ARO 14305.1-A-EL

INTERACTIVE AIDS FOR CARTOGRAPHY AND PHOTO INTERPRETATION

12

Semiannual Technical Report

Covering the Period 2 November 1977 to 11 May 1978

June 1978

LEVEL II

A043418

By: G. J. Agin
H. G. Barrow
R. C. Bolles
M. A. Fischler
T. D. Garvey
J. H. Kremers
L. Quam
J. M. Tenenbaum
H. C. Wolf

ARPA Order No. 2894-5
Contract No. DAAG29-76-C-0057
Effective Date: 12 May 1976
Expiration Date: 9 October 1978

Harry G. Barrow (415) 326-6200, Ext. 5089
Principal Investigator

Approved for public release; distribution unlimited

This research was sponsored by the Defense Advanced Research Projects Agency of the Department of Defense and was monitored by the U.S. Army Research Office under Contract DAAG29-76-C-0057.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the United States Government.

SRI International
333 Ravenswood Avenue
Menlo Park, California 94025
(415) 326-6200
Cable: SRI INTL MNP
TWX: 910-373-1246

INTERACTIVE AIDS FOR CARTOGRAPHY AND PHOTO INTERPRETATION

Semiannual Technical Report

Covering the Period 2 November 1977 to 11 May 1978

June 1978

By: G. J. Agin, H. G. Barrow, R. C. Bolles, M. A. Fischler, T. D. Garvey, J. H. Kremers, L. Quam, J. M. Tenenbaum, H. C. Wolf

ARPA Order No. 2894-5
Contract No. DAAG29-76-C-0057
Effective Date: 12 May 1976
Expiration Date: 9 October 1978

Harry G. Barrow (415) 326-6200, Ext. 5089
Principal Investigator

This research was sponsored by the Defense Advanced Research Projects Agency of the Department of Defense and was monitored by the U.S. Army Research Office under Contract DAAG29-76-C-0057.

Approved:

Peter E. Hart, Director
Artificial Intelligence Center

SRI International

ABSTRACT

The central scientific goal of the ARPA Image Understanding Project research program at SRI International is to investigate and develop ways in which diverse sources of knowledge may be brought to bear on the problem of interpreting images. The research is concerned with specific problems that arise in processing aerial photographs for such military applications as cartography, intelligence, weapon guidance, and targeting. A key concept is the use of a generalized digital map to guide the process of image analysis.

The objectives, methodology, and current status of our research are described in this report. In the present phase of our program, the primary focus is on developing a "road expert," whose purpose is to monitor and interpret road events in aerial imagery. We also describe herein the details of an innovative procedure for tracking road segments and finding potential vehicles in imagery of approximately 1-3 feet per pixel ground resolution. This work represents an important contribution towards achieving the indicated goals.

CONTENTS

ILLUSTRATIONS

ACKNOWLEDGMENTS

The work reported here was performed by G.J. Agin, H.G. Barrow, R.C. Bolles, M.A. Fischler, T.D. Garvey, J.H. Kremers, L. Quam, J.M. Tenenbaum, and H.C. Wolf.

The research was supported by the Defense Advanced Research Projects Agency of the Department of Defense and was monitored by the U.S. Army Research Office under Contract No. DAAG29-76-C-0057.

I DETECTING AND INTERPRETING ROAD EVENTS IN AERIAL IMAGERY

A. Introduction

Research at SRI International under the ARPA Image Understanding Program was initiated to investigate ways in which diverse sources of knowledge might be brought to bear on the problem of analyzing and interpreting aerial images. The initial phase of research was exploratory in nature and identified various means for exploiting knowledge in processing aerial photographs for such military applications as cartography, intelligence, weapon guidance, and targeting. A key concept is the use of a generalized digital map to guide the process of image analysis.

The results of this earlier work were integrated in an interactive computer system called "Hawkeye."[1] This system provides necessary basic facilities for a wide range of tasks and a framework within which specialist programs can be integrated. Research is now concentrated upon development of a program capable of expert performance in a specific task domain: road monitoring. The following sections of this report present an overview of this new effort.

B. Objective

The primary objective of this research is to build a computer system that "understands" the nature of roads and road events. It should be capable of performing such tasks as:

* Finding roads in aerial imagery
* Distinguishing vehicles on roads from shadows, signposts, road markings, etc.

----

1 H.G. Barrow, "Interactive Aids for Cartography and Photo Interpretation," Semiannual Technical Report to ARPA, Contract DAAG29-76-C-0012, Artificial Intelligence Center, SRI International, Menlo Park, California, (October 1977).

* Comparing multiple images and symbolic information pertaining to the same road segment, and deciding whether significant changes have occurred.

It should be capable of performing the above tasks even when the roads are partially occluded by clouds or terrain features, or are viewed from arbitrary angles and distances, or pass through a variety of terrains.

C. Approach

To achieve the above capabilities, we are developing two "expert" subsystems: the "road expert" and the "vehicle expert." The road expert knows mainly about roads, how to find them in imagery, and what things belong on them. It works at low-to-intermediate resolution (e.g., from 1 to 20 feet of ground distance per image pixel) and has the ability to distinguish vehicles from other road detail. The vehicle expert will work on higher-resolution imagery and can identify vehicles as to type. We are concentrating our initial efforts on the road expert and therefore will limit our discussion to this component of our system.

Among the specific tasks to be performed by the road expert are the following:

* Place an image in correspondence with the map data base
* Establish the precise location of known roads in the image
* Determine the visibility of the located road segments
* Mark the road centerline and lane boundaries
* Detect anomalous regions on and along the road pavement
* Determine which anomalies are potential vehicles.

The image/map correspondence task will be accomplished primarily by using roads as landmarks; thus, the first 3 tasks will interact strongly with one another. These tasks will be performed at approximately 20 feet/pixel resolution so that a reasonably wide field of view (10 to 100 square miles) can be processed at one time.

After visible portions of roads have been located, individual sections will be selected for detailed analysis. With resolution increased to approximately 1-3 feet/pixel we can find the road centerline and lane boundaries, starting with the initial estimate obtained in the low-resolution step. We will then detect anomalous regions on and along the road pavement, and finally decide which of these regions are vehicles. Since road anomalies will cause problems in tracking a nominally homogeneous road surface, Tasks 4 through 6 will be integrated to some extent.

The above tasks will be supported by information about road condition and general structure from a symbolic data base. For example, if prior photographic coverage of the area being analyzed is available, the problem of anomaly classification can be simplified by determining if a similarly shaped anomaly had been found in the same general location over some extended period of time. Additional examples of how data-base knowledge and stored models can aid in the analysis process include: using the time of day in discriminating shadows from objects of interest; utilizing the general shape and width of the road (obtained from a map) as an aid in road tracking; providing relevant information on the anticipated size, shape, and road orientation of potential vehicles.

A central theme of this effort is to consider roads as a knowledge domain. In particular, we plan to address the question of how apriori knowledge can be directly invoked by the image-processing modules (what type of knowledge, how should it be represented, and what are the mechanisms for its use?). To achieve our goal of building a very-high-performance system, we plan to develop explicit models of the image structures we will be dealing with, as well as models of the decision procedures embedded in the image-processing algorithms -- so that the algorithms can evaluate their own performance. Finally, we must develop an overall control structure, whose function will be to coordinate analysis across a number of levels of resolution, and to integrate multisource information.

D. Progress

Working programs already exist that are capable of performing most of the major tasks to be performed by the road expert. For example, the problems of image map correspondence and road marking (at low resolution) can be dealt with in an integrated way using the road-tracking and chamfer-matching techniques reported on in our description of the Hawkeye system[1]. Figures 1-6 illustrate this capability.

Figure 1 presents an overview of one of our test sites in the San Pablo area north of Oakland, California.

Figure 2 shows a section of the road network that appears in Figure 1, as located by our road tracking algorithm.

Figure 3 shows our data base representation of this road network superimposed upon the data in Figure 2 as a set of point landmarks. The lack of correspondence is due to our initial error in estimating the location of the point in space from which the image in Figure 1 was acquired.

Figure 4 displays an intermediate stage in the iterative calibration process in which the camera model is being refined using "goodness-of-fit" data provided by the chamfer-matching technique.

Figure 5 shows the final result, in which accurate calibration has been achieved.

Figure 6 shows that the road follower can operate successfully even in a rather cluttered urban environment; thus, our ability to locate roads and utilize them for achieving image map correspondence is not restricted to rural areas.

Some new work now in progress on an extension of the chamfer-matching technique allows a more precise estimate of the mismatch between a reference map and a sensed image, and thus should enable a more reliable and accurate final calibration. The basic idea is to associate individual points in an image and reference map, rather than simply measure the distance between a point in the reference map and

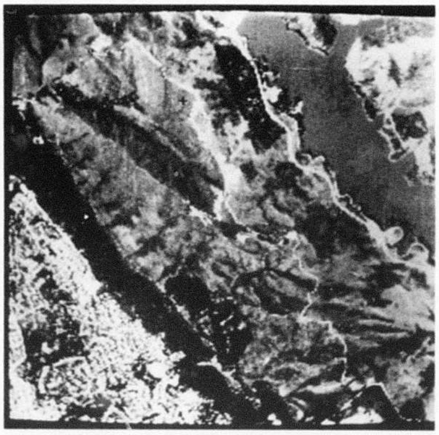

FIGURE 1 A TEST SITE IN THE SAN PABLO AREA

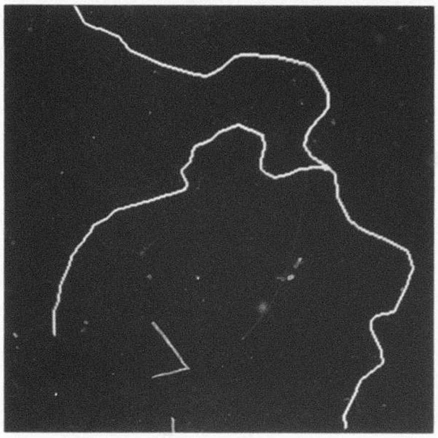

FIGURE 2 RESULTS OF ROAD TRACKING

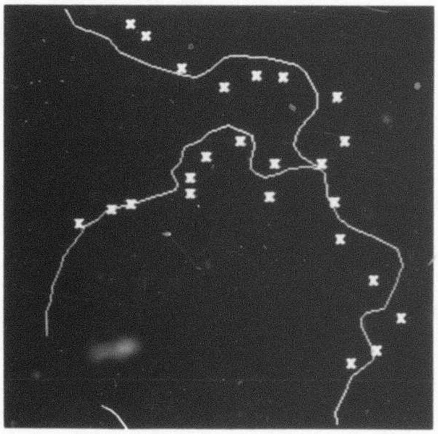

FIGURE 3 ROAD POINTS FROM THE DIGITAL MAP

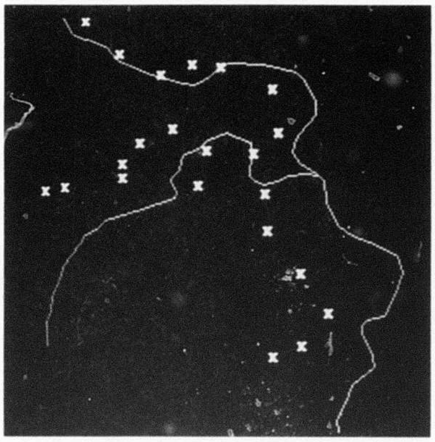

FIGURE 4 ITERATIVE MATCHING IN PROGRESS

FIGURE 5 RESULT OF MATCHING IMAGE AND MAP

FIGURE 6 ROADS FOUND IN AN URBAN AREA

some unspecified nearest point in the corresponding image. It turns out that very little additional computation is required beyond that needed for the older chamfer-matching algorithm.

This technique, which will be described in detail in a forthcoming technical note, is illustrated in Figures 7-11.

Figure 7 is a high altitude image of the San Francisco Bay Area. Figure 8 is an outline of the land-water boundaries in Figure 7, extracted by an edge follower.

Figure 9 shows the map data base representation of these boundaries, superimposed upon the data of Figure 8. The mismatch reflects our initial calibration error.

Figure 10 shows an intermediate stage in the iterative calibration process. The links between the map and image boundaries are automatically inserted by the chamfering algorithm, and their total length is a function of the existing mismatch.

Figure 11 shows the final result of the calibration process.

Figures 15 through 19, discussed in the next section of this report, illustrate some of our new capabilities in following roads and marking potential anomalies in higher resolution imagery. For example, Figure 15a depicts a portion of our test site near Palo Alto, California; Figure 15b shows our ability to accurately locate the road centerlines and mark potential anomalies; Figure 19 shows that we can even track poorly defined dirt roads.

All of the above capabilities, notwihstanding rather impressive performance, represent merely an initial stage in the development of the high-performance road expert we have defined as our goal in this effort. Our current programs are "low-level," in the sense that they still cannot communicate with one another, or modify their performance on the basis of context or self-evaluation. In almost all cases, their level of performance is expected to improve substantially as we integrate the individual modules and modify them to accept data-base support.

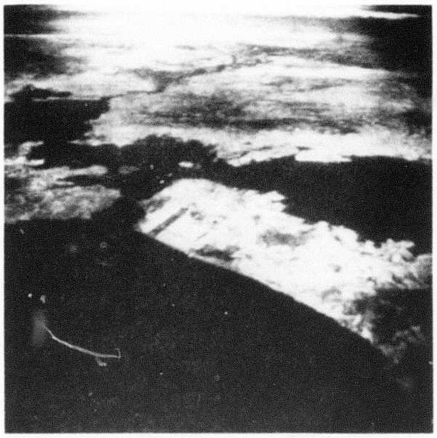

FIGURE 7 SAN FRANCISCO BAY AREA

FIGURE 8 RESULTS OF TRACING COASTLINES

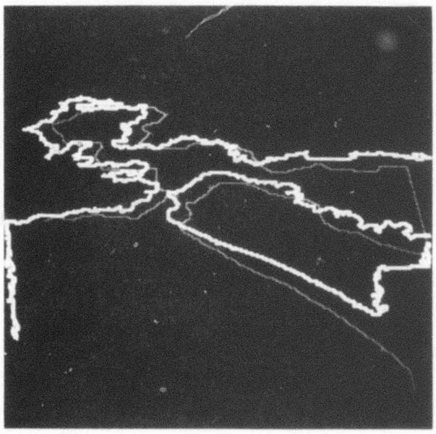

FIGURE 9 COASTLINES FROM THE DIGITAL MAP

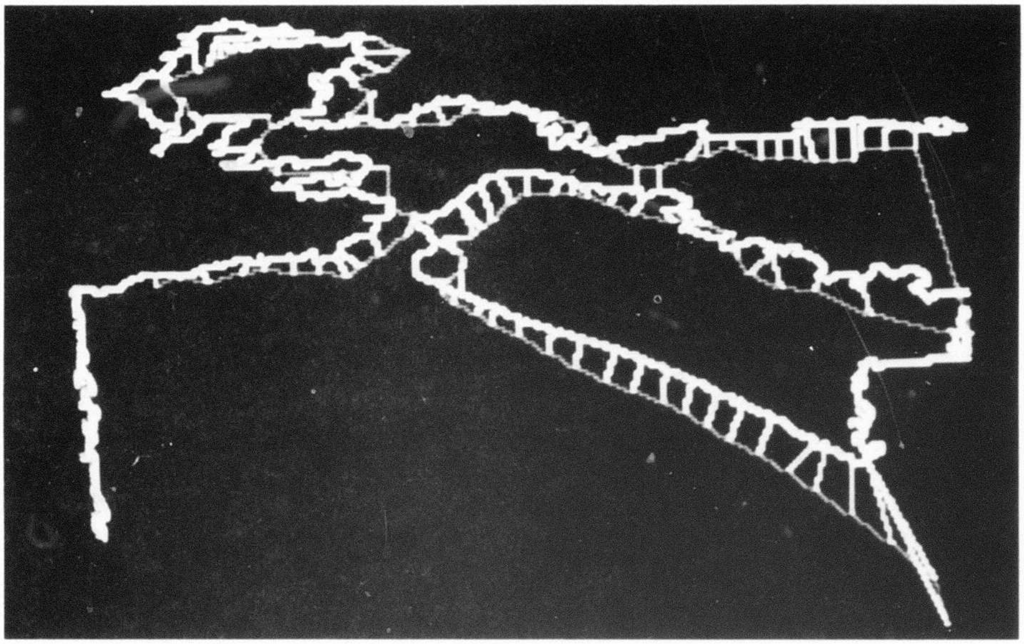

FIGURE 10 ITERATIVE MATCHING IN PROGRESS

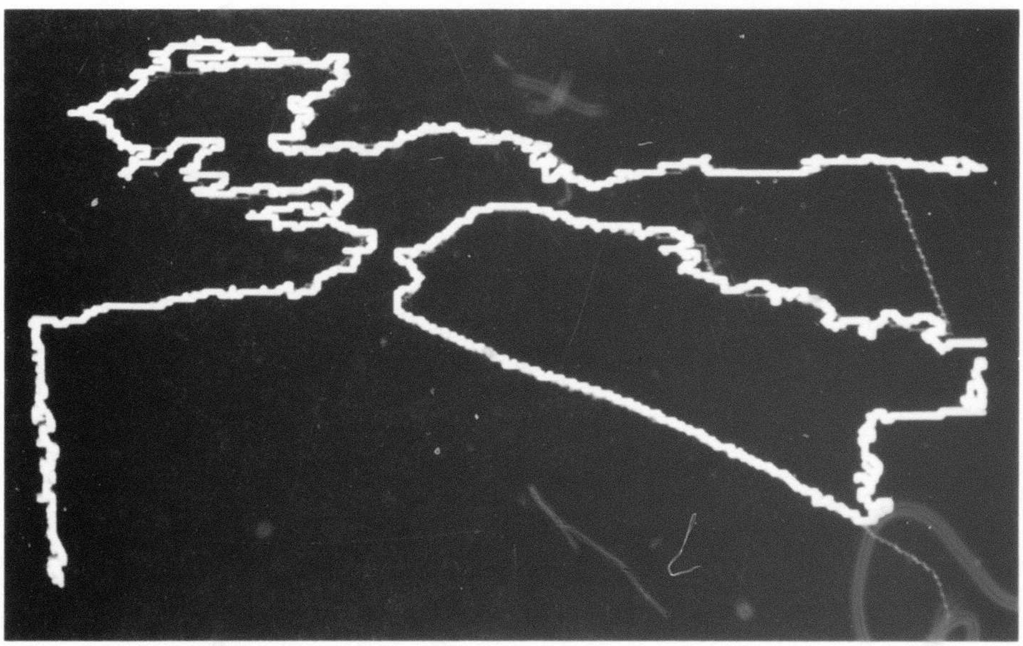

FIGURE 11 RESULT OF MATCHING IMAGE AND MAP

There are still a number of problems that we must contend with. In Figure 17, the road centerline tracker is shown to have successfully followed a curved section of roadway, but it gets somewhat confused when the pavement material changes abruptly; here it marks a section of the road as a potential anomaly. Certainly, such a problem could be alleviated by using data-base knowledge of the fact that such a road surface transition does occur at the given location. Even a previous image (with a suitable time differential), showing a similarily located anomaly, might suffice to largely eliminate such problems. However, in Figure 12 we see two images of the same area taken approximately six months apart. In the earlier image we find two arrows painted on the road surface; six months later there is a moving vehicle positioned directly over one of the arrows just at the moment the photograph was taken. Any simple heuristic routine which attempts to use data-base knowledge in this situation will almost certainly confuse the vehicle with the assumed road marking.

In the image shown in Figure 13, we see that the six-month interval between the acquisition times of the two images has produced some very significant photometric changes (reversals in intensity between the road surfaces and their surrounds), landscape changes (a creek with bordering trees to the left of the curved road segment looks dramatically different in the two images), and even the improbable movement of what is apparently a fair-sized building from its initial location near the creek to a point perhaps 60 feet away in the later image.

The point we wish to make here is that access to data-base knowledge does not necessarily provide a simple solution to the difficult problems of image interpretation. If used indiscriminately, such knowledge can cause more problems than it eliminates. We believe that we can make an important contribution by determining what type of apriori knowledge can "pay" for itself, i.e., will provide a gain in performance commensurate with its complexity, reliability, acquisition cost, etc.

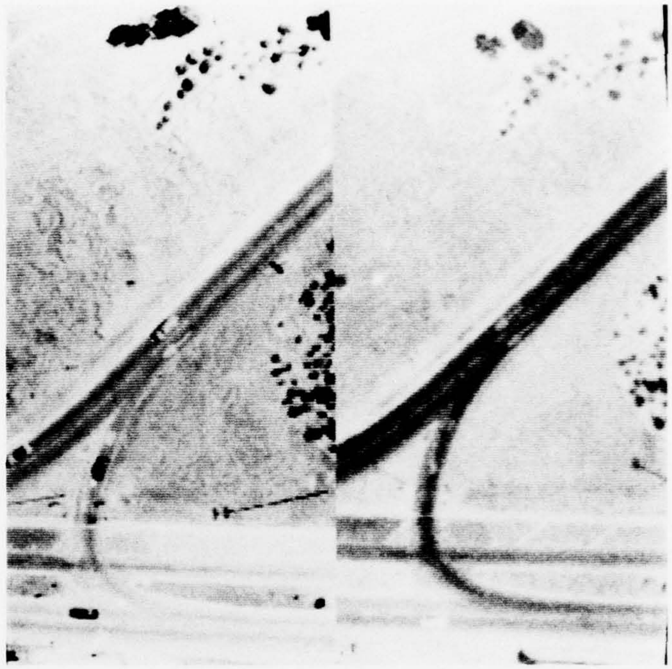

FIGURE 12 TWO IMAGES SEPARATED BY SIX MONTHS

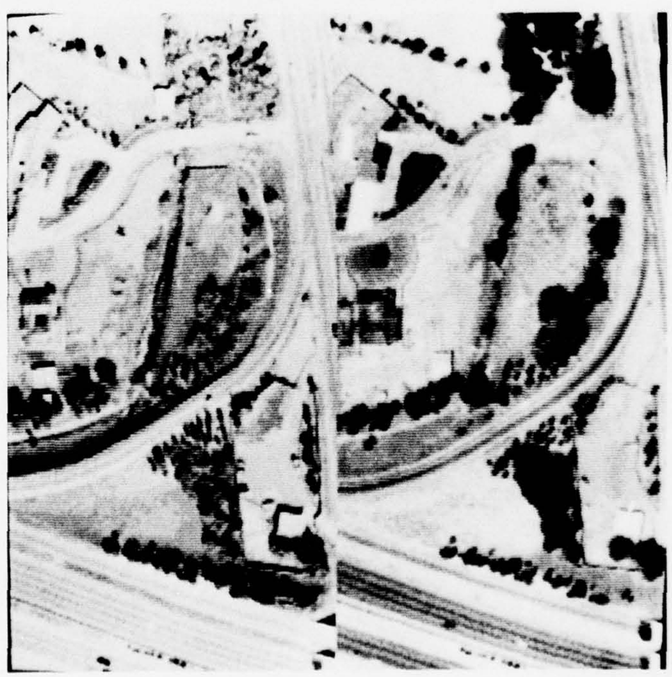

FIGURE 13 PHOTOMETRIC CHANGES IN SIX MONTHS

The images presented in the above examples are part of a data base we are building in order to support our experimental work. We have acquired multiple photographic coverage of five distinct sites (Figures 1 and 14) scattered around the San Francisco Bay Area. This imagery, most which has been scanned, shows road detail at resolutions ranging from 1 to 20 feet of ground distance per image pixel, and reflects significant variations in time and viewing position.

II ROAD TRACKING AND ANOMALY DETECTION

This section of the report describes a new procedure for tracking road segments and detecting potential vehicles in aerial imagery of approximately 1 to 3 feet per pixel ground resolution. The road tracking algorithm discussed here is currently initiated by manually specifying the position of the center and the direction of the road fragment we wish to analyze. The nominal road width could be supplied by the user, by the data base, or by an image analysis function that can determine the width of a road fragment. The road tracker produces two forms of output: a point list describing the track of the road center, and a binary image of all points in the road that are anomalous and might represent to vehicles. In the complete road-expert system, this image will then be analyzed to screen false alarms and interpret the remaining anomalies.

A. Algorithm Description

Figure 14 shows a representative road scene digitized at a ground resolution of approximately 2 feet per pixel and containing segments of a multilane freeway, with a few vehicles and road surface markings (painted arrows and words in the extreme left lane). The variations in road surface materials, centerlines, and intralane wear patterns correspond linearly to the road itself. The vehicles and other anomalies, however, stand out as being quite different from the pattern of the road.

These observations lead to a simple model for the photometry of most of the road surface. The model predicts that successive road reflectance profiles extracted perpendicularly to the direction of the road and centered upon it should have approximately the same appearance. Deviations from the model, defined as anomalies, are caused by road

FIGURE 14 A FREEWAY INTERSECTION TEST SITE

surface markings, vehicles, shadows, changes in road width and constituent materials, and by other influencing factors. The first attempt at a road-tracking algorithm exploited this model. Successive road reflectance profiles (RRP) extracted perpendicular to the direction of the road showed a high degree of correlation, which suggested that road tracking could be accomplished by using a cross-correlation-based approach. The location of the correlation peak was used to maintain alignment with the road center and to generate a model for the road trajectory. However, this first attempt turned out to be unsatisfactory because anomalies perturbed the location of the correlation peak and, concomitantly, other small errors in locating the correlation peak accumulated.

To overcome these problems, six refinements were introduced:

* A cumulative road reflectance profile model
* Bounded-difference alignment
* Masked alignment
* Match-function peak interpolation
* Anomaly detection
* Trajectory extrapolation.

The cumulative road reflectance profile model was introduced to reduce the tendency of alignment errors between successive RRPs to cause an increasing drift away from the road center. Each new RRP is aligned with the current RRP model, rather than with the preceding RRP. The RRP model is an exponentially weighted average of all RRPs previously encountered in a road segment (excluding anomalous points), as expressed in the following equation:

$$RRPmodel(t+1,i) = K*RRPmodel(t,i) + (1-K)*RRP(t+1,i)$$

K determines the "half-life" of the model. The choice of exponential weighted average rather than some other form of average was empirical. The model is initialized to the first RRP.

Bounded-difference alignment is a technique which is less prone than conventional cross correlation alignment to cause misalignment of RRPs. This is accomplished by computing the mean bounded absolute difference between two RRPs for a variety of alignments according to the following formula:

```
find k which minimizes summation over i of
    min( KMAD, abs[ RRPmodel(t,i)-RRP(t,i+k) ] )
```

where KMAD is set to a multiple (like 1.5 or 2) of the expected mean absolute difference between RRPs which do not contain anomalies.

Masked alignment further reduces the effect of anomalies on the alignment peak by eliminating from the alignment function those points whose absolute difference exceeds a threshold similar to KMAD. Match-function peak interpolation uses parabolic interpolation to refine to sub-pixel accuracy the estimate of the alignment peak location.

Anomaly detection is accomplished by comparing the aligned RRP with the RRP model. Corresponding pixels that disagree by more than a threshold (similar to KMAD) are marked as anomalies. Anomaly detection can be done densely, so that every pixel on the road is tested against the RRP model. This differs from RRP alignment, which is performed only as needed to determine the course of the road.

Parabolic extrapolation of the locations of previous road centers is used to predict road trajectory. The trajectory prognosis is used to guide the tracker past areas where the match function peak is unsatisfactory or an inordinate portion of the road consists of anomalies.

$$x(l) = a*l^2 + b*l + c$$
$$y(l) = d*l^2 + e*l + f$$

where l is length of road path, and a-f are coefficients determined by a least-squares approximation of the road course over the few (typically six) preceding RRP alignment points.

Steps for the refined tracking algorithm are given below:

(1) Based on past road center points and directions, extrapolate the position of the road center K feet ahead, using the parabolic model.

(2) Extract the road reflectance profiles (RRP) along a line perpendicular to the direction of the road at the extrapolated center point.

(3) Use mean bounded absolute-difference alignment to determine displacement of the current RRP with respect to the RRP model.

(4) Generate a mask indicating the positions of anomalous pixels that deviate from the RRP model.

(5) Use masked alignment to locate more accurately the proper alignment of the RRP with the RRP model.

(6) Use match-function peak interpolation to refine the alignment.

(7) Detect the anomalies by comparing each pixel in the vicinity of the current RRP with the RRP model.

(8) Update the RRP model, using only the "good" points of the current RRP at the optimum alignment. Updating is done using the exponentially weighted average.

(9) Repeat steps 1-8 until the edge of the image is encountered or the RRP model becomes invalid (see the following section).

B. Experimental Results

In the experiments shown here, the road tracker was interactively started by indicating the following information for each road segment:

<X0,Y0>	center of road lane
theta0	direction of road at <X0,Y0>
w	nominal width of road

The freeway example in Figure 15 conforms well to the above road model, as shown by the overlay results in Figure 15b. The bright lines indicate the road trajectory and the bright blobs indicate potential anomalies.

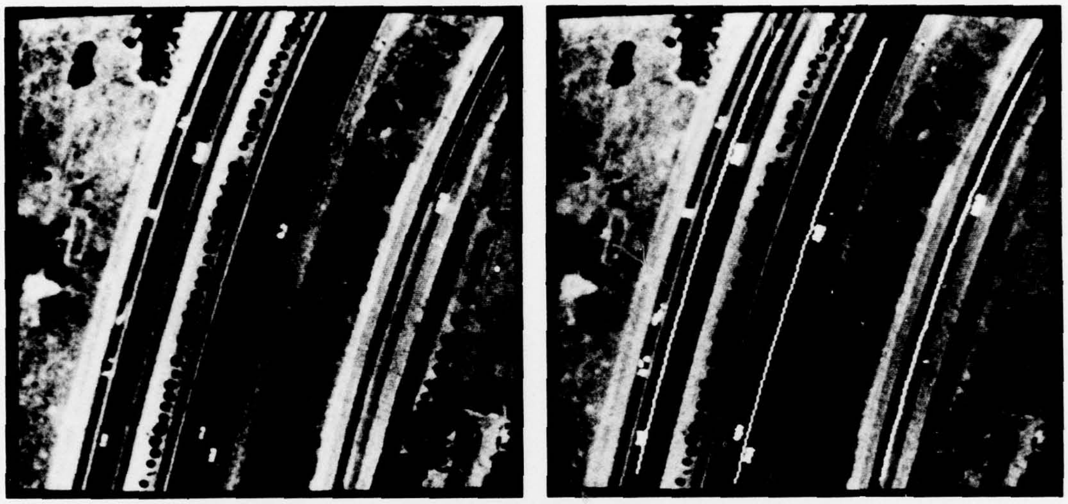

(a) (b)

FIGURE 15 A SECTION OF FREEWAY

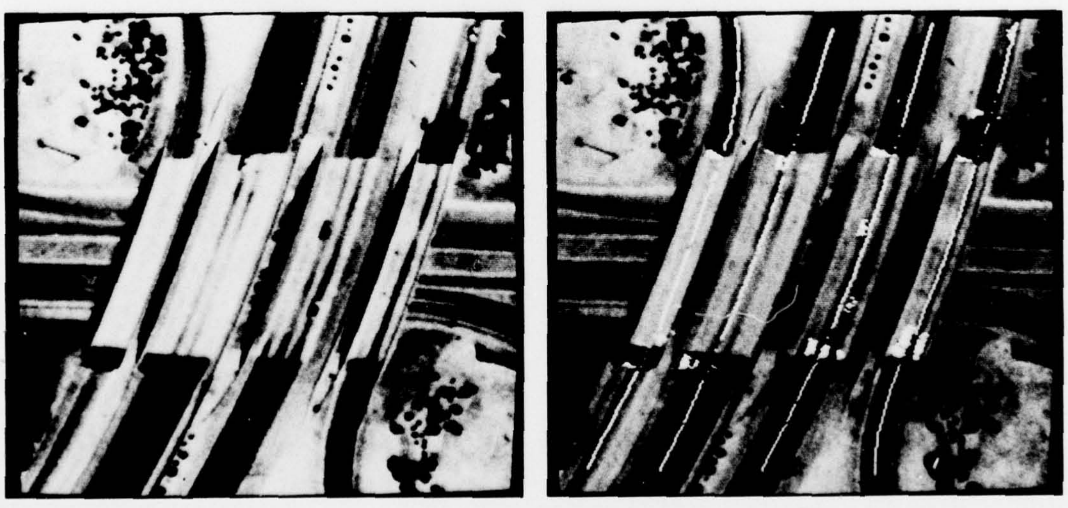

(a) (b)

FIGURE 16 A FREEWAY OVERPASS

The simplistic model in which a road consists of well-correlated reflectance profiles clearly breaks down in the example shown in Figure 16a, where the road surface changes from concrete to asphalt on the overpass. Certainly the RRP model generated for the asphalt will not match the intensities in this different road surface.

When the tracker encounters the surface change a high percentage of the pixels in the RRP will be anomalous (Figure 16b). When this occurs, the tracker extrapolates in anticipation and tries to reacquire the road. If the road is not reacquired within the length of the longest expected anomaly, the tracker then assumes that a pavement transition has taken place and establishes a new RRP model.

The four vehicles shown in Figure 16 were detected, but most of the anomalies marked therein are due to road surface changes. A later section will discuss basic changes in the control structure of the current program to eliminate the false alarms resulting from the surface changes.

Figure 17 shows results for a freeway interchange on-ramp loop. This example is interesting because the road curves rather tightly, and the road surface changes at approximately the same place where the road trajectory changes from a circular arc to a straight line.

Figure 18 illustrates a very complicated example of road forks, variation in lane width, and intersections. For the lanes tracked all vehicles and at least portions of the road surface marks (arrows and words) were found. In a developed road-expert system, the data base should help significantly in handling the complexities of this image through its knowledge of the location of forks, intersections, lane-width changes, and the like. This information will facilitate the task of interpreting causal factors in RRP model changes.

In marked contrast to the situation in most of the previous figures, Figure 19a shows a rather poorly defined dirt road with little evidence of wear patterns. Figure 19b shows the successful results of the road tracker. Most of the anomalies marked were due to shadows cast by sparsely foliaged trees.

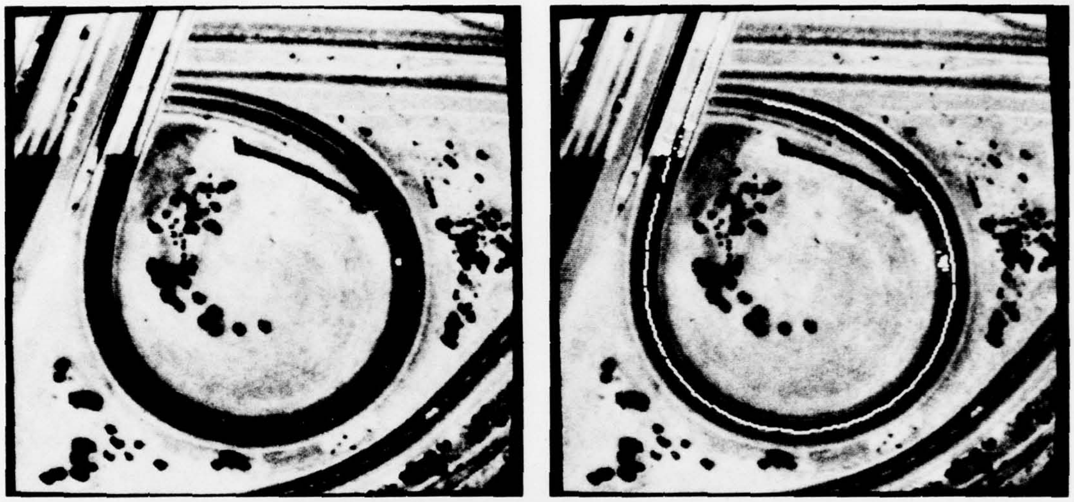

(a) (b)

FIGURE 17 AN INTERCHANGE ON-RAMP

(a) (b)

FIGURE 18 A BUSY INTERSECTION

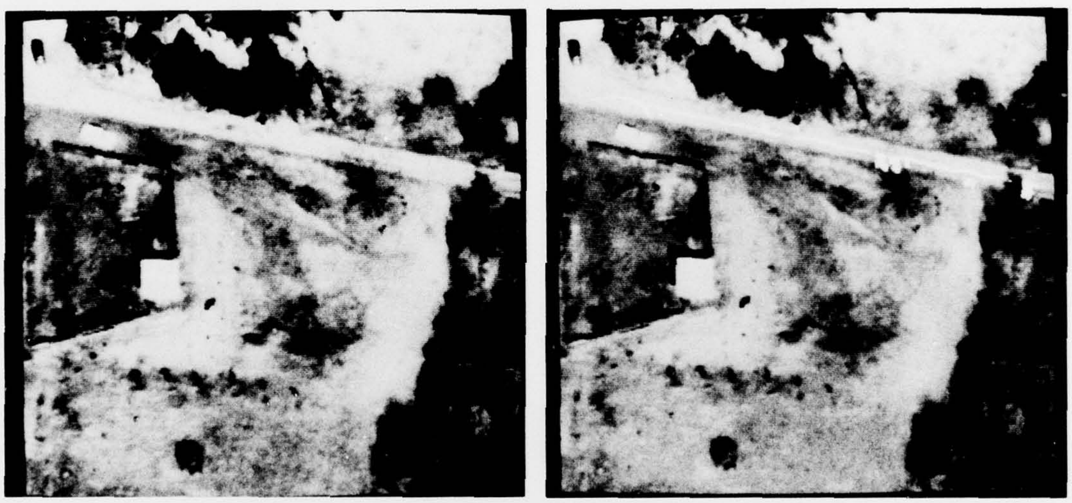

(a) (b)

FIGURE 19 A DIRT ROAD

C. Concluding Discussion

The preceding examples demonstrate both the capabilities and limitations of the present tracking algorithm. The algorithm has shown surprising ability to cope with a wide variety of road situations, including total change in the road surface. The use of bounded-difference and masked-alignment techniques nearly eliminates perturbation of the road track by anomalies. Trajectory extrapolation enables the tracker to reacquire the road after detecting that the road surface has changed. All results were obtained using the identical program and the same detection and threshold criteria; no attempt was made to "fine-tune" the parameters individually for each example.

One defect of the present algorithm is the attempt to accomplish too much in one pass along the road. In particular, anomaly marking, in the present system, begins before road-surface changes have been detected. The false alarms induced by this defect can be eliminated either by backtracking when a road transition is found, or by performing the detailed anomaly detection in a second pass along the road, utilizing the road-course and surface-change knowledge generated by the tracker.

The road tracker presently operates as an independent module. As a component of a larger road-expert system, it will be initiated from the output of a map-guided road-detection algorithm operating on lower-resolution imagery. Data-base knowledge can also be used in the tracking algorithm to increase reliability and reduce false alarms in anomaly detection. Such knowledge might consist of previous imagery of the same area or geometric information about the location of road forks, intersections, overpasses, surface changes, lane-width changes, and other parameters. To exploit such knowledge, it is necessary to establish geometric correspondence between the image and the data-base coordinate system. If, for example, a road anomaly corresponds to a known surface marking on the map or appeared in the same place in previous images, it is probably a surface marking rather than a vehicle. Similarly, the use of an illumination model can help to distinguish shadow-casting objects from surface markings.

We have acquired and digitized images taken under diverse viewing conditions with seasonal variations. We have also developed algorithms to introduce clouds and cloud shadows into these images to simulate realistic situations in which visibility is impaired and surface lighting modified. This will permit testing of the tracking algorithms under controlled conditions of nonvisibility of road segments (due to clouds) and major photometric differences between key features in images of the same area. The use of a map data base will be essential to guide the interpretation of such images; as noted earlier, however, determining how such apriori knowledge should be used is a major focus of our research effort.

With the enhancements and improvements that are planned, it should be possible to track roads and detect potential vehicles with very high hit rates and low false-alarm rates, even when operating with difficult imagery. This capability is a central component of an overall road-monitoring system.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

19

1. REPORT NUMBER
14305.1-A-EL
18 ARO

2. GOVT ACCESSION NO.

3. RECIPIENT'S CATALOG NUMBER
9

4. TITLE (and Subtitle)
6
Interactive Aids for Cartography and Photo Interpretation.

5. TYPE OF REPORT & PERIOD COVERED
Semiannual
Technical Report.
Nov 77 - 11 May 78

6. PERFORMING ORG. REPORT NUMBER

7. AUTHOR(s)
10

G. J. Agin	M. A. Fischler	L. Quam
H. G. Barrow	T. D. Garvey	J. M. Tenenbaum
R. C. Bolles	J. H. Kremers	H. C. Wolf

8. CONTRACT OR GRANT NUMBER(s)
15
DAAG29-76-C-0057,
ARPA Order-2894

9. PERFORMING ORGANIZATION NAME AND ADDRESS
SRI International
Menlo Park, California 94025
332 500

10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS

11. CONTROLLING OFFICE NAME AND ADDRESS
U. S. Army Research Office
P. O. Box 12211
Research Triangle Park, NC 27709

12. REPORT DATE
11
June [redacted]78

13. NUMBER OF PAGES
23
12 31p.

14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)

15. SECURITY CLASS. (of this report)
unclassified

15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

The findings in this report are not to be construed as an official Department of the Army position, unless so designated by other authorized documents.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

The central scientific goal of the ARPA Image Understanding Project research program at SRI International is to investigate and develop ways in which diverse sources of knowledge may be brought to bear on the problem of interpreting images. The research is concerned with specific problems that arise in processing aerial photographs for such military applications as cartography, intelligence, weapon guidance, and targeting. A key concept is the use of a generalized digital map to guide the process of image analysis. The objectives, methodology, and current

410 281

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified

20. ABSTRACT CONTINUED

status of ~~our~~ the research are described in this report. In the present phase of ~~our~~ the program, the primary focus is on developing a "road expert," whose purpose is to monitor and interpret road events in aerial imagery. ~~We~~ also described ~~herein~~ are the details of an innovative procedure for tracking road segments and finding potential vehicles in imagery of approximately 1-3 feet per pixel ground resolution. This work represents an important contribution toward~~s~~ achieving the indicated goals.