J. Alfred Stringham*
President
Land Care, Inc.
P.O. Box 381
Boonville, N.Y. 13309

Automated Imagery Processing

Requirements beyond real-time.

Introduction

Reconnaissance imagery provides the basis for a wide variety of targeting, intelligence, maps, charts, and weapon delivery materials. After joining Rome Air Development Center (RADC) in June 1957, I was engaged in developing technology to use effectively the vast array of reconnaissance and cartographic imagery collected by the U. S. Air Force. The current trends of stated and assumed requirements would seem to indicate the Defense Department would like to be able to see everything in the world on a real-time basis. Furthermore, it should be possible to display information in real-time to top commanders so they can direct individual actions, such as a rifle shot or bomb drop. Of course the technology required to do this and organize all the options of what decision-makers might view provides a challenge to the Defense R & D Establishment which will keep image processing and data handling people busy for decades while consuming a measurable percentage of the Gross National Product. If without serious thought we imagine this a desirable and achievable objective, we may fail to develop a human command and control structure under which our elected leaders can trust lower echelons to carry out policy within clearly defined limits.

Conversely, our elected leaders saturated with options and dependent on real-time control may fail to develop and explain understandable policy guidelines for lower echelons facing difficult situations. I'm afraid it will always be easier to demand more information than to create effective policy and courses of action. The unintentional delegation of these functions to the presumed wonders of electronics, communications, and data processing could create the Achilles heel of democracy.

Abstract. *In an effort to ensure timeliness, electronic imagery, and background interpretive data to military decision makers, a* Visualizer *is conceived (but not yet engineered) to aid in the process. Not only is the concept intended to be more effective than present practice, but also to make smaller demands on manpower.*

Basically, man makes plans and decisions founded on mental pictures which can be powerfully augmented by images of situations or scenes as they were, are, or might be. Timeliness is having an image of the proper scene when it fits properly into the thought process and represents a great potential for decoupled technology . . . a term I use to describe tools to augment individual action without linking them together to function effectively. A bicycle, pocket computer, or running shoes are examples of *decoupled* aides to man, whereas telephones, time-shared computers, and electronic trains represent *coupled* technology. As I recently terminated 17 years with the Defense R & D Establishment, I want to take this opportunity to reflect on some past experience and comment on where future requirements for image processing technology might be directed.

Some Observations from the Past

Aerial reconnaissance imagery provides a substantial portion of the Intelligence and targeting information for military planning

*Resigned as Chief of Rome Air Development Center's Reconnaissance and Mapping Branch on March 29, 1974. Comments and observations expressed in this article *do not* represent opinion of Rome Air Development Center, The U. S. Air Force, or the Dept. of Defense. Criticisms and comments are welcomed.

and operations. During World War II, when the art and science of aerial photography came into its own, we made significant use of imagery for everything from targeting to bomb damage assessment. Post-mortem review of World War II photography revealed approximately 25 percent of the *imaged targets* and items of intelligence interest had been detected during the interpretation process. Similar figures were derived during the Korean conflict from tactical reconnaissance imagery. We therefore succeeded in deriving 25 percent of the pertinent information imaged at manual interpretation rates of 20 to 50 square miles per hour.

During the 1960s, our collection capability made spectacular advances as we applied sensor technology to the problems of finding a shadowy enemy in the jungles of Southeast Asia. What had been a predominantly fair-weather, daytime capability was extended to night with infrared sensors and foul weather with a variety of radar sensors. The aerial camera itself was immensely improved as we went from 10 to 15 lines/mm resolution to 80 to 120 lines/mm adding the dimension of color in the process. These improvements allowed us to image tens of thousands of square miles daily with ground resolutions in inches creating an overwhelming problem for the limited staff of skilled interpreters.

The combination of tremendous quantities of high-resolution detail and an enemy skillful in camouflaging his activities probably decreased the *percentage of imaged* intelligence actually detected to what I estimate at 1 to 5 percent. I define this as the ratio between items detected and the total that would have been derived by an individual totally familiar with the area and the adversary's tactics. During our interpretation experience in Southeast Asia with photography and other imaging systems such as infrared, we repeatedly found that less than 10 percent of imaged elements had been found by operational personnel working under time-pressure with little or no collateral intelligence. I'm sure we missed many significant items through our ignorance so the ratio between actual and a theoretically *perfect* solution would be far lower.

I must add that our tactical reconnaissance/intelligence in Southeast Asia was the best in history, but note that it was achieved by a tremendous expenditure of collection resources. In addition, we began to find that much of the information being derived from high-resolution photography was too late for effective action against a mobile adversary. This led to a series of real-time sensors associated with strike systems such as the AC-47, AC-130, and AC-119 gunships. Deployment of these awesome attack systems required planning based on reconnaissance to assure acquisition of lucrative targets with their narrow field sensors such as Forward Looking Infrared (FLIR) and Low Light Level TV (LLLTV). Sending a gunship to a segment of route where a low resolution Moving Target Indicators (MTI) radar had detected moving objects was typical of this. The targets were then reacquired and struck with other sensors ranging from the human eye to FLIRs.

Our first reaction to the problem of utilizing high-resolution imagery in massive quantities was typical of pursuing an elegant technical solution to the wrong problem. During the 1960s, with a strong push from the pattern recognition technologists, many efforts were starting which investigated approaches toward automatically recognizing targets imaged on high-resolution photography in the hopes of raising the percentage of imaged targets detected a few points. In retrospect, this might be viewed as solving the problem of over-eating by developing and swallowing a machine to digest the results of over-consumption.

In general, my observation is that the important intelligence derived from high-resolution photography is not compatible with automatic detection. For example, we can design and build machines to recognize trucks on photography because we can define their characteristics somewhat precisely.

However, the most valuable information derived from photography is frequently not defined *a priori*. Our interpreters in Vietnam, for example, made tremendous use of modern photographic technology in breaking the back of the enemy seige of the Khe Sanh. In that instance, every mission was carefully analyzed by individuals concentrating on the immediate area. They began to detect subtle indications of enemy trenches, positions, and supply caches. These in turn became targets for air strikes that, among other things, defoliated the adjacent areas so more detail was imaged on subsequent missions. The team of Air Force interpreters became so familiar with the area around Khe Sanh that every movement or activity imaged had a high probability of detection. Although I would estimate the percentage of targets *imaged* around Khe Sanh that were found during interpretation exceeded 75 percent, I doubt that many significant elements would have been discovered through automatic techniques because their characteristics were never well-defined, and in many instances, individually unique.

What we have concluded is that high-resolution photography is an excellent means for finding things and clues of activity we don't know and can't define in detail. For targets that we can define (trucks, objects painted with Soviet OD, etc.) with detailed parameters *a priori*, we are better off designing sensors that collect what we need and are coupled to well-conceived exploitation systems. Of course, in the process, we might as well focus on eliminating the constraints of weather and darkness in order to observe at our option.

Although many of the first generation "real time" reconnaissance systems produce imagery (FLIR, LLTV) the trend is toward sensors like Side-Looking Radar (SLR) or MIT radars. These provide broad area coverage with the image non-existent or at most an interim step in the detection process. RADC recently demonstrated on-line real-time reduction of SLR imagery through automatic-change detection, presenting the analyst with an Order of Battle superimposed on a map, not high-resolution imagery. Where we are investigating augmenting inflight detection of selected targets, we are exploring the design of the sensor-exploitation package as a unit. This technology is currently being pursued on a cooperative basis between RADC (exploitation) and Air Force Avionics Laboratory (sensor) and should eventually lead to a new generation of real-time target acquisition-strike systems.

An important aspect of real-time operations which we tend to overlook is the decision-making cycle. In talking about real-time weapon delivery, one should remember that many *a priori* decisions have already been made—such as evaluating the cost of error—will we hit our own troops or allies? Can we afford to expend weapons on decoys, etc.? Therefore, not only are the targets to be detected and struck well-defined, but also the background conditions thoroughly restricted to certain areas, trails, roads, etc. Simplistic augments on the problem of manually interpreting vast quantities of high-resolution detail as justification for automatic image analysis really lead to elegant technical solutions to the wrong problem.

Our experience in improving the exploitation of high resolution photography leads in a different direction. Because, as I mentioned, the clues are frequently subtle, everything possible to augment the interpreter's knowledge of the area or problem is beneficial. Application of modern data-processing technology has been a tremendous aid in this area by giving interpreters on-line access to the vast quantities of collateral intelligence that make clues imaged on high-resolution sensors meaningful. In addition, this has allowed us to manage much more effectively our collection so that we get only what we need.

In the past, we frequently flew repeated missions because there was no convenient means of knowing what we had collected in the past or what we had or should have learned from it. Under the direction of one of my RADC associates, Bill Moore, the Air Force achieved a tremendous improvement in its ability to exploit imagery rapidly and fully with a program known as PACER (Program Assisted Console Evaluation and Review). This was a combination of computer hardware and software giving the interpreter on-line access to a tremendous quantity of collateral information. Furthermore, the interpreter's results were immediately available to collection managers and significantly reduced the cycle time. This has in turn actually *reduced* the amount of imagery required although improving the quality of the information produced.

The Tactical Information Processing and Interpretating System (TIPI) which grew out of a RADC program will provide much of this capability to the Tactical Reconnaissance Forces. It is important to note that neither of these capabilities involves any form of replacing the interpreter's function of perceiving and identifying targets or subtle clues to enemy operations. They merely augment his access to important relevant information and make his routine image-management functions more efficient. For example, although the thought of missing 75 percent of the targets imaged on photography is disturbing, the fact that tactical interpreters frequently spent 40 to 60 percent of the allocated time with a reconnaissance mission plotting and filling out paperwork should be horrifying.

I'm sure many of you are aware of the fact that Air Force funding of R & D for Automatic Imagery Pattern Recognition has dropped significantly during the past four years. We are now focusing on different image processing problems. Little work is addressing the question of automatically finding targets on extremely high-resolution imagery based on spatial characteristics, but rather is being focused on interactive imagery analysis concepts. The automation efforts are directed toward sorting targets (carefully defined) from background on the basis of combined collection-exploitation systems, using spectral/spatial and collateral factors. Although our support of attempts to detect or automatically recognize numerous classes of targets on high-resolution photography is rapidly converging on zero, R & D to augment the interpreter in numerous other ways continue.

One very powerful tool recently demonstrated under RADC's sponsorship by my friend Bob Lillestrand is the ability digitally to match or warp imagery of various types into each other extremely rapidly. This capability, that can be applied at collection rates, allows to normalize new imagery in many ways and integrate it with collateral materials such as digital cartographic data. Although it's possible to design equipment able to screen large quantities of imagery automatically for activity clues, the trend is not to collect such imagery in the first place. Imagery stores a permanent picture of the past, and many areas are imaged repeatedly. Appropriate comparison gives us a chronological history of an area providing important clues to the present and future activity. The problem is to develop convenient means of accessing and utilizing this *image history* during exploitation process. The ability to match and normalize imagery from different sensors and orientation automatically provides a great potential in providing a rapid *time-history* of an area or scene.

Imagery frequently contains information highly dependent on the knowledge of the observer. In some instances, developing improved mechanisms to distribute high-resolution pictures of a new enemy fighter to a skilled aeronautical designer is more important than any device or technique which we can provide the conventional image interpreter. I have outlined a possible concept to implement this process in the final section. Those things that automate imagery management functions (plotting both coverage and individual targets, quality assessment, cloud cover determination, etc.) make possible a distribution system appropriate for decentralized use.

Everyone can point to examples where an automatic pattern-recognition device might have been of great value. But these examples usually cover a period when we were frantically trying to exploit the results of over-collection. We currently plan to collect only what we need by using a different series of wide-area detection sensors to focus our high-resolution systems, such as photography and modern infrared scanners. By collecting less and augmenting the human interpreter-analysts we are digesting more, although using fewer resources. Our basic objective is to solve the problem of overweight by consuming only what is needed, not automating the digestive process. Unfortunately, organizations as complex as the Air Force don't have the marvelous mechanism of the human body to establish exactly what is essential to survive.

However, our experience during the past decade has gone a long way toward making it possible to spend less of our allocated resources on imagery collection while substantially improving the quality and timeliness of our intelligence and targeting information. The Air Force reconnaissance community is constantly reviewing ways to establish optimum timeliness and quality of information available to military commanders. Complete, efficient exploitation of carefully managed imagery collection programs is a key element in accomplishing this. However, the means of getting the image or information to the planners, decision-makers, and operators at the right time is another area requiring careful thought.

Some Thoughts for the Future

The theme of this panel is to identify User Requirements as related to Automated Imagery Processing Technology. The Defense Department in general, and the Air Force in particular, have been and will be major collectors and consumers of imagery. Perhaps we should differentiate the ultimate consumers of imagery data from the brokers. Developers of imagery-processing technology frequently view the interpreters as the user and press for requirements. This is tantamount to asking a coal miner to identify the requirements for efficiency using coal. As a consumer for heating, etc., a miner's opinion will probably be no better or worse than the average citizen's, unless, of course, one wants to develop mining equipment—not more effective use of coal for energy, etc. In reality, image interpreters and photogrammetrists, cartographers, and our vast variety of intelligence analysts are extractors of information for planners, operators, and decision-makers. In many instances, the future thrust of image processing should be to bypass their function—and provide useable image products to a vast array of actual users.

Before proceeding, it's perhaps worth reviewing another technology that provides analogous lessons to this area: automated computation. When I began work at RADC, electronic digital computers were just being developed and I had a large vacuum-tube prototype beside my desk. Today I carry more computational power in my pocket.

During the early and mid sixties, we passed through a period when numerous centralized computer facilities were set up to serve the needs of scientists, engineers, accountants, etc. With these facilities, came legions of specialists—programmers, system analysts, keypunch operators, etc., who interfaced centralized computing power with users. Today mini- and micro-mini-computers have reached the stage where the man with computation problems can perform them himself except for large linked efforts like orbital calculations.

Currently, our imagery exploitation facilities are becoming an expensive combination of specialized viewing equipment and collateral data base and computational support (Figure 1). Ultimately, annotated images of intelligence and targeting value are displayed for decision on briefing boards, viewgraphs, etc. The next aspect is the matter of timeliness of data. Current requirements often indicate the ultimate is to show the target area as it is now to the decision maker or planner. This of course requires real-time sensors, data links, etc., and obviously a monumental investment in collection capability. In reality, my experience would indicate *timeliness* is often related more to giving the user an image of the scene if it fits appropriately into his continuum of activity even if based on dated imagery weeks, months, or years old.

Perhaps we should review the actual functions interpreters and intelligence specialists usually perform. First, they plot the coverage or certify the automated plot along with an assessment of quality. Second, they answer some predetermined questions ranging from counting various types of aircraft on a field to high-lighting undefined possibly significant activity. Their textual reports are augmented pictorially so that the users can see the situation for themselves. In some instances, our user, because of specialized knowledge, will derive information from an image interpreters missed. For example, an aeronautical engineer will tell things about an imaged aircraft performance characteristics which no normal interpreter can. Third, they may compile or update maps, produce radar predictions, or produce a variety of other image related products.

All of this says that one should perhaps strive to decentralize the use of imagery and put into the hands of planners, decision-makers, etc., those tools that will allow them easily to reformat reconnaissance and mapping products for the problem at hand. Perhaps we need to develop a small, relatively inexpensive *Visionizer* (Figure 2) which could be carried in the user's pocket like our small electronic digital computers. Input to the *Visionizer* would be digital imagery and cartographic information on cassettes or some form of storage that could be transformed and displayed at varying scales, levels of detail, contrast, etc.

The *Visionizer* would allow an individual to view, say 4×10^6 spatial elements representing maps and spatial images of places and things. Attached would be a computer similar to those currently in use. The user could put in a code for the scene coordinates, perspective, desired scale, detail, etc., and get a spatial picture. It is my contention that an image available during the planning or decision cycle when the individuals actually want it will be far more valuable than *real-time* information fit into the planning-decision cycle on an "as available or cost effective management" basis. The input to a user's *Visionizer* could be updated periodically so that it always contained the latest imagery available consistent with the needs and priority of the user.

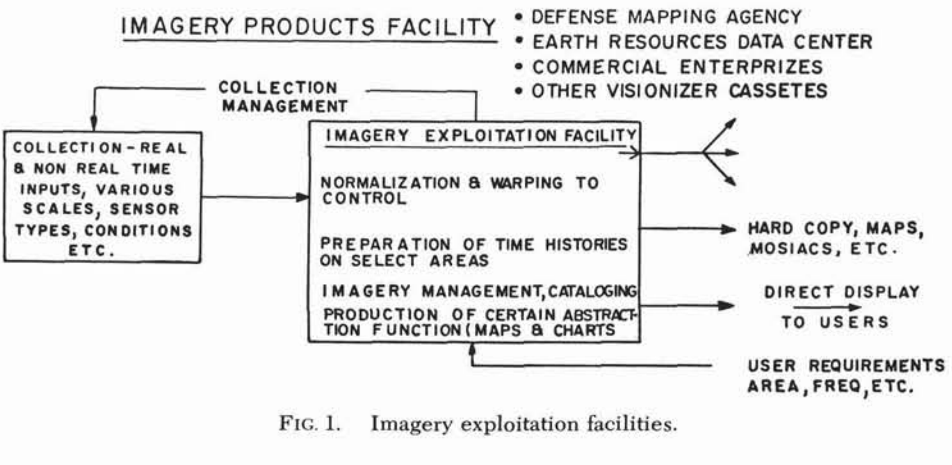

FIG. 1. Imagery exploitation facilities.

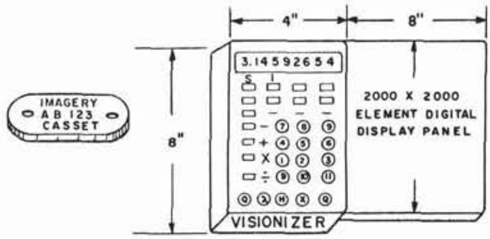

FIG. 2. Individualized pocket *Visionizer*.

Needless to say, this conceptual *Visionizer* assumes a good deal of technology not currently available, such as preparing input (warping, normalizing, abstracting, etc.), high density, small volume store (1×10^{12} bits on a cassette) and foldout-panel high-resolution (color(?)) display technology.

The concept of the *Visionizer* is that a decoupled technology to augment individual planning, decision and action—not tie him continuous to a centralized image exploitation facility. One can imagine such a device in use by pilots, field army personnel, high-level decision-makers, low-level planners, and numerous private taxpaying citizens. The variety and frequency of update would vary but, in essence, the definition of interpreter would change from one who transforms reconnaissance into words, view-graphs, and briefing aids, to anyone who uses imagery or digital maps to plan, move, target, fly, shoot, analyze, or decide.

In essence, my message is that the value of imagery perhaps depends more on astute distribution in a form or forms which will have wide application rather than automatically deriving reams of information that nobody really needs. The concept of the *Visionizer* and the technology required to reduce a massive quantity and variety of imagery to easily stored and carried cassettes could represent augmentation of what man envisions.

PROLOGUE TO THE FUTURE (?)

In conclusion, before you enterprisers rush to bombard the Defense R & D Establishment with proposals to study, investigate, develop, and design technology for the *Visionizer*, or your better schemes to accomplish the function, stop and think. A huge commercial *Visionizer* market exists among customers ranging from hunters to real estate agents. Challenge your management or private investors for venture capital —innovate—produce an effective, affordable product at the cost of today's HP-45 pocket computer. Five years from now I hope to buy one on the open market. If I do, DOD will be buying them by the thousands, irrespective of whether the inevitable Interdepartmental Committee on Standardized, Reliable, Maintainable, Compatible *Visionizer* specs are met or not. In conclusion, I believe the *Visionizer* concept has tremendous military and commercial potential. The major Defense Department requirement is for enterprising men and women who can present them capability as an off-shelf item. Good luck enterprisers, I'm betting on you.