, 2-D grid plot arbitrarily colored by region.

30 *VALIDATION*

Despite the lack of a truth image to which the approach could be tuned and its results evaluated, these MSI techniques are clearly able to isolate significant regions of the merged images. As such, these figures demonstrate the suitability of this MSI approach to GC images and hence its overall domain robustness.

5. SUMMARY AND CONCLUSIONS

MSI techniques can be divided into three categories based on the relative information content of the original images with that of the desired representation: detail enhancement, data enhancement, and conceptual enhancement techniques. Detail enhancement techniques are used to increase the level of image detail or contrast by applying an appropriate inverse mapping function to reduce/remove “blurring” distortions. Data enhancement techniques improve the accuracy of the data rather than either increase or decrease its level of detail. Conceptual enhancement techniques are used when the original image contains more detail than is desired, making it difficult to easily recognize objects of interest.

This paper presents two data enhancement techniques. The first resolves the wrap-around ambiguity problem common to most LRCs. This technique is implemented and tested on real data and found to be very effective for single wrap ambiguities. (Multi-wrap ambiguity was not considered, though a similar solution should be as effective.) The second data enhancement technique increases a sensor’s measurement precision by applying a “discriminating” smoothing filter to each local pixel neighborhood of the image. This filter is sensitive to local noise, discontinuities and the sensor’s known precision. This method is applicable to virtually any imaging sensor and is successfully demonstrated on a LRC’s data.

Conceptual enhancement requires segmenting an image into meaningful clusters which correspond to conceptual entities. The approach presented in this paper is based on the premise that given appropriate sensors, clusters of pixels can be quickly and accurately formed by examining the inter-pixel continuities implied by the local pixel neighborhoods. Hence, this research is based on the potentially faster and more robust formation of “clusters from pixels” rather than the slower process of extracting “clusters from images.” This approach is applicable to cluster formation regardless of sensor and/or application domain—whether one is isolating shapes in CCD images, forms in LRCs, or conductivity anomalies in GC data. As such, this approach constitutes a “general-purpose” MSI methodology.

A suite of conceptual enhancement techniques is designed and applied to the distance and reflectance images of an LRC. These techniques are objectively optimized and compared by an automated parameter optimization system using Genetic Algorithms. The best performing technique, according to this stochastic optimization system, is able to isolate nearly all of the significant regions of the learned image. The domain robustness of this optimized technique is validated on an unknown LRC image in that nearly all of the significant regions of the new image are correctly isolated. Application robustness of this basic approach is validated on an extremely dissimilar sensor and application domain—ground conductivity data of buried waste. Since truth for GC data is unavailable, the conceptual enhancement techniques could not be tuned for this sensor/application domain. However, manual tuning is able to demonstrate the approach’s ability to integrate GC images and isolate significant regions.

32 *SUMMARY AND CONCLUSIONS*

Future research includes: time optimizing these various MSI techniques; optimizing and validating the conceptual enhancement techniques' abilities over a suite of images, so as to permit their use in fully autonomous robotics applications; extending the techniques by permitting heterogeneous match functions across sensor modalities and extending the set of combination functions being investigated; and genetically searching the space of potential match functions upon which to base clustering.

ACKNOWLEDGMENTS

Research sponsored by the Engineering Research Program, Office of Basic Energy Sciences, and by the Office of Technology Development, U.S. Department of Energy, under contract DE-AC05-84OR21400 with Martin Marietta Energy Systems, Inc.

REFERENCES

1. M. Beckerman, "Phase I Report on Sensor Technology, Data Fusion and Data Interpretation for Site Characterization," ORNL/TM-11971 (1991).
2. A. Kak and S. Chen, "Spatial Reasoning and Multi-Sensor Fusion," *Proceedings of the 1987 Workshop*, Morgan Kaufmann Publishers, Inc., San Mateo, California, 1987.
3. M. Beckerman, "A Bayes-Maximum Entropy Method for Multi-Sensor Data Fusion," *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, 1668-1674, Nice, France, May 10-15, 1992.
4. "3-D Laser Imaging System User's Guide," Odetics Publication No. 8533031-PR, March 29, 1990, Odetics, Inc., 1515 S. Manchester, Ave., Anaheim, California.
5. Datacube, Inc., 300 Rosewood Dr., Danvers, Massachusetts.
6. J. E. Baker, "An Analysis of the Effects of Selection in Genetic Algorithms," Ph.D. Thesis, Computer Science Dept., Vanderbilt University, Nashville, Tennessee, May 1989.
7. L. B. Booker, Editor, *Proceedings of the Fourth International Conference on Genetic Algorithms*, Morgan Kaufmann Publishers, Inc., San Manteo, California, 1991.
8. J. J. Grefenstette, "A User's Guide to GENESIS," Version 5.0, direct distribution: gref@aic.nrl.navy.mil, October 1990.

INTERNAL DISTRIBUTION

- | | |
|--------------------|---|
| 1. B. R. Appleton | 26. S. A. Raby |
| 2. S. M. Babcock | 27. K. Rahmani |
| 3-7. J. E. Baker | 28. D. B. Reister |
| 8. A. L. Bangs | 29. P. F. Spelt |
| 9. M. Beckerman | 30. T. Steinel |
| 10. B. L. Burks | 31. F. J. Sweeney |
| 11. P. L. Butler | 32. E. C. Uberbacher |
| 12. R. J. Carter | 33. M. A. Unseren |
| 13. O. Doerum | 34. R. C. Ward |
| 14. J. R. Einstein | 35-36. Laboratory Records
Department |
| 15. C. W. Glover | 37. Laboratory Records,
ORNL-RC |
| 16. W. R. Hamel | 38. Document Reference
Section |
| 17. J. P. Jones | 39. Central Research Library |
| 18. H. E. Knee | 40. ORNL Patent Section |
| 19. R. C. Mann | |
| 20. E. M. Oblow | |
| 21-25. F. G. Pin | |

EXTERNAL DISTRIBUTION

41. Office of Assistant Manager, Energy Research and Development, Department of Energy, Oak Ridge Operations, Oak Ridge, TN 37831
42. Dr. Peter Allen, Department of Computer Science, 450 Computer Science, Columbia University, New York, NY 10027
43. Dr. Wayne Book, Department of Mechanical Engineering, J. S. Coon Building, Room 306, Georgia Institute of Technology, Atlanta, GA 30332
44. Professor Roger W. Brockett, Wang Professor of Electrical Engineering and Computer Science, Division of Applied Sciences, Harvard University, Cambridge, MA 02138
45. Prof. John J. Dorning, Department of Nuclear Engineering and Physics, Thornton Hall, McCormick Rd., University of Virginia, Charlottesville, VA 22901
46. Dr. Steven Dubowsky, Massachusetts Institute of Technology, Building 3, Room 469A, 77 Massachusetts Ave., Cambridge, MA 02139
47. Dr. Avi Kak, Robot Vision Lab, Department of Electrical Engineering, Purdue University, Northwestern Ave., Engineering Mall, West Lafayette, IN 47907
48. Dr. James E. Leiss, Rt. 2, Box 142C, Broadway, VA 22815-9303
49. Dr. Oscar P. Manley, Division of Engineering, Mathematical, and Geosciences, Office of Basic Energy Sciences, ER-15, U.S. Department of Energy-Germantown, Washington, DC 20545
50. Prof. Neville Moray, Department of Mechanical and Industrial Engineering, University of Illinois, 1206 West Green St., Urbana, IL 61801
51. Dr. Wes Snyder, Department of Radiology, Bowman Gray School of Medicine, 300 S. Hawthorne Dr., Winston-Salem, NC 27103
52. Prof. Mary F. Wheeler, Rice University, Department of Mathematical Sciences, P.O. Box 1892, Houston, TX 77251

53-54. Office of Scientific Technical Information, P.O. Box 62, Oak Ridge,
TN 37831