

R-2057/2-PR
November 1976

Image Correlation: Part II Theoretical Basis

H. W. Wessely

A report prepared for
UNITED STATES AIR FORCE PROJECT RAND

Rand
SANTA MONICA, CA. 90406

The research described in this report was sponsored by the United States Air Force under Contract No. F49620-77-C-0023— Monitored by the Directorate of Planning, Programming and Analysis, Deputy Chief of Staff, Research and Development, Hq USAF.

Reports of The Rand Corporation do not necessarily reflect the opinions or policies of the sponsors of Rand research.

R-2057/2-PR
November 1976

Image Correlation: Part II Theoretical Basis

H. W. Wessely

A report prepared for
UNITED STATES AIR FORCE PROJECT RAND

Rand
SANTA MONICA, CA. 90406

PREFACE

This is the second part of a two-part report that describes research initiated at Rand in mid-1975 under the Project RAND research project "Target Acquisition." The subject is "map matching" or image correlation to achieve autonomous target acquisition and terminal guidance for missiles (both strategic and tactical), with particular emphasis on the acquisition phase.

Part I* presents an analysis of the probabilities of correct and false acquisitions, extends it to include the effects of a number of common error sources, and describes computer simulations based on data samples from real scenes. Part II provides a more general and more rigorous analytical approach. Some of the conclusions derive jointly from both phases of the study, but Part II is published separately because it is addressed to readers with a theoretical and mathematical interest in the subject.

Both reports should be of interest to defense and industrial project managers and engineers involved in the development of missile guidance, particularly those concerned with current or future correlator programs.

* H. H. Bailey, F. W. Blackwell, C. L. Lowery, and J. A. Ratkovic, *Image Correlation, Part I: Simulation and Analysis*, The Rand Corporation, R-2057/1-PR, November 1976.



SUMMARY

Image correlation or "map matching" makes possible a type of weapon guidance that provides autonomous target acquisition and tracking. This study analyzes the image correlation process, both theoretically and by using computerized simulations; primary emphasis is on the often neglected but crucial acquisition phase. (The requirement for achieving adequate terminal tracking accuracy in weapon delivery has been and is being studied extensively elsewhere and, for this report, is considered to be of secondary importance; it is simply assumed that, if necessary, operational systems could accommodate a software change to maximize tracking accuracy *after* the initial acquisition has been accomplished.)

The essential step in image correlation guidance is to find the position of "best fit" between two similar but nonidentical images or "maps": a *sensor* image of the terrain surrounding a desired target, obtained in real time as the weapon approaches, and a previously prepared *reference* image of roughly the same area. The match point is found by systematically displacing one map relative to the other and computing, for each of the many possible displacements, the value of a comparison function or "metric" that, ideally, has an extremum (max or min) value at the match point. The particular displacement, suitably scaled, that produces the extremum becomes the correction signal for the guidance system.

Unfortunately, precisely because the two maps are not identical--owing to detector noise, real changes in the scene, geometrical distortions, and several other causes that are discussed in Part I^{*} of this report--the displacement that produces the extremum does not *always* correspond to the correct match point. It only does so *on the average*. Accordingly, this analysis of Part I focuses on two topics of principal concern: (a) the probability of achieving a correct match (conversely,

* H. H. Bailey, F. W. Blackwell, C. L. Lowery, and J. A. Ratkovic, *Image Correlation, Part I: Simulation and Analysis*, The Rand Corporation, R-2057/1-PR, November 1976.

the probability of a "false lock" or, in military terms, a gross error) and (b) the selection of an appropriate comparison metric to maximize (a).

Part I of this report describes the results of a direct approach to determining the probability of a correct match, P_c , first for random scenes with Gaussian statistics, and then for real scenes in the presence of noise and various errors. Two commonly used comparison metrics are calculated by using the so-called Product (a sum of products that is related to classical correlation) and MAD (mean absolute difference) algorithms, respectively. It is shown that in all cases P_c increases with the size of the data sample and with the elemental signal-to-noise ratio (S/N), and decreases (slowly) with increasing search area. It is also shown that at low S/N, the Product algorithm is the preferred one (i.e., it leads to higher probabilities of correct lock), but at high S/N, the MAD algorithm is preferred. However, when geometrical errors, such as synchronization (an effect peculiar to digital systems in which the cells of the two maps are staggered by some unknown fraction of a picture element), rotation, and scale factor (magnification), are present, the effective value of S/N is significantly reduced. Thus the higher values that would render the MAD or similar algorithms attractive are seldom realized in practice.

The fundamental nature of the map-matching problem is reexamined in this portion of the study (Part II), and the degree of theoretical justification for the use of the various comparison metrics is also investigated. Since the problem is basically one in statistical decision theory, it is shown that the optimum solution is achieved by computing the likelihood ratio for each comparison and then choosing the match point at the place where the likelihood ratio is maximum. Unfortunately, that computation requires a knowledge of the N-dimensional joint probability distributions--functions that are unknown and, in a practical sense, unmeasurable. Hence, one must resort to approximations. These usually take the form of maximizing or minimizing one of several functions, herein called "metrics." In much current work, these are chosen almost arbitrarily and therefore must be subjected to essentially experimental validation. By considering two-picture-element scenes, such that

the likelihood ratio and several of the commonly used metrics can be expressed in simple algebraic form and discussed in geometrical terms, the essential features of the various metrics are explained and compared with the likelihood ratio. In this way heuristic arguments are developed that support the use of the Product algorithm when S/N is low and the MAD algorithm when S/N is high. The notion, borrowed from signal detection theory, that the classical correlator (i.e., the Product algorithm) is optimum for this application is shown to be erroneous; the two problems are fundamentally different.


Two major conclusions were derived from this study. First, on the basis of the empirical results of the simulation studies described in Part I, it appears that the theoretically predicted values of P_c for random scenes with Gaussian statistics can be considered as an approximate lower bound for the probability of acquisition in a real application. The quantitative relationships between P_c and various system parameters that have been derived for random scenes, and that have been largely confirmed by simulation testing with real scenes, can therefore be used to carry out various design tradeoffs, including a balancing of the costs of a tighter overall P_c requirement with the loss of those weapons that fail to acquire. The theoretical model of the random Gaussian scene is known to be not completely realistic, but it appears to err on the conservative side. Thus, a "floor" for P_c can be established, which should permit the flight test performance of future properly designed systems to be somewhat better (i.e., to exhibit fewer gross errors) than is predicted by the theory.

The second conclusion is that there ought to be better algorithms than those that have usually been used in the past. Since (a) there is at present little theoretical basis for the commonly used comparison metrics, and (b) most real terrain contains features beyond those describable by simple Gaussian statistics, it seems both reasonable and not inconsistent with theory to search for more efficient ways to carry out the initial map-matching or target-acquisition function. In particular, drastic preprocessing to extract special features of a given scene, using techniques currently being developed and exploited in the field of pattern recognition, should lead to more efficient algorithms.



CONTENTS

PREFACE	iii
SUMMARY	v
Section	
I. INTRODUCTION	1
II. STATISTICAL DECISION THEORY	2
Extremal Metrics	3
Heuristic Interpretation of the Likelihood Ratio ...	5
Likelihood Ratio for Gaussian Distributions	9
III. A COMPARISON OF THE LIKELIHOOD RATIO FOR GAUSSIAN DISTRIBUTIONS WITH OTHER METRICS	13
The Normalized Inner Product (NProd) Metric	13
The Difference Squared Metric	15
The Mean Absolute Difference (MAD) Metric	15
Probability of Correct Acquisition	15
IV. SUMMARY AND CONCLUSIONS	18
Appendix	
LIKELIHOOD RATIO FOR CORRELATED BIVARIATE GAUSSIAN DISTRIBUTIONS	19



I. INTRODUCTION

Map matching is fundamentally a problem in statistical decision theory. During the acquisition phase it is necessary to decide whether or not the sensor and reference maps are matched. But because of noise, errors, distortions, etc., possibly in each map, there is no foolproof way to make certain that the maps have been matched correctly. The most that can be done is to determine the probability that the maps are in a certain geometrical relationship to each other based on the available data. The optimum map-matching system, therefore, is by definition the system whose output is the set of a posteriori probabilities describing each possible relationship. Once these probabilities are determined and costs are assigned to each kind of wrong decision, the decision rule is implicitly defined by the requirement that some measure of the total cost, e.g., the average total cost of a decision, be a minimum. The cost assignment and the details of the resulting decision rule do not affect the underlying structure of the map-matching system. The decision rule, whatever its form, is necessarily based on the values of the a posteriori probabilities. Hence, the primary task is to determine these probabilities.

II. STATISTICAL DECISION THEORY

The basic relationship between a priori and a posteriori probabilities is given by Bayes' formula. For simplicity of exposition, suppose that only two possibilities are of interest, namely that the maps are, or are not, matched. Let these two conditions be denoted by the symbols S (signal) and B (background), respectively. If a sensor map consisting of N data elements is represented by the N-dimensional vector x , then Bayes' formula for the matched hypothesis states that

$$P(S|x)P(x) = P(x|S)P(S) , \quad (1)$$

where $P(S|x)$ = the a posteriori conditional probability that the maps are matched, given that the sensor map has the value x ,

$P(x)$ = the a priori probability that the sensor map has the value x ,

$P(x|S)$ = the conditional probability that the sensor map has the value x , given that the maps are matched,

$P(S)$ = the a priori probability that the maps are matched.

An analogous statement with S replaced by B holds for the unmatched hypothesis. Since there are only two possibilities, the probability that the sensor map has the value x is simply

$$P(x) = P(x|S)P(S) + P(x|B)P(B) . \quad (2)$$

The likelihood ratio $L(x)$ is defined by

$$L(x) = P(x|S)/P(x|B) , \quad (3)$$

so that, after substituting (2) and (3) into (1), the a posteriori probability of a match can be written as

$$P(S|x) = \frac{1}{1 + \frac{P(B)/P(S)}{L(x)}} . \quad (4)$$

Since the probability that the map assumes precisely the value x is usually zero, it is to be understood that the indeterminate form of the likelihood ratio is to be replaced by the ratio of the respective probability densities $p(x|S)$ and $p(x|B)$. The a priori probabilities can be considered to have known constant values that describe the estimates of the position of the flight vehicle and the direction of the sensor's line of sight when a sensor map is obtained. Since the a posteriori probability is a simple monotonic function of the likelihood ratio, it is clear that a knowledge of the likelihood ratio's value is equivalent to a knowledge of the a posteriori probability. Hence, the optimum map-matching system is nothing more than the system that computes the value of the likelihood ratio for each comparison between a sensor map and a reference map.

In a purely formal sense, Eq. (4) represents a complete solution to the problem. The essential difficulty in applying Eq. (4) to a practical situation, however, is that the N-dimensional joint probability functions needed to form the likelihood ratio are usually unknown. A typical sensor map may consist of tens to thousands of spatially correlated element values with a non-Gaussian joint probability distribution. In the absence of a theoretical framework to describe the form of the distribution, the only recourse for determining the distribution is direct measurement, which, for maps containing even as few as three or four elements, is a practical impossibility.

Because the probability distributions needed to form the likelihood ratio are generally unknown, and, in a practical sense, essentially unmeasurable, another approach to map matching is required. The approach most frequently encountered in the literature is one based on the extremals of a metric.*

EXTREMAL METRICS

A metric F is a class of pairs of the form $(x, F(x))$, where x is

*The term "metric" as used here simply refers to a scalar function of a vector. There is no requirement that the metric satisfy the so-called triangle inequality as in the study of functions of a real variable.

a vector and $F(x)$ is a number. The first member of the pair is often called the object and the second member, its value. Many metrics have extremal properties, i.e., for some particular vector, the value of the metric is either a maximum or a minimum. Familiar examples of such metrics are given below.

The Normalized Inner Product (NProd) Metric

The normalized inner Product, sometimes called the normalized Product (NProd), is defined by

$$F(x) = \frac{x \cdot y}{\|x\| \|y\|} , \quad (5)$$

where y is some given vector, $x \cdot y$ denotes the inner Product, and $\| \cdot \|$ denotes the norm of a vector defined by

$$\|x\| = (x \cdot x)^{\frac{1}{2}} \quad \text{and} \quad \|y\| = (y \cdot y)^{\frac{1}{2}} . \quad (6)$$

The extremal property of this metric follows from the familiar Cauchy-Schwartz inequality. Thus, $|F(x)| \leq 1$ for all vectors $x \neq y$, with equality occurring only when $x = Cy$, where C is an arbitrary constant.

The Difference Squared Metric

This metric is defined by

$$F(x) = \|x - y\|^2 , \quad (7)$$

where y is some given vector. Thus, $F(x) \geq 0$ for $x \neq y$, with equality only when $x = y$.

The Mean Absolute Difference (MAD) Metric

This metric is defined by

$$F(x) = \sum |x_i - y_i| , \quad (8)$$

where x_i and y_i denote i th components of the vectors x and y , respectively, and the sum extends over all components. As with the difference squared, $F(x) \geq 0$ for $x \neq y$, with equality only when $x = y$.

If the vectors x and y are associated with the sensor and the reference maps, respectively, then a physically reasonable basis for attempting to adapt such metrics to the problem of map matching immediately suggests itself. If it is assumed that when the sensor's line of sight (i.e., the map center) coincides with the target, the sensor map "most nearly" resembles the reference map, then it is reasonable to assume that the value of the metric in that case will be "closest" to the extremal value. Because of noise and other distortions, however, a perfect match is never to be expected. Thus, some acceptance threshold different from the extremal value is implied. If it is known beforehand that the sensor map is contained within the ensemble of maps available for comparison, i.e., the reference map, the decision rule may simply decide that the extremal-producing map defines the match point. If the sensor map is not known to be contained within the reference map, the decision rule may require the value of the extremal to be above some prescribed threshold value.

The obvious question that arises when such metrics are used is, "What metric should one use? or, stated differently, How is an optimum or at least a "good" metric selected? The answer clearly depends on the map statistics--which have been completely ignored in the above definitions of metrics. Any attempt at an analytical formulation of this question leads directly back to the original statement of the map-matching problem as a problem in statistical decision theory. The optimum metric is the likelihood ratio defined by Eq. (3), since it essentially defines the most that can be known, namely the a posteriori probabilities that the maps are matched given the available data. Thus, if a completely ad hoc approach to map matching is to be avoided, some sort of theoretical analysis based on the likelihood ratio must be undertaken.

HEURISTIC INTERPRETATION OF THE LIKELIHOOD RATIO

Some insight as to what constitutes a good metric can be obtained if one considers the likelihood ratio for the extremely simple case of

maps that consist of only two elements. This simplification permits the use of elementary geometrical concepts to illustrate the essential features of various metrics. The extension of the concepts to larger maps is straightforward.

A single two-element map is thus represented by a pair of values (x_1, x_2) . The ensemble of possible sensor maps for the unmatched condition is then described by its joint probability density function $p(x_1, x_2|B)$. For simplicity, let the mean value of this distribution be assumed zero. The two values comprising a sensor map will generally be correlated, so that the intersection of a horizontal plane with the density function will produce a somewhat oval-shaped curve, as shown by the dashed curve of Fig. 1. (The exact shape of this curve of




Fig. 1 — Schematic representation of the map-matching problem

intersection depends, of course, on the particular form of the joint probability density function.) Thus, the maps within this boundary contour occur with a certain frequency determined by the height of the intersecting plane. If the map values are highly correlated, the oval tends to be narrow; if the correlation is weak, the oval tends toward a circle.

Suppose for the moment that no errors of any kind are present in either map. If the values associated with the sensor map are (s_1, s_2) for the matched condition, the joint probability density function of the sensor's output is a two-dimensional Dirac delta function, namely,

$$p(x|S) = p(x_1, x_2|S) = \delta(x_1 - s_1)\delta(x_2 - s_2) . \quad (9)$$

The likelihood ratio in this simple case is then

$$L(x_1, x_2) = \frac{\delta(x_1 - s_1)\delta(x_2 - s_2)}{p(x_1, x_2|B)} . \quad (10)$$

Thus, if the joint probability density function $p(x_1, x_2|B)$ is well behaved, the likelihood ratio itself has the form of a delta function. The decision rule, which identifies the match point with the maximum value of the likelihood ratio, then tells us to decide that a match is present whenever precisely the pair (s_1, s_2) occurs. The rule, of course, simply expresses the obvious. We know exactly what to expect when a match is present. Although the occurrence of the pair may not be an unambiguous indication of a particular match point, we have no means for refining our decision. Hence, the best we can do is to assume that a match is present. The situation is completely analogous to finding a specific two-digit number in a table of numbers, for example. If we know that the table is accurate, we decide that we have found the number when precisely the reference two-digit pair occurs. In the absence of other information, every such pair must be considered a match point.

Now suppose that both the sensor and reference maps are subject to error because of noise, distortions, etc. For simplicity, suppose further that all of these errors have zero mean value, are statistically independent of the uncorrupted map values, and can be considered to be additive in their effect. If all of these effects are simply called "noise," with a joint probability density function $p(x_1, x_2 | N)$, then, if the noise is uncorrelated and has a zero mean value, the intersection of the noise distribution with a horizontal plane will be a circle, as shown at the origin in Fig. 1. With these assumptions, the joint probability density function $p(x_1, x_2 | BN)$ of the corrupted sensor maps for the unmatched condition is $p_{BN} = p_N * p_B$, where the asterisk denotes the convolution operation and where the arguments of the functions have been dropped and the conditions incorporated into subscripts to simplify writing. Similarly, the joint probability density function $p(x_1, x_2 | SN)$ of the sensor's output when the maps are matched is $p_{SN} = p_N * p_S = p_N(x_1 - s_1, x_2 - s_2 | N)$. Therefore, horizontal cuts through these two probability density functions result in contour curves somewhat like the solid curves of Fig. 1 that are centered on the origin and the point (s_1, s_2) .

If the noise is small compared with the variance of the uncorrupted sensor map values, the maximum value of the likelihood ratio will occur in the neighborhood of the reference map values (s_1, s_2) . A horizontal cut through the likelihood ratio near its maximum value will then result in a closed curve somewhat like that shown bounding the shaded area in Fig. 1. The occurrence of a pair of map values close to the reference values is then an indication of a probable match. Thus, for the assumptions stated above, a metric that defines a family of closed contours about the reference map values qualitatively resembles the likelihood ratio. It is to be emphasized, however, that the assumption that all the errors affecting the maps can be lumped together into a single, statistically independent, probability density function with random noise-like properties is a simplification made here for illustrative purposes only.

LIKELIHOOD RATIO FOR GAUSSIAN DISTRIBUTIONS

If it is assumed that all the distributions have a Gaussian form, then an explicit representation of the likelihood ratio is obtained by simple substitution. Let it therefore be assumed that the ensemble of possible sensor map values for the unmatched condition has the joint probability density

$$p(x_1, x_2 | B) = \frac{\exp \left[-(x_1^2 - 2r_B x_1 x_2 + x_2^2) / 2\sigma_B^2 (1 - r_B^2) \right]}{2\pi\sigma_B^2 (1 - r_B^2)^{1/2}} \quad (11)$$

where σ_B is the root-mean-square variation of the uncorrupted sensor map values and r_B is the correlation between the values x_1 and x_2 . Similarly, let it be assumed that

$$p(x_1, x_2 | N) = \frac{\exp \left[-(x_1^2 + x_2^2) / 2\sigma_N^2 \right]}{2\pi\sigma_N^2} \quad (12)$$

where σ_N is the root-mean-square variation of the noise. Then by forming p_{SN} and p_{BN} and substituting into Eq. (3), one can show that the logarithm of the likelihood ratio is proportional to the quadratic form

$$F(x_1, x_2) = \sigma_N^2 (x_1^2 - 2r_{BN} x_1 x_2 + x_2^2) - \sigma_{BN}^2 (1 - r_{BN}^2) \left[(x_1 - s_1)^2 + (x_2 - s_2)^2 \right] \quad (13)$$

where

$$\sigma_{BN}^2 = \sigma_B^2 + \sigma_N^2 \quad (14)$$

and

$$r_{BN} = \frac{\sigma_B^2}{\sigma_{BN}^2} r_B \quad (15)$$

Since in practice the correlation is greater than zero and less than one, one can show by elementary analytical geometry that this equation defines an ellipse with axes rotated 45 deg with respect to the x_1, x_2 axes. The center of the ellipse is displaced from the uncorrupted map values (s_1, s_2) , but the displacement tends to zero as the ratio of σ_N/σ_B tends to zero, i.e., at high values of the signal-to-noise ratio. When the correlation is zero, the ellipse reduces to a circle. The likelihood ratio thus assumes its largest values when the *ratio* of the map values is approximately in the ratio of the reference map values, and when the values of the sensor map are separately approximately equal to their corresponding reference map values.

When the background variation is small in comparison with the noise, the sensor's output consists essentially of either the desired signal plus noise, or noise alone. This situation is depicted in Fig. 2. The problem of discriminating between these two cases is the classical problem of *signal detection*. Its solution is the well-known matched filter that, for the case of Gaussian noise, is the classical correlator, i.e., the unnormalized product metric. This result follows immediately from the likelihood ratio (see Eq. (13)) when σ_B and r_B are set equal to zero. In that case, the ellipse defined by the likelihood ratio degenerates into the straight line defined by

$$F(x_1, x_2) = s_1 x_1 + s_2 x_2 . \quad (16)$$

The problem of map matching has sometimes been confused with the problem of signal detection. As a result, the solution to the signal-detection problem, i.e., the classical correlator, has sometimes been erroneously interpreted as the theoretically optimum solution to the map-matching problem. The starting point in many analyses of map matching has therefore been the calculation of the classical correlation function. The fundamental difference between the two problems, however, is readily apparent when Figs. 1 and 2 are compared. In the case of map matching, one needs to distinguish between signal plus noise and background plus noise. In the case of signal detection, it is necessary to distinguish between signal plus noise and noise alone.




Fig. 2—Schematic representation of the signal-detection problem

The solution of either problem lies in the calculation of the likelihood ratio, which, for additive Gaussian noise, takes the form of Eq. (13) and the shaded ellipse of Fig. 1 for map matching and of Eq. (16) and the straight line of Fig. 2 for signal detection.

The more general case in which the noise associated with the matched condition is different from that for the unmatched condition is discussed in the appendix. Noise in both cases is often partly due to geometrical misalignments in angle or scale. In such cases, the statistics of the sensor's map for the matched condition will strongly depend on the properties of the map in the vicinity of the nominal

match point. It is shown in the appendix that when the conic is non-degenerate, the locus of the likelihood ratio can be either an ellipse, hyperbola, or parabola, depending on the relative variances of the sum and difference of the pixel values for the matched and unmatched conditions. In most cases of practical interest, however, neither variance for the matched condition is expected to exceed its corresponding variance for the unmatched condition. Thus, as the analysis presented in the appendix shows, the locus of the likelihood ratio is always an ellipse, which agrees with the simplified discussion above.

In summary, therefore, we have the following two extremes: when the signal-to-noise ratio is very large, the likelihood ratio reduces to a delta function, so that the optimum metric is one that requires each component of the sensor map vector to fall within an arbitrarily small distance from its corresponding reference map value. This is essentially the MAD metric. When the signal-to-noise ratio is very small, the likelihood ratio approaches the unnormalized Product metric. These conclusions, although based on an analysis of a two-picture-element map, are readily extended to maps of arbitrary size.

III. A COMPARISON OF THE LIKELIHOOD RATIO FOR
GAUSSIAN DISTRIBUTIONS WITH OTHER METRICS

It is interesting to compare the metric defined by the likelihood ratio for the Gaussian case discussed in Section II with some of the metrics mentioned earlier in that section.

THE NORMALIZED INNER PRODUCT (NProd) METRIC

In two dimensions, the normalized inner Product metric has the form

$$F(x_1, x_2) = \frac{x_1 s_1 + x_2 s_2}{(x_1^2 + x_2^2)^{1/2} (s_1^2 + s_2^2)^{1/2}}, \quad (17)$$

where the given vector, of course, is the reference map vector (s_1, s_2) . The values of the metric that exceed some prescribed threshold F_0 are those within the angular sector $\pm \cos^{-1} F_0$ centered about the line through the origin and the point (s_1, s_2) , as shown in Fig. 3. The normalized inner product assumes its largest values when the *ratio* of the sensor map values x_2/x_1 is equal to the ratio of the reference map values s_2/s_1 . Thus, this metric is appropriate when no a priori knowledge of absolute map values is assumed. The only assumption is that the two data sets are in a linear relationship.

If the statistics of a sensor map are assumed to be stationary and ergodic, the average norm for a sensor map containing N elements approaches the ensemble average. Thus,

$$\lim_{N \rightarrow \infty} \|x\|/N^{1/2} = \sigma_B, \quad (18)$$

where σ_B is the root-mean-square variation of the sensor map values averaged over an infinite ensemble. If one assumes that the sensor

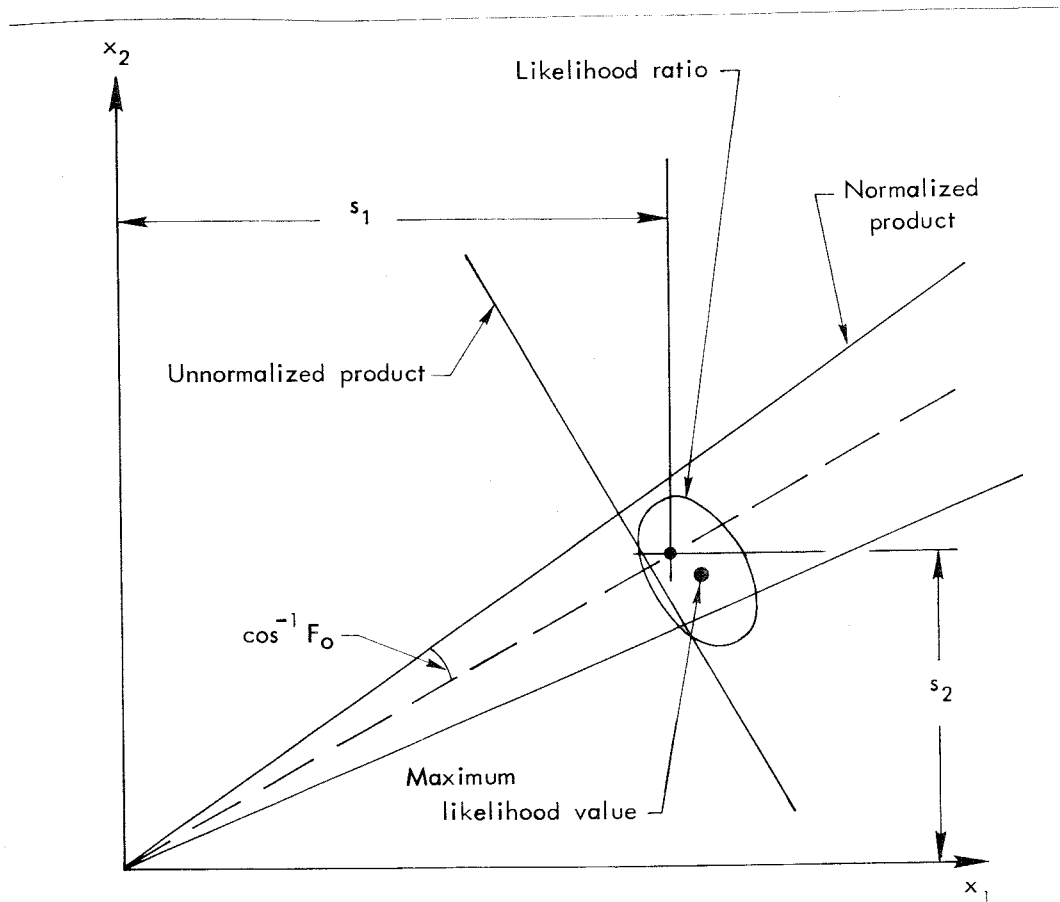


Fig. 3 — Comparison between metrics

map contains so many elements that Eq. (18) can be considered an equality for a finite value of N , then the average norm associated with both the sensor and reference maps can be considered constant for all maps. Except for a constant factor, Eq. (17) reduces to the classical correlator or the unnormalized Product (Prod) metric defined by

$$F(x) = \frac{1}{N} (x \cdot y) , \quad (19)$$

where y is the reference map vector. For this case, computations using either metric will give approximately equal performance.

The distinction between the use of the Product and normalized Product metrics is well illustrated in the calculations of Sections II, III, and IV of Part I^{*} of this report. In Sections II and III, the probability of acquisition is computed over an infinite ensemble so that the unnormalized Product metric can be used. In Section IV, calculations are made for a finite sample of maps. Hence, the different metrics produce different results.

THE DIFFERENCE SQUARED METRIC

The locus of the difference squared metric defined by Eq. (7) is a circle centered at (s_1, s_2) , as shown in Fig. 4. Thus, the acceptance region is similar to that defined by the likelihood ratio for the special case of uncorrelated map values, but the location of the center of the circle is not at the optimum point. The displacement between centers is minor, however, when the signal-to-noise ratio is large. The metric takes no account of the possible correlation between map values as does the likelihood ratio.

THE MEAN ABSOLUTE DIFFERENCE (MAD) METRIC

The locus of the mean absolute difference (MAD) is a family of squares rotated 45 deg with respect to the x_1, x_2 axes, also shown in Fig. 4. The properties of this metric are thus similar to the difference squared metric. The acceptance region is a small symmetrical area surrounding the reference map values. This metric does not make use of the correlation between sensor map values.

PROBABILITY OF CORRECT ACQUISITION

When the desired map is assumed to be present in the ensemble of maps, and when the extremal value of the ensemble is assumed to define the matched condition, the probability of correct acquisition $P_c(s)$

* H. H. Bailey, F. W. Blackwell, C. L. Lowery, and J. A. Ratkovic, *Image Correlation, Part I: Simulation and Analysis*, The Rand Corporation, R-2057/1-PR, November 1976.




Fig. 4 — Further comparison between metrics

for a maximizing metric can be defined as

$$P_c(s) = \int_{-\infty}^{\infty} p(F|S) \left[\int_{-\infty}^F p(F'|B) dF' \right]^Q dF, \quad (20)$$

where $p(F|S)$ is the probability density of the metric given that a match is present, $p(F|B)$ is the probability density of the metric given that a match is not present, and Q is the total number of out-of-register comparisons. Equation (20) is equivalent to Eq. (3) of Part I. A simple change in the limits of integration defines the probability of correct acquisition for a minimizing metric. The calculation of this probability requires a knowledge of the probability density functions of

the metric, which, in turn, depend on the multidimensional joint probability density function of the sensor map values. This joint probability density function is generally unknown, so that some sort of an assumption regarding the form of the distribution of the metric's values must be made in order to use Eq. (20) to estimate the probability of correct acquisition. The calculations described in Sections II and III of Part I of this report are based on the assumption that these distributions are Gaussian.

IV. SUMMARY AND CONCLUSIONS

The theoretical solution to the problem of map matching is to compute the likelihood ratio. Unfortunately, the probability density functions needed to form the likelihood ratio are generally unknown, and, in a practical sense, essentially unmeasurable. Hence, some sort of theoretical model of the statistics of maps must be developed or simply assumed if an approach based on the likelihood ratio is to be pursued.

The common approach to map matching based on the extremals of a metric has little theoretical foundation beyond the simple notion of continuity. The extremal property of a metric generally represents a universal statement about the class of all functions. Moreover, the extremal value is associated with a particular function or a narrowly defined class of functions whose values are precisely specified at every point. This corresponds to the ideal error-free (noiseless) sensor map, which, of course, occurs with probability zero. By continuity, however, it is reasonable to assume that when the sensor map is close to the ideal, the value of the metric is close to the extremal.

When the various map errors can be treated as additive noise, a heuristic interpretation of the likelihood ratio suggests the properties of a good metric. When the noise is small compared with the map variance, and when the absolute amplitude of the sensor map is calibrated to that of the reference map, a metric that requires an approximate equality between corresponding sensor map and reference map values should be used. The MAD algorithm is such a metric. But when the noise is large compared with the map variance, and when no assumption concerning absolute signal levels is made, a metric requiring an approximate proportionality between corresponding sensor map and reference map values should be used. The normalized Product algorithm is such a metric.

The above conclusions, although based on an analysis of a simple two-picture-element map, are readily extended to maps of arbitrary size.

The sometimes encountered statement that the classical correlator, i.e., the Product algorithm, is a theoretically optimum approach to map matching has limited validity. The statement often derives from a confusion between the problem of map matching and the problem of signal detection.

Appendix

LIKELIHOOD RATIO FOR CORRELATED BIVARIATE GAUSSIAN DISTRIBUTIONS

Suppose that the joint probability densities for the matched and unmatched conditions are denoted by $p(x,y|SN)$ and $p(x,y|BN)$, respectively, where

$$p(x,y|SN) = \frac{\exp \left[- \frac{(x - s_1)^2 - 2r_{SN}(x - s_1)(y - s_2) + (y - s_2)^2}{2\sigma_{SN}^2(1 - r_{SN}^2)} \right]}{2\pi\sigma_{SN}^2(1 - r_{SN}^2)^{\frac{1}{2}}} \quad (A.1)$$

and

$$p(x,y|BN) = \frac{\exp \left[- \frac{x^2 - 2r_{BN}xy + y^2}{2\sigma_{BN}^2(1 - r_{BN}^2)} \right]}{2\pi\sigma_{BN}^2(1 - r_{BN}^2)^{\frac{1}{2}}}, \quad (A.2)$$

where, in general, $\sigma_{SN}^2 \neq \sigma_{BN}^2$, and $r_{SN} \neq r_{BN}$.

The likelihood ratio $L(x,y)$ is defined by

$$L(x,y) = \frac{p(x,y|SN)}{p(x,y|BN)}. \quad (A.3)$$

Its logarithm, in terms of the two probability densities defined by Eqs. (A.1) and (A.2), is

$$\begin{aligned} \ln L = & \frac{\sigma_{BN}^2(1 - r_{BN}^2)^{\frac{1}{2}}}{\sigma_{SN}^2(1 - r_{SN}^2)^{\frac{1}{2}}} + \frac{x^2 - 2r_{BN}xy + y^2}{2\sigma_{BN}^2(1 - r_{BN}^2)} \\ & - \frac{(x - s_1)^2 - 2r_{SN}(x - s_1)(y - s_2) + (y - s_2)^2}{2\sigma_{SN}^2(1 - r_{SN}^2)}. \end{aligned} \quad (A.4)$$

If Eq. (A.4) is rewritten as the quadratic form

$$Ax^2 + Bxy + Cy^2 + Dx + Ey + F = 0 , \quad (\text{A.5})$$

the coefficients of the quadratic terms are easily found to be

$$A = \frac{1}{\sigma_{BN}^2 (1 - r_{BN}^2)} - \frac{1}{\sigma_{SN}^2 (1 - r_{SN}^2)} = C , \quad (\text{A.6})$$

$$B = \frac{-r_B}{\sigma_{BN}^2 (1 - r_{BN}^2)} + \frac{r_S}{\sigma_{SN}^2 (1 - r_{SN}^2)} . \quad (\text{A.7})$$

From elementary analytical geometry, the acute positive rotation angle, θ , that transforms the conic to its principal axes is given by

$$\tan 2\theta = \frac{B}{A - C} . \quad (\text{A.8})$$

If B as defined by Eq. (A.7) is assumed to be different from zero, it follows immediately that the principal axes of the conic are rotated 45 deg with respect to the (x,y) coordinate axes.

If the conic is assumed to be nondegenerate, i.e., if

$$\Delta = \begin{vmatrix} 2A & B & D \\ B & 2C & E \\ E & D & 2F \end{vmatrix} \neq 0 , \quad (\text{A.9})$$

then, from elementary analytical geometry, the type of conic is determined by the following rules:

1. If $B^2 - 4AC = 0$, the conic is a parabola.
2. If $B^2 - 4AC < 0$, and Δ and $A + C$ disagree in sign, the conic is an ellipse.

3. If $B^2 - 4AC < 0$, and Δ and $A + C$ agree in sign, the conic is an imaginary ellipse, i.e., no real locus.
4. If $B^2 - 4AC > 0$, the conic is a hyperbola.

From physical considerations we know that a real locus always exists, so that the condition $B^2 - 4AC < 0$ alone is sufficient to conclude that the conic is an ellipse. If W is defined as

$$W = B^2 - 4AC, \quad (\text{A.10})$$

then, since $A = C$ from Eq. (A.6), we have

$$W = B^2 - 4AC = (B - 2A)(B + 2A). \quad (\text{A.11})$$

By combining Eqs. (A.6) and (A.7), it is easy to show that

$$B - 2A = - \left[\frac{1}{\sigma_{BN}^2(1 - r_B)} - \frac{1}{\sigma_{SN}^2(1 - r_S)} \right] \quad (\text{A.12})$$

and

$$B + 2A = \frac{1}{\sigma_{BN}^2(1 + r_B)} - \frac{1}{\sigma_{SN}^2(1 + r_{SN})}. \quad (\text{A.13})$$

From elementary probability theory a quantity of the form $\sigma^2(1 + r)$ is recognized as one-half the variance of the sum of two correlated random variables. Similarly a quantity of the form $\sigma^2(1 - r)$ is recognized as one-half the variance of the difference of two correlated random variables. Using these facts, the relationship between the type of conic and the properties of the two probability densities is immediately evident from Eqs. (A.10), (A.11), (A.12), and (A.13). When the variances of the sum and difference of the pixel values for the matched condition are both either less than, or greater than, the corresponding variances for the unmatched condition, the conic is an ellipse. When the respective variances for the matched and unmatched conditions are in an

opposite algebraic relationship--i.e., in the matched condition one variance exceeds while the other is less than its corresponding variance for the unmatched condition--the conic is a hyperbola.

Of the various possibilities, the case in which the variances in the matched condition are each less than their corresponding variances in the unmatched condition is the principal case of practical interest. Although the variability of the sensor map for the matched condition depends on the statistics of the map in the vicinity of the match point, its variation on the average should not exceed that associated with the unmatched condition. Hence, the acceptance region defined by the likelihood ratio will be elliptical in most cases of practical interest.

The center of the conic can be found either by completing the squares in Eq. (A.4) or by finding the coordinates of the extremum of the likelihood ratio by using elementary differential calculus. If the conic is transformed to its principal axes by the transformation

$$u = (x - y)/2^{\frac{1}{2}} \tag{A.14}$$

and

$$v = (x + y)2^{\frac{1}{2}}, \tag{A.15}$$

it can be shown that the coordinates of the center (u_o, v_o) are given by

$$u_o = \frac{-(s_1 + s_2)}{2^{\frac{1}{2}} Q \sigma_{SN}^2 (1 + r_{SN})} \tag{A.16}$$

and

$$v_o = \frac{-(s_2 - s_1)}{2^{\frac{1}{2}} P \sigma_{SN}^2 (1 - r_{SN})}, \tag{A.17}$$

where

$$P = B - 2A \tag{A.18}$$

and

$$Q = B + 2A . \tag{A.19}$$

Thus, the conics for different values of the likelihood ratio are all concentric.

Finally, it can be shown by completing the squares in Eq. (A.4) that the major and minor axes of the conic are always in a constant ratio, namely

$$u : v = |P|^{\frac{1}{2}} : |Q|^{\frac{1}{2}} . \tag{A.20}$$