

Title: **The Ineffectiveness of the Correlation Coefficient for Image Comparisons**

Author(s): Eugene K. Yen
Roger G. Johnston

Submitted to:

<http://lib-www.lanl.gov/la-pubs/00418797.pdf>



Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; therefore, the Laboratory as an institution does not endorse the viewpoint of a publication or guarantee its technical correctness.

LAUR #96-2474

The Ineffectiveness of the Correlation Coefficient for Image Comparisons

Eugene K. Yen and Roger G. Johnston, Ph.D., CPP

Vulnerability Assessment Team
Los Alamos National Laboratory
MS J565, Los Alamos, New Mexico 87545
phone: 505-667-7414 fax: 505-665-4631
email: roger_johnston@lanl.gov

The Ineffectiveness of the Correlation Coefficient for Image Comparisons

ABSTRACT: Pearson's r linear correlation coefficient is widely used for comparing images. Image processing experts appear to be cognizant of its serious limitations. This information, however, has not been communicated well to non-experts (including those working on security applications), who sometimes use the correlation coefficient without being aware of its problems. We discuss the disadvantages of the correlation coefficient and show specific examples of strikingly poor performance in the context of security.

KEYWORDS: correlation coefficient, tamper detection, tags, seals, image comparison

Introduction

Pearson's correlation coefficient, r , is widely used in statistical analysis, pattern recognition, and image processing [1-8]. Applications for the latter include comparing two images for the purposes of image registration, object recognition, and disparity measurement. For monochrome digital images, the Pearson correlation coefficient is defined as [1,4,7]:

$$r = \frac{\sum_i (x_i - x_m)(y_i - y_m)}{\sqrt{\sum_i (x_i - x_m)^2} \sqrt{\sum_i (y_i - y_m)^2}}$$

where x_i is the intensity of the i th pixel in image 1, y_i is the intensity of the i th pixel in image 2, x_m is the mean intensity of image 1, and y_m is the mean intensity of image 2.

The correlation coefficient has the value $r=1$ if the two images are absolutely identical, $r=0$ if they are completely uncorrelated, and $r=-1$ if they are completely anti-correlated, for example, if one image is the negative of the other.

The correlation coefficient is used for security applications such as surveillance, treaty verification, tamper detection using security seals, and tagging [9-14]. (Seals leave unerasable evidence of entry or tampering; tags are devices or procedures that uniquely fingerprint an object.) Typically, the correlation coefficient is used to compare two images of the same object (or scene), taken at different times [11-14]. The r value indicates whether the object has been altered or moved.

In theory, we would obtain an r value of 1 if the object is intact, and a value of less than 1 if alteration or movement has occurred. In practice, distortions in the imaging system, pixel noise, slight variations in the object's position relative to the camera, and other factors produce an r value less than 1, even if the object has not been moved or physically altered in any manner [13]. In our experience with security applications, typical r values for two digital images of the same scene, one recorded immediately after the other using the same imaging system and illumination, range from 0.95 to 0.98. This is the same range as that reported by Palm and DeVolpi [13] for scanning electron microscope images.

Usually there needs to be an empirical definition of a threshold r value that indicates a breach of security. In other words, it is necessary to determine the minimum r value needed to conclude with confidence that the image is unchanged. For several security applications for which we are familiar, the chosen threshold values for r range from 0.30 to 0.85, depending on the application [9-14].

The purpose of this paper is to offer simple examples of the poor performance of the correlation coefficient for image comparison, particularly with security applications in mind. We know of no similar presentation. The problems and limitations of the correlation coefficient have been discussed previously [2, 3, 13-17], but briefly and abstractly, without specific image examples. Furthermore, the knowledge that the correlation coefficient often performs poorly does not seem to have been communicated well to non-experts. We conducted a survey of 50 image processing texts, introductory through advanced, that discussed the correlation coefficient. Only 3 of the 50 texts discussed potential problems and limitations of the correlation coefficient, and then mostly in terms of its computational intensiveness, rather than performance problems.

Strengths and Weaknesses of the Correlation Coefficient

One of the obvious advantages of Pearson's correlation coefficient is that it condenses the comparison of 2 (often large) two-dimensional images down to a single scalar, r . Additionally, the correlation coefficient is completely invariant to linear transformations of x or y [1, 4, 13]. As a result, r is insensitive (within limits) to uniform variations in brightness or contrast across an image. Such spatially uniform variations can be caused, for example, by differences in the brightness of the light source over time, by changing background light levels, or by variations in the gain of the imaging system.

Despite its advantages, the correlation coefficient has many problems and limitations. The most widely recognized disadvantage is that it is computationally intensive. This often limits its usefulness for image registration (that is, orienting and positioning two images so they overlap). The correlation coefficient is also extremely sensitive to the image skewing, pincushioning, and vignetting that inevitably occur in imaging systems. Such distortions are particularly prevalent in scanning electron microscope pictures [13] because of the non-linearities and complexities of electron optics. (Image skewing occurs when an image slants in one direction. Pincushioning means that the edges of the image are concave. Vignetting is a reduction in the image intensity near the edges due to optical light collecting considerations.)

Another problem often overlooked in practical applications is that r is undefined--due to division by zero--if one of the test images has constant, uniform intensity. Users, especially for automated security applications, must be careful that the computer or micro-processor doing the r calculations does not permit an undefined value to default to $r=0$ or $r=1$.

Other problems with the correlation coefficient include possible bias [16], complexities of interpretation [1], over-sensitivity to pixel noise and gain variations [2, 5, 15], difficulties in dealing with perspective or with moving illumination sources [2,15, 17], undesirable behavior for images containing too much fine structure or too little [2, 15], and trouble in dealing with images having strong spatial disparity gradients [2].

There is an additional problem with the correlation coefficient that does not appear to have been discussed or demonstrated in detail elsewhere. An image can be greatly modified, without this being detected, as long as the local mean and/or histogram of pixel intensities are relatively unchanged. This is demonstrated below.

Examples of Poor Performance

Figure 1(a) shows an 8-bit (monochrome) 512 X 512 scanning electron microscope image of a metal surface. The complexity of the surface can be used as a unique identifier, or tag [13]. For figure 1(b), we have used a computer to artificially modify the image in figure 1(a) such that the letters "LANL" replace a portion of the original image. The intensity for the letters was chosen such that the mean intensity of the region being overwritten did not change. Even though the human eye can immediately judge that figure 1(a) and 1(b) are dramatically different, the correlation coefficient reports that these two images are essentially identical: $r=0.94$. This is about the same correlation coefficient one gets by re-recording a second image of the surface in figure 1(a).

Figures 2(a) and 2(b) show an even more disturbing example of the failure of the correlation coefficient to detect changes in an image. Recognizing that individual pieces are missing is critical for applications such as reflective particle tags (RPTs). These consist of small, highly reflective particles attached to an object. RPTs uniquely fingerprint the object based on their complex spatial distribution [12].

Figure 2(a) is a monochrome 8-bit video image with 512 X 512 pixels. The image shows a collection of plastic-coated paper clips spread out randomly on a surface. Figure 2(b) shows a second image of the scene except that one of the paper clips has been physically removed prior to recording this second image. The dark paper clip missing from figure 2(b) can be seen slightly up and to the right of center in figure 2(a). The correlation coefficient for figure 2(a) vs. figure 2(b) is $r=0.98$.

Now one way to improve the performance of the correlation coefficient--and make it less sensitive to image skewing, pincushioning, vignetting, or imperfect registration--is to compute r over a subset region of the complete image [13]. Even this approach, however, can still produce poor performance. If only the region inside the rectangle in figure 3 is used to compute r , the value is still quite large: $r=0.86$. The correlation coefficient is thus still reporting an almost unchanged image, even though the missing paper clip is a major feature inside the (200 X 200 pixel) rectangle!

Finally, in figure 4(a) and 4(b), the correlation coefficient fails to be useful for Elvis sightings. The two images are clearly different, yet $r = 0.94$. Image 4(b) was created by a superposition of figure 4(a) and an image of Elvis, such that the resulting local mean and histogram of pixel intensities were largely unchanged.

Concluding Remarks

In summary, the correlation coefficient often fails to find differences in images that are widely disparate. In the case of a security system utilizing the correlation coefficient, an adversary can modify an object or scene quite dramatically and yet still go undetected, especially if he approximately preserves the local intensity mean and/or intensity histogram. Of particular significance for security applications, the correlation coefficient often fails to detect missing objects within an image. Performance often improves only modestly if the correlation coefficient is computed for subset windows of the entire image.

Even when the correlation coefficient does perform acceptably, there are usually better algorithms for image comparison. Typically, the optimum choice of algorithms depends critically on general characteristics of the relevant images, and details of the application. One fact often overlooked is that the use of human vision with a blink comparator [10, 18] can often dramatically outperform even very sophisticated computer algorithms.

Acknowledgments

This work was performed under the auspices of the United States Department of Energy. Anthony Garcia, Phil Jacobson, and Chuck Mansfield provided valuable assistance and input.

References

1. J.L. Rodgers, J. L. and W.A. Nicewander, "Thirteen Ways to Look at the Correlation Coefficient", *American Statistician* 42, 59-66 (1995).
2. M. Jenkin, A.D. Jepson, and J.L. Tsotsos, "Techniques for Disparity Measurement", *CVGIP: Image Understanding* 53, 14-30 (1991).
3. R.Y. Wong, E.L. Hall, and J. Rouge, "Hierarchical Search for Image Matching", *Proceedings of the IEEE Conference on Decision Control*, pp. 405-408 (December, 1976).
4. M. James, Pattern Recognition, John Wiley and Sons, New York (1988), pp. 36-40.
5. D.I. Barnea and H.F. Silverman, "A class of Algorithms for Fast Digital Image Registration", *IEEE Transactions on Computers* C-21, 179-186 (1972).
6. E.H. Hall, Computer Image Processing and Recognition, Academic, New York (1979), pp. 480-485.
7. W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling, Numerical Recipes in Pascal, Cambridge University Press, New York (1989), pp. 532-534.
8. J. Lee, "A Cautionary Note on the Use of the Correlation-Coefficient", *British Journal of Industrial Medicine* 49, 526-527 (1992).
9. R.G. Johnston and A.R.E. Garcia, "Vulnerability Assessment of Security Seals", *Journal of Security Administration* 224, 24-29 (1995).
10. R.G. Johnston, "The Real Deal on Seals", *Security Management* 41, 93-100 (1997).
11. D.S. Kupperman, A.C. Raptis, J.T. Dusek, R.G. Palm, S.H. Sheen, and S.E. Dorris, "Assessment of Tamper-Revealing Ceramic Seals", *Argonne National Laboratory Report ANL/ACTV-90/5, DE91007793*, (September, 1990).
12. G. Staehle (Editor), "DOE's Tags and Seals Program", *Verification Technologies*, U.S. Department of Energy Report DOE/DP/OAC/VT-92B, (October, 1992).
13. R.G. Palm and A. DeVolpi, "Plastic-Casting Intrinsic-Surface Unique Identifier (Tag)", *Argonne National Laboratory Report ANL/ACTV-94/1* (April 1995).
14. A. DeVolpi, "Understanding Correlation Coefficients in Treaty Verification", *Argonne National Laboratory Report ANL/ACTV-91/4 REVISED* (February 1993).
15. M.A. Crombie, "Coordination of Stereo Image Registration and Pixel Classification", *Photogrammetric Engineering and Remote Sensing* 49, 529-532 (1983).
16. I.J. Good and E.P. Smith, "The Possible Bias of the Pearson Chi-Squared Test in Non-Equiprobable Cases", *J. Statist. Comput. Simul.* 19, 79-95 (1984).
17. R.A. Schowengerdt, Techniques for Image Processing and Classification in Remote Sensing, Academic Press, New York (1983), pp. 27-35, 47.

18. R.G. Johnston and A.R.E. Garcia, "Simple, Low-Cost Ways to Dramatically Improve the Security of Tags and Seals". CD-ROM Proceedings of the IAEA Symposium on International Safeguards, Vienna Austria, October 13-17, 1997.

Figure Captions

Figure 1. Comparing two images using the correlation coefficient. Both images are 8-bit, 512 X 512 scanning electron microscope images of a metal surface. The image (b) is the same as (a) except that the letters "LANL" have been overwritten. The pixel intensity of the letters is approximately equal to the local mean intensity of the original image (a). The correlation coefficient for image (a) vs. (b) is $r=0.94$. The correlation coefficient reports little difference.

Figure 2. The correlation coefficient can be relatively poor at detecting missing objects in an image. Both (a) and (b) are 512 X 512 8-bit video images of plastic-coated paper clips. For (b), one of the paper clips has been picked up and the image re-acquired. The missing paper clip is not detected: the correlation coefficient for (a) vs. (b) is $r=0.98$.

Figure 3. Reducing the window size only modestly helps the correlation coefficient detect the missing paper clip. If the correlation coefficient is recomputed for image (a) vs. (b) in figure 2, limited to the subset (200 X 200) rectangle shown in this figure, $r=0.86$.

Figure 4. Scanning electron microscope images. The correlation coefficient fails to definitively detect the ghostly appearance of Elvis in (b). For (a) vs. (b), $r=0.94$.

Figure 1a

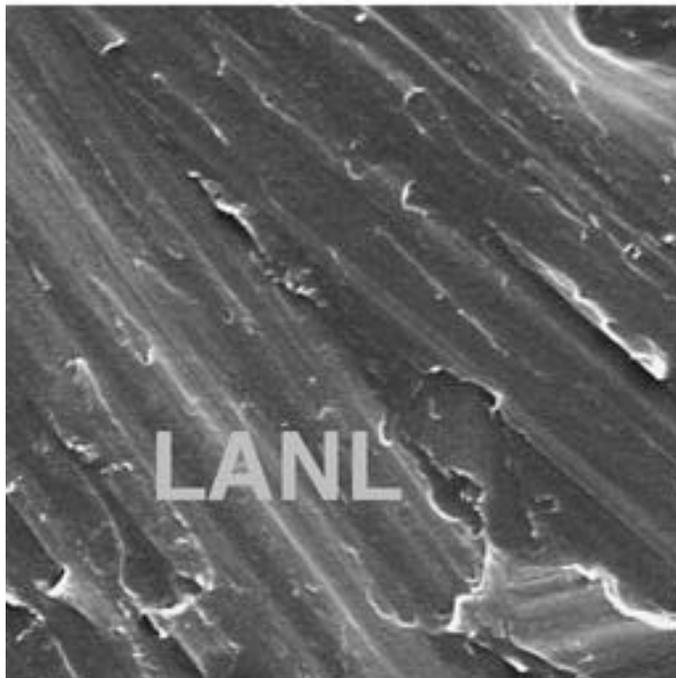
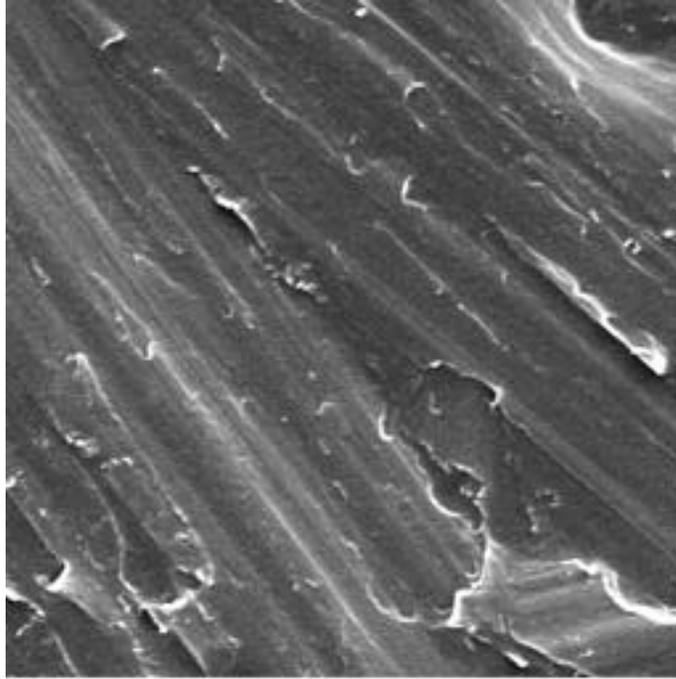


Figure 1b

Figure 2a

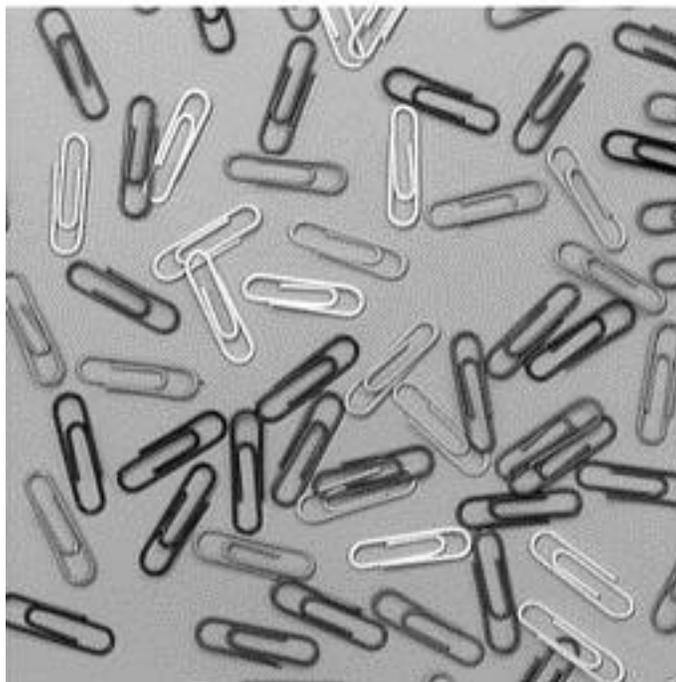
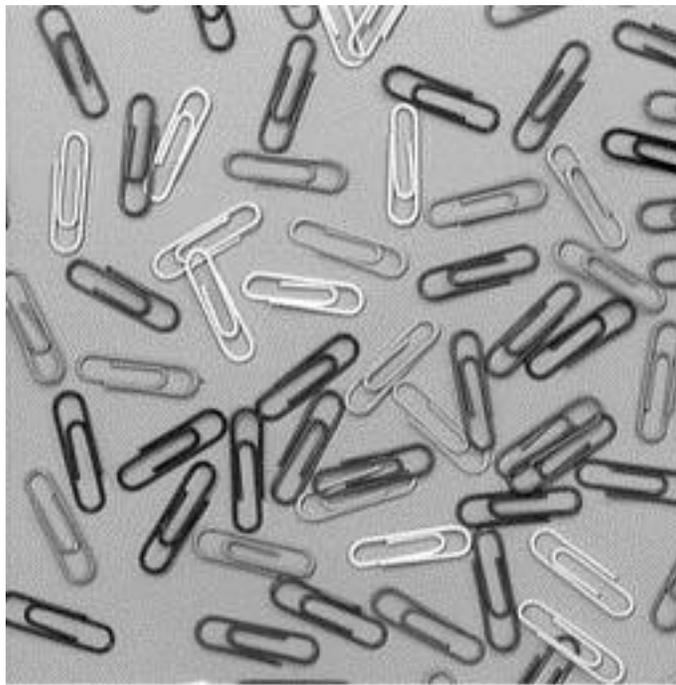


Figure 2b

Figure 3

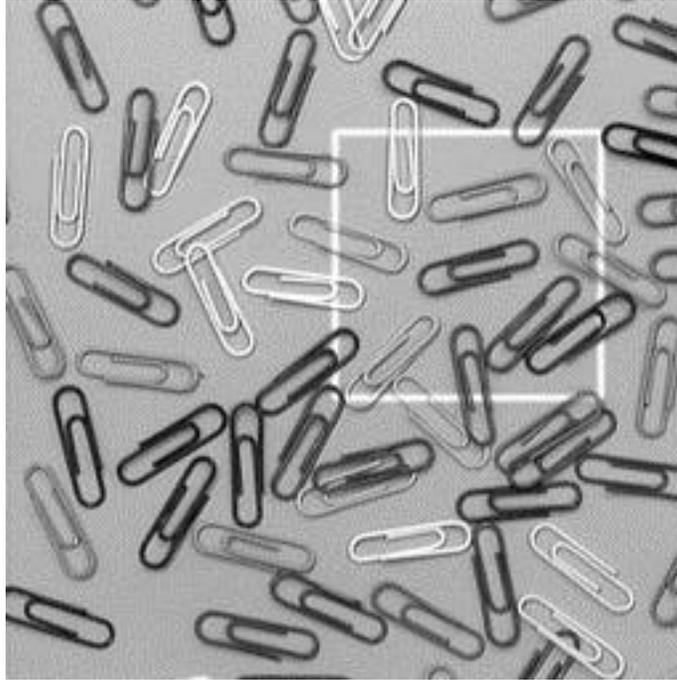


Figure 4a

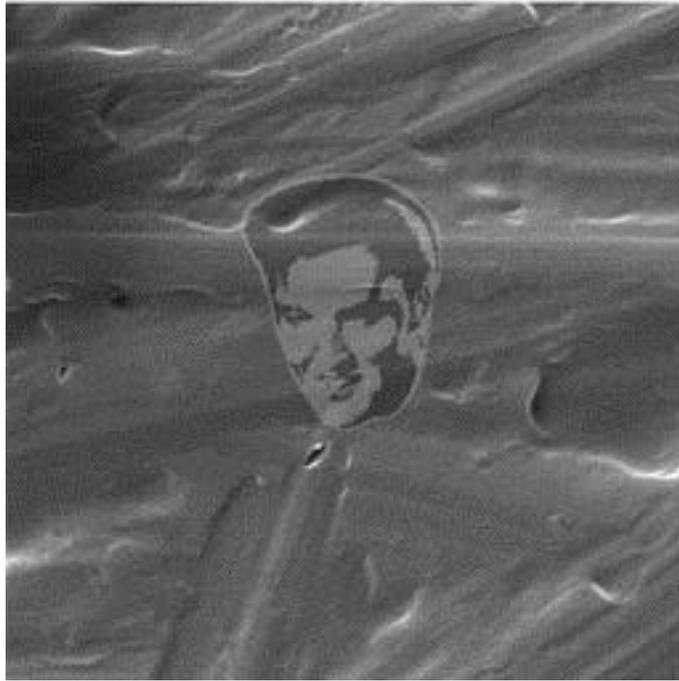
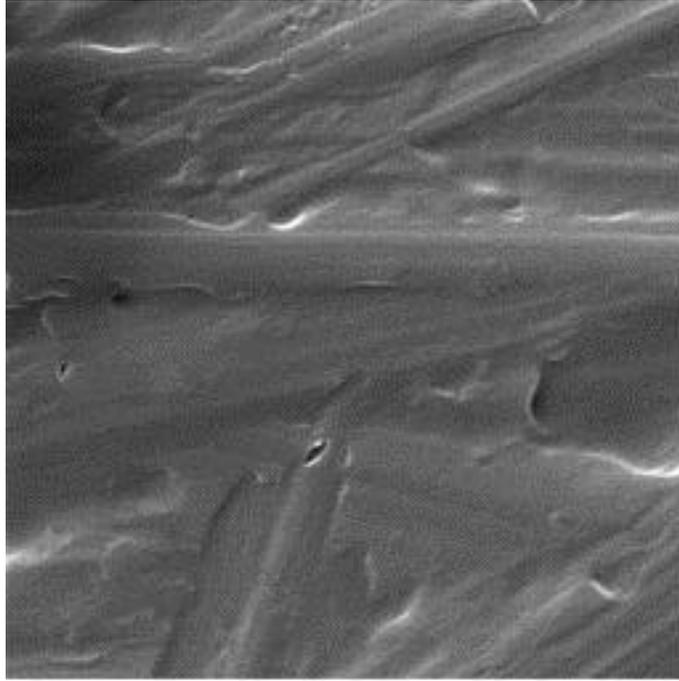


Figure 4b