Journal of Imaging Science and Technology® 57(4): 040502-1–040502-7, 2013.

Noise and Background Resistance of Correlation Image Alignment for Military Surveillance

Oscar J.G. Somsen and Fok Bolderheij
Netherlands Defense Academy, Den Helder, Netherlands
E-mail: ojg.somsen@nlda.nl

Abstract. *Electro-optical devices are used for military applications to detect, identify and track targets. Typically, video information is presented to an operator. With an increasing availability of such devices data volumes are becoming large, and the need for automated analysis is becoming more urgent. In a military setting, this typically involves detecting and identifying targets as early as possible, i.e., when little visual information is available. The identification can be facilitated by combining the video stream into single enhanced images that provide more information for the operator. Using simulated and basic experimental images the authors study alignment in the aforementioned context and find that basic correlation is a potentially useful technique. Problems with background variation can be overcome and good alignment can be obtained even with severe noise. The authors illustrate how alignment quality responds to various parameters, which will help in the development of practical applications. © 2013 Society for Imaging Science and Technology.*
[DOI: 10.2352/J.ImagingSci.Technol.2013.57.4.040502]

INTRODUCTION

The application of electro-optical imaging devices such as video is a developing field. In particular, military applications have long been dominated by radar and sonar although night vision and infrared play an important role. At present, however, electro-optical sensors are becoming more significant, also in the visible region, for example to identify objects that have been detected by other means or to detect objects, such as sea-skimming missiles, that can escape other sensors. Optical devices are essentially passive and do not betray one's presence to other parties. With the increasing availability of high volumes of optical image data there is an increasing opportunity as well as an increasing need for automated detection and recognition of objects.[4,6]

Early warning, which is often a priority in military surveillance, leads to special considerations for image acquisition and analysis techniques. There is a need to detect potential targets as early as possible, before they may become a threat. Also, it is important to rapidly identify and characterize a detected target. The sooner information is available, the better an appropriate response can be prepared. In such cases very little optical information is typically available either because the object is at a large distance (detection or identification may be attempted when the object spans only a few pixels) or because detection is hindered, e.g., by atmospheric conditions or the background.

To facilitate early identification of an object one may combine multiple images into a single enhanced image. Even when a rapid identification is required, at least a few seconds of video stream will be available. Combining multiple images can help to distinguish an object from the background, remove coverage by foreground features and reduce noise. It may even be used to obtain an image of an object moving behind an extensive foreground such as foliage. When both the object and the conditions are very steady, the image quality can be enhanced simply by a longer exposure. Otherwise it is essential to align the object within the set of images before combining them into a single enhanced image. Image registration and restoration may even be considered a joint process.[14,17]

Combination of multiple images is not by itself a new technique. Perhaps best known is the combination of images into a mosaic panorama[2] or the application to motion analysis.[13] Multiple images are also used to obtain super-resolution,[14,19] e.g., for satellite images or even conversion of TV signals to HDTV. The alignment or registration of these images is often carried out with the help of (multi-scale) features.[3,11] Typically these applications involve registration and stitching of large sections or entire images. Correlation is also used especially to correct for small alignment errors or when feature extraction is not possible.[7,10,23,24] Refinements have been developed to prevent erroneous alignment of background or illumination features.[1,18] Optical flow detection uses gradients to detect displacements down to the sub-pixel level.[7,20]

As described above, early identification implies that the object's image is relatively featureless. Also, the target will typically be meters to tens of meters in size and will be located at a distance of several kilometers or more. Thus, it will span a relatively small angle of the order of 1 to at most 10 mrad. Enlarged images may be obtained with a telephoto lens or a detection algorithm may extract sub-images that contain the object for further analysis. In many applications, there will be no other objects within that range that may confuse alignment. Also, many backgrounds such as sea or sky can be relatively constant on this scale. Due to the above, it may be useful to consider the basic correlation algorithm to align these images and present an enhanced image to the human observer.

Received Feb. 24, 2013; accepted for publication Aug. 15, 2013; published online Jul. 1, 2013. Associate Editor: Yeong Ho Ha.
1062-3701/2013/57(4)/040502/7/$20.00

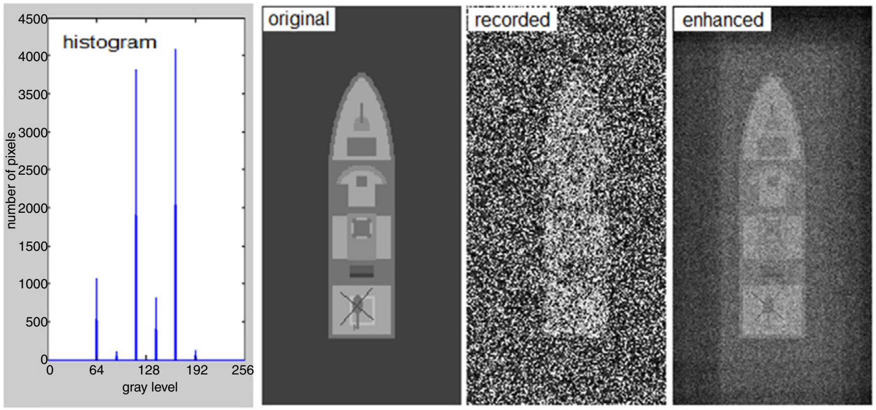

Figure 1. Left: gray-scale histogram of the test image (without margin). Right: illustration of a typical enhancement procedure (see text) with (left to right) the original image (50 × 200 pixel object surrounded by a 50 pixel margin), one of the test images (object displaced and noise added) and the final enhanced result (the dark areas are an artifact of the procedure caused by the fact that the blank areas in the aligned images were filled with black).

In this article we mainly investigate correlation alignment and how it is affected by various parameters. Particular in our approach is the type of long range early identification image for which correlation alignment could work well. Using simulated images we investigate how the quality of the alignment is affected by parameters for noise, resolution, background, etc. We study the underlying mechanisms which will help in the acquisition and analysis of experimental images. After describing the method we will first study a basic artificial image and white noise, then consider variability of the background and finally analyze simple experimental images.

METHODS

In this article we consider a basic correlation alignment. Briefly, the correlation of two gray-scale images is obtained by displacing one image with respect to the other, multiplying the values of corresponding pixels and summing the result over all pixels. If the two images are identical, it can be shown that the correlation is maximal when they are aligned. Therefore, we shall define the optimal alignment as the displacement for which the correlation is maximal. Note that this is not always correct. Images may be constructed that are obviously misaligned to the human eye while still satisfying the above criterion. By determining when and how these cases occur it may be possible to prevent them.

Numerically, the correlation is defined as follows. Each image is represented by an $N \times M$ array $A(x, y)$ of gray levels, typically between 0 (black) and 255 (white). With two images A and B of the same size, their correlation is $C(\mathrm{d}x, \mathrm{d}y) = \sum_x \sum_y A(x, y)B(x+\mathrm{d}x, y+\mathrm{d}y)$. The summation is carried out over all pixels, while dx and dy define the displacement between the images. We assume periodicity ($B(x \pm N, y) = B(x, y \pm M) = B(x, y)$) so that the correlation is efficiently calculated with fast Fourier transforms. With a sufficiently constant background the assumed periodicity will not affect results. The correlation may be refined by subtracting the average from each image and/or normalizing the images. This reduces the sensitivity for global variation in background or illumination.[7] We do not do this here because subtracting a constant from an image and/or division by a fixed number has no effect on the displacement for which correlation is maximal. We only need to worry about background or illumination variation within each window. This will be considered below.

Note that we consider only translational displacement. Rotational alignment may be carried out with rotated copies of the reference, selecting the one that aligns best. While this requires much computation compared to translation, it is not a practical problem especially with low-resolution images. The reason that we have not done this is that in order to obtain quantitative statistics we had to align many thousands of image pairs. Since the errors for rotational alignment are likely to be qualitatively similar to those for translation we considered only the former. Since we are using small sub-images, more complex alignment that includes resizing or even deformation[7,10] has not been considered here.

The basis of our test image is a gray-scale top view drawing of a navy vessel ($w \times l = 50 \times 200$ pixels), as shown in Figure 1. As indicated by the histogram, a few distinct gray-scale levels were used in this case. The level 64 corresponds to background pixels at the bow of the ship. A margin is added around the original test image, also with the same gray level. While most features in the object are of the order of tens of pixels wide, some narrower features are also included. For instance, the rotor blades of the helicopter are only one pixel wide in this image.

Test images are generated by making a randomly displaced copy of the above and then introducing noise. This may come in the form of, e.g., additive or salt and pepper noise,[3] clutter,[8] turbulence[5,9] or blurring.[4] Preliminary results with additive noise showed that a very large magnitude was often needed to see any effect. Therefore, we switched to salt and pepper noise which is created by resetting a random fraction of the pixels. Unless stated otherwise the value of the reset pixels is randomly set to 0 (black) or 255 (white) with equal probability. This type of noise occurs under very low light conditions. However, it essentially models that the information in a fraction of the pixels is lost. It may also

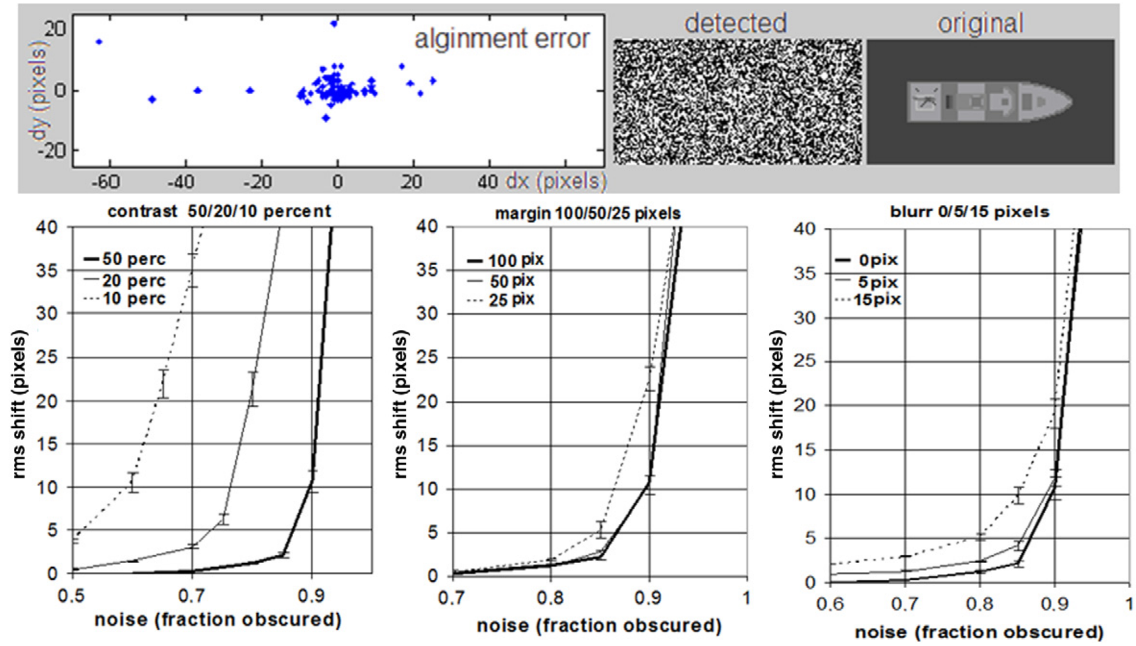

Figure 2. Obscuring by noise (see text). Top (left to right): distribution of noise induced shifts (in pixels), one of the experimental images (90% of the pixels obscured by noise) and the original (reference) image (50 pixel margin). Bottom: images of rms displacement against noise fraction showing the effects of (left to right) contrast levels (at 50% (thick), 20% (thin) or 10% (dotted), margin width (100 (thin), 50 (thick) or 25 (dotted) pixels) and image blurring (using a Gaussian point spread function with rms width of 0 (thick), 3.5 (thin) or 10.6 (dotted) pixels). In all images, the error bars represent 3–10 runs of 100 individually simulated images.

represent clutter, e.g., by precipitation or even foreground. The probability at which each pixel is reset, i.e., the fraction between 0 and 1, defines the amount of noise.

While the main focus of this article is the alignment, Fig. 1 shows one example of the image enhancement that can be obtained with a set of aligned images. The details have been presented elsewhere.[16] In this case twenty-four test images were created by randomly placing the object within the image and creating noise by resetting 90% of the image. As shown, all but the largest features have been obscured by noise. After aligning all test images to a reference, their average image was obtained. The resulting enhanced average image not only shows a finer grained noise with lower amplitude but also many finer details of the original image, including the helicopter rotors. While the circumstances were rather ideal, using identical images, a constant background and random noise, this result shows that considerable enhancement can be obtained by correlation alignment even in the presence of severe noise.

ALIGNMENT WITH NOISY IMAGES

Figure 2 presents the alignment results and its dependence on various parameters for the image and noise. To obtain these results we created a large number of displaced noisy copies of the reference as described above and used correlation to estimate the displacement. The difference between the actual and the estimated displacement is shown. The top frame shows a typical case where the rms error is 1.4 pixels vertically and 7.5 pixels horizontally. A number of outliers are also observed. If these stand out from a set of images it may be possible to remove them and thus improve the alignment quality. The error is relatively small (<4% of the object size) considering the excessive noise. In this particular case 90% of the pixels were obscured. As can be seen, this makes the individual images unintelligible to the human eye. The rms error determines how much detail can be recovered if sufficient images are available.

In the bottom images of Fig. 2, the rms alignment error is plotted against the percentage of pixels obscured by noise. The 90% obscuring in the example may occur only under extreme circumstances. In all cases the alignment shows a switch-like behavior. At low noise levels shifts of a low number of pixels occur. Typically, the decrease is relatively steep. In most cases the error decreases from 10 to below 1 pixel within a three-fold increase of the non-obscured pixel fraction. When the noise increases above some threshold level the alignment deteriorates rapidly and is soon lost completely. Unfortunately, the change is not so sudden that an unambiguous threshold level can be assigned. Below we obtain it somewhat arbitrarily with an error of 10 pixels, approximately 10% of the object's size in the image. The threshold may be taken as a measure of the object's "noise resistance".

To investigate how the alignment is affected by various parameters, we first varied the contrast of the object. As shown in Fig. 1, the object spans a range of $192-64=128$ gray levels. This is obtained by multiplying an original drawing by 0.5 and adding 64, so that the background is gray instead of black. In this series we reduced the multiplication factor to 0.2 and 0.1. The results (bottom left) show that the error increases considerably when the contrast is reduced. The threshold at which the shift reaches 10 pixels decreases

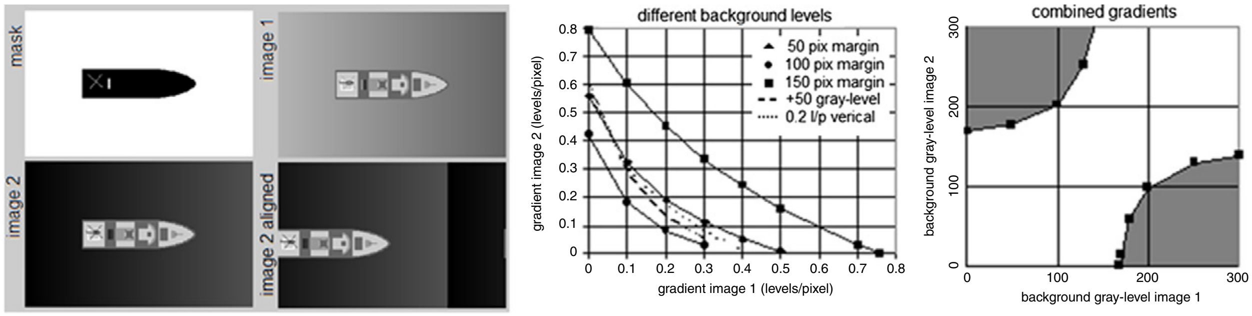

Figure 3. Background effects on image alignment. Left: typical case with mask (top left), image 1 (top right), image 2 (bottom left) and (mis)aligned copy of image 2 (bottom right). Opposing gradients (0.2 levels/pixel) were used and 100 gray levels were added to image 1. Center: threshold levels with gradient background and a margin of 50 (squares), 150 (circles) or 100 pixels (triangles). For the latter case we also show an added constant (50 gray levels, dashed) and vertical gradient (0.2 levels/pixel, dotted). Right: threshold levels with constant background around the object. Alignment is lost in the shaded area.

from 0.9 to 0.6 when the contrast is reduced from 50 to 10%. This result may be compared to the areas within the image that are covered with the object and with noise. The latter increases fourfold (from 10 to 40% of the image) while the former decreases with a factor of 1.5 (from 90 to 60% of the image). The ratio of the two areas increases with a factor of $4 \times 1.5 = 6$, which agrees nicely with the fivefold reduction of contrast.

As the second parameter, we considered the margin around the object. This parameter is important when windows with the detected object are copied from the video frames. When a large margin is used, the risk is lower that the object is (partially) outside the image. On the other hand, a large margin means that more noisy data are added, which may reduce the quality of the alignment. By default we use a margin of 50 pixels. The graph in Fig. 2 (bottom middle) shows what happens when the margin is altered to 25 or 100 pixels. Remarkably, the increase to a 100 pixel margin does not reduce the quality of the alignment. This is indeed remarkable since in that case the image is largely filled with background. The object covers only 10% of the image. This does not seem to affect the alignment at all. In fact, the results show that a decrease of the margin to 25 pixels makes the alignment more susceptible to noise. Apparently a minimal amount of background area is necessary to average out the noise correlation. While an exact number would require more extensive simulation we suggest for now that a margin covering 80–90% of the image does not appear to be a problem for the alignment procedure.

In order to understand the importance of structures within the object for its alignment, we considered various degrees of blurring. For this, the image of the object was correlated with a Gaussian point spread function before the noise was added. This will smooth out structures within the object's image, especially those that are narrower than the point spread function. The results (bottom right) are as one would expect: with increased blurring the alignment is more easily disturbed by the addition of noise. Although an analysis of independent correlation functions is not strictly allowed because the image is not a linear sum of object and noise, it may help to analyze these results. Due to the blurring, the correlation function of the object (with the reference) widens and therefore becomes lower while the correlation function of the noise (also with the reference) is unaffected. When the object's correlation function becomes lower than that of the noise, the alignment is lost.

ALIGNMENT WITH BACKGROUND VARIATION

With correlation alignment it is important to consider the possibility that the backgrounds are not constant and will therefore correlate as well as the objects. It is not possible to distinguish these two types of correlation or even cross-correlation of the background from one image with the object from the other. If the magnitude of the background variation is too large, this will lead to a shift or even loss of the alignment of the objects. Since the background typically covers a larger area than the object, even a relatively small variation may affect the alignment.

In order to study how the alignment is affected by the background, we add a linear function of the pixel coordinates. This is a reasonable assumption since the object spans a small angle, as discussed in the introduction. A typical result is shown in Figure 3 (left). In one image the background becomes darker to the right and in the other to the left. An additional constant background was added to the latter. In the original images, the object was (exactly) in the center, but in the "optimal" alignment, the object is shifted to the edge and somewhat down. This is clearly not correct, but the light background areas do indeed overlap better than in the original images. Apparently, it is the backgrounds that have been aligned in this case rather than the objects. The vertical shift is likely due to cross-correlation.

In all the simulations, we found a strong switch-like effect. If the magnitude of the background signal is small, the objects remain perfectly aligned. When it increases above a threshold level, the alignment is lost completely. Typically, the object shifts to the edge of the image, as in the example, and remains there when the background becomes stronger. In some cases the object is shifted some 10–15 pixels over the edge. The sudden shift indicates that the correlation peaks of both the object and the background are relatively narrow.

To study the effect of background gradient on both images we added a fixed gradient to the background of one image and increased the gradient of the other until the

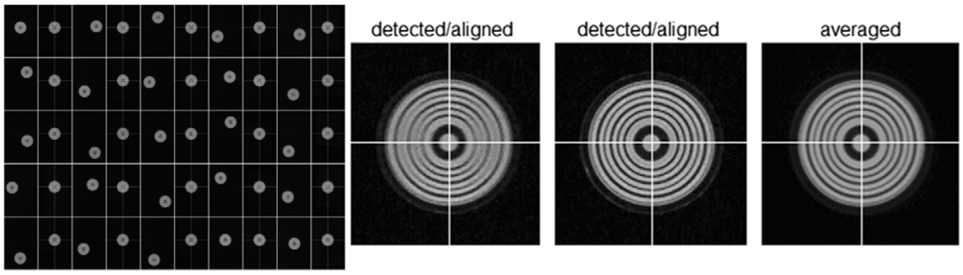

Figure 4. "Align and average" procedure with recorded images of a printed circular object against a black background. Left: the reference and 24 recorded images (odd columns) and the final result with the 24 aligned images (even columns). A crosshair has been included in the aligned images for visual inspection of the alignment quality (see text). Other images (left to right): enlarged copies of two recorded images and the final result.

alignment was lost. The threshold lines (center plot) are roughly straight, indicating that the sum of the two gradients is the determining factor. The threshold is somewhat higher when only one background shows a gradient because the misalignment is then due only to cross-correlation. The threshold is around 0.8 levels/pixel for a 50 pixel margin and decreases when the margin is widened. Minor shifts occur when the overall background level is changed or when a vertical gradient is added. Thus, a background change of $200 \times 0.8 = 160$ gray levels causes loss of alignment. In view of the small angle ($\sim$mrad) that is typically spanned by the object, it is unlikely that such a large gradient will occur in practice. Of course, the results depend on the type of image that is being used, so a higher sensitivity may occur in some cases.

It may appear that addition of only a constant background should have no effect on the alignment. After all, the correlation of an image with a constant is not affected by displacement. However, since the background only affects the area around the object, the above is not completely true. As shown in the right plot of Fig. 2, adding 170 gray levels to either background causes loss of alignment. This makes the background so much brighter that it correlates better with the object of the other image than the object itself. As can be seen, the effect is reduced when a constant background is also added to the other image. Again, the required threshold is large compared to the range (0–255) of gray levels and will not often occur in practice.

EXPERIMENTAL IMAGE

Figure 4 shows correlation alignment of a set of simple experimental images. For this initial case noise and background variation were avoided as much as possible. The object was a white cardboard circle (d = 15 cm) with printed black circles. Circular symmetry was used on purpose so that only translational alignment was necessary. The object was placed at various positions on a background of black cloth and photographed at 1.5 m distance with a compact camera. The pictures were taken in typical indoor lighting conditions without flash. As a result, the exposure time was several seconds. Some images are slightly blurred by vibrations of the camera. From each picture a window of 420×270 pixels was roughly selected by hand which always included the object, but not at the center. The object measures approximately 100 pixels across.

One of the experimental images was used as a reference and the other images were aligned to that image. To check the performance we drew a crosshair at the center of the reference image and copied it to the same fixed position on all aligned images. Inspection by eye indicates that the aligned objects are off by at most one pixel except for the most severely blurred cases. The final enhanced image was obtained by taking the average of the 24 aligned images. The result is smoother than the original images, but does not show any features that were obscured in the recorded images. The thinnest line in the object was the first dark circle from the outside. This is visible in the final result but also visible in most of the recorded images.

We must conclude that while the correlation alignment worked very well, the effects on the final enhanced image are limited. Nevertheless, some enhancement can be observed in Fig. 4. In the blurred images, parts of the rings are obscured. This problem is not present in the final resulting picture. At least, the recorded images of lower quality do not affect the final result. Perhaps more important is, however, that a small amount of noise can be observed in the recorded images and also some staining in the white and black areas of the object. Apparently, this noise did not affect the alignment, and both the noise and the staining have been reduced in the final resulting image.

DISCUSSION

In the above we studied correlation alignment of multiple images. We focus on long range imaging with objects typically several meters to tens of meters in size at a distance of several kilometers, which will thus span an angle of around 1–10 mrad. We furthermore assume that the object has been detected and roughly located in a series of images so that windows can be extracted which contain the object and a limited amount of background. After alignment, a single enhanced image can be obtained, e.g., for identification of the object. The object will be relatively featureless and distorted by noise but (due to the small size) the background may be relatively constant. Because of this, correlation appears to be the useful tool to perform the alignment.

Using simulated images we established that accurate correlation alignment can be obtained, at least in theory, even when the noise is so severe (90% of the pixels obscured) that the object cannot be distinguished by the naked eye. This may occur only in extreme cases. In a more realistic scenario, contrast and detail are reduced by absorption, scattering and turbulence in the atmosphere, which indeed reduce the noise resistance of the image. The ratio of noise- and image-covered areas at the threshold level may be proportional to the contrast in the image. To our surprise we found that including more background in the window may actually reduce alignment errors at least when the background is constant. Possibly, this due to averaging of the random noise. Alignment may be most accurate when 80–90% of each image is occupied by background.

The resistance of correlation alignment to a background gradient can be quite strong. We found that alignment is lost only when the background gradient across the object is of the same order as the gray level range within it. Surprisingly, miss-alignment can also occur if a background gradient is present in only one of the images (i.e., even if a reference with a constant background is used) and even without gradients with a sufficiently strong difference in background gray level between the images. In view of the small angle typically spanned the above may not usually occur, unless the contrast of the object is very small. In that case it may be necessary to reduce background variation, e.g., with high-pass filters[1] or local normalization.[10]

A theoretical treatment is complicated because neither noise nor background is additive, making it difficult to compare what happens with the images and the correlation. The correlation of the objects in the two images consists of a main peak that represents the displacement between them. Side peaks occur if different parts of the object are similar. The width the main peak depends on the details in the image. The correlation of the noise will form a homogeneous random spike pattern. Alignment is lost if one of the spikes is higher than the main peak. Alignment error occurs by spikes on top of a relatively broad alignment peak. This is in line with our observations: low noise intensities lead to minor alignment errors, while the alignment is lost if the noise reaches a threshold value. The background effects may be considered with a similar model. While the addition of object, noise and background correlation is not strictly true it does provide a useful description of our observations.

The purpose of the work presented in this article was not so much to provide the best possible alignment enhancement, but rather to investigate which parameters affect the outcome and quality of the results. Nevertheless, it may be useful to try to compare the performance to that of similar and other techniques. Our main performance result is that it may be possible to align images to within a few pixels (50×200 pixel object) when up to 90% of the pixels are obscured by noise and the object can no longer be recognized by the naked eye. This appears to be affected mainly by the amount of contrast in the image and not so much by other parameters such as blurring, the amount of background included in the image and even background gradients. Besides contrast, the performance relies on the fact that the underlying image of the object is the same so that information from the entire image can be used for the alignment. A little differently from our original thoughts, our present approach may therefore be most useful not so much for extremely long range imaging, where blurring but also image deformation is caused by air turbulence,[5,9] but rather when the object is hidden by foreground clutter such as, for example, precipitation or (on land) small objects or foliage.[8,22]

Correlation alignment and registration is routinely applied to medical images[23] and remote sensing.[24] However, this typically involves alignment of the entire image, rather than a small sub-image, and techniques such as elastic alignment are applied to compensate for small image deformation in the detector. Perhaps a better example is made with face alignment studies.[25,26] Yigang et al.[25] were able to obtain alignment to within 0.5 pixels in a set of 100 images (49×49 pixels) of a dummy head using a robust alignment procedure with sparse decomposition. This advanced analysis does require a rather long computation time (24 min for the set). In Tzimiropoulos et al.,[26] artificially distorted facial images are aligned to within 1 pixel using robust alignment based on gradient ascent. It should be noted that these studies focus more on image displacement and warping and less on resistance to clutter. As far as the comparison is justified, these results still outperform ours, although we also obtain 1 pixel rms shifts in lower noise cases.

We have studied the basic characteristics for correlation alignment in long range target identification. Rotational correlation can be included with some computational effort. Combined with translation this may be useful in practical applications where the image of the object itself is relatively constant.[4,6] More fundamental are changes of aspect ratio that result from a maneuvering target. An extension of our approach is then to find groups of images, within the data, that correlate among themselves. Distortions may be compensated,[12,21] while optical flow[19,20] or patch tracking[15] may detect independent displacement within the object's image. Our results have shown that correlation alignment may be effective when little information is available in the image and have illustrated how the alignment quality is affected by various parameters such as noise strength and the amount of background included in the image. This information should be useful in the analysis of experimental images and the development of practical applications.

ACKNOWLEDGMENTS

The authors are very grateful to two unknown reviewers. Their comments helped us considerably with the focus of this article and literature.

REFERENCES

[1] P. J. Burt, "Smart sensing with a pyramid vision machine," Proc. IEEE **76**, 1006–1015 (1998).

[2] G. Bond and R. Seyfarth, "Improving image alignment in aerial image mosaics via error estimation of flight attitude parameters," Comput. Sci. Eng. **2**, 86–91 (2012).

[3] A. Buades, B. Coll, and J. M. Morel, "A review of image denoising algorithms, with a new one," Multiscale Model. Simul. **4**, 490–530 (2005).
[4] X. H. Chen and N. A. Schmid, "Automatic target recognition performance losses in the presence of atmospheric and camera effects," J. Electron. Imaging **19**, (2010), Article nr 023016.
[5] P. Chervet and A. Rorblin, "High-altitude cloud effects on airborne electro-optical sensor performance," J. Atmos. Oceanic Technol. **23**, 1530–1538 (2006).
[6] C. Corbane, L. Najman, and E. Pecoul, "A complete processing chain for ship detection using optical satellite imagery," Int. J. Remote Sens. **31**, 5837–5854 (2010).
[7] G. D. Evangelidis and E. Z. Psarakis, "Parametric image alignment using enhanced correlation coefficient maximization," IEEE Trans. Pattern Anal. Mach. Intell. **30**, 1858–1865 (2008).
[8] U. Grenander and A. Srivastava, "Probability models for clutter in natural images," IEEE Trans. Pattern Anal. Mach. Intell. **23**, 424–429 (2001).
[9] S. Gepshtein, A. Shtainman, B. Fishbain, and L. Yaroslavsky, "Restoration of atmospheric turbulent video containing real motion using elastic image registration," Proc. European Signal Process. Conf. 477–480 (2004).
[10] M. Irani and P. Anandan, "Robust multi sensor image alignment," *Proc. Sixth Int'l. Conf. on Computer Vision* (1998), pp. 959–966.
[11] H. P. D. S. Jun and S. G. Sun, "Automatic target detection using binary template matching," Opt. Eng. **44**, (2005), Article nr 036401.
[12] H. S. Li, T. A. Shen, and X. L. Huang, "Approximately global optimization for robust alignment of generalized shapes," IEEE Trans. Pattern Anal. Mach. Intell. **33**, 1116–1131 (2011).
[13] X. Lu and R. Manduchi, "Fast image motion segmentation for surveillance applications," Image Vis. Comput. **29**, 104–116 (2011).
[14] S. C. Park, M. K. Park, and M. G. Kang, "Super-resolution image reconstruction: a technical overview," IEEE Signal Process. Mag. **20**, 21–36 (2003).
[15] F. Rothganger, S. Lazebnik, C. Schmid, and J. Ponce, "3D object modeling and recognition using affine-invariant patches and multi-view spatial constraints," Proc. IEEE Comput. Soc. Conf. on Comput. Vis. Pattern Recognit. **18–20**, 272–277 (2003).
[16] O. J. G. Somsen and F. Bolderheij, "Object alignment for military optical surveillance," Proc. Int. Conf. on Pattern Recognit. Comput. Vis. 273–276 (2011).
[17] N. A. Woods, N. P. Galatsanos, and A. K. Katsaggelos, "Stochastic methods for joint registration, restoration, and interpolation of multiple undersampled images," IEEE Trans. Image Process. **15**, 201–213 (2006).
[18] G. Yang, C. V. Stewart, M. Sofka, and C.-L. Tsai, "Alignment of challenging image pairs: refinement and region growing starting from a single keypoint correspondence," IEEE Trans. Pattern Anal. Machine Intell. (2007).
[19] L. Yaroslavsky, B. Fishbain, G. Shabat, and I. Ideses, Super-resolution in turbulent videos: making profit from damage (2007).
[20] L. P. Yaroslavsky, G. Shabat, B. Fishbain, and I. A. Ideses, "Super-resolution of turbulent video: potentials and limitations," Proc. SPIE **6812**, (2009).
[21] Z. Yi, C. Zhiguo, and X. Yang, "Multi-spectral remote image registration based on SIFT," Electronics Lett. **44**, 107–108 (2008).
[22] O. J. G. Somsen, "Analysis of a basic multiple-image algorithm for removal of foreground objects," *Proc. Third Int'l. Conf. on Digital Information and Communication Technol. and its Applications* (2013) pp. 248–255.
[23] Z. S. Saad, D. R. Glen, G. Chen, M. S. Beauchamp, R. Desai, and R. W. Cox, "A new method for improving functional-to-structural MRI alignment using local pearson correlation," Neuroimage **44**, 839–848 (2009).
[24] S. R. Lee, "A coarse-to-fine approach for remote-sensing image registration based on a local method," Intl J. Smart Sensing and Intell. Syst. **3**, 690–702 (2010).
[25] Y. Peng, A. Ganesh, J. Wright, W. Xu, and Y. Ma, "RASL: Robust alignment by sparse and low-rank decomposition for linearly correlated images," Proc. IEEE Conf. Comput. Soc. Conf. on Comput. Vis. Pattern Recognit. 763–770 (2010).
[26] G. Tzimiropoulos, S. Zafeiriou, and M. Pantic, "Robust and efficient parametric face alignment," Proc. IEEE Conf. Comput. Soc. Conf. on Comput. Vis. Pattern Recognit. 1847–1854 (2011).