

Superresolution reconstruction for moving point target detection

Judith Dijk

TNO Defence, Security and Safety
Electro-Optics Group
P.O. Box 96864, 2509 JG
The Hague, The Netherlands
E-mail: judith.dijk@tno.nl

Adam W. M. van Eekeren

TNO Defence, Security and Safety
Electro-Optics Group
P.O. Box 96864, 2509 JG
The Hague, The Netherlands
and
Delft University of Technology
Faculty of Applied Sciences
Quantitative Imaging Group
Lorentzweg 1, 2628 CJ
Delft, The Netherlands

Klamer Schutte

Dirk-Jan J. de Lange

TNO Defence, Security and Safety
Electro-Optics Group
P.O. Box 96864, 2509 JG
The Hague, The Netherlands

Lucas J. van Vliet

Delft University of Technology
Faculty of Applied Sciences
Quantitative Imaging Group
Lorentzweg 1, 2628 CJ
Delft, The Netherlands

1 Introduction

In surveillance applications moving targets need to be detected at a very early stage. Electro-optical surveillance systems observe missiles or other incoming threats as moving point targets. At maximum detection range these point targets will have a low signal-to-noise ratio with respect to the background. Furthermore, the background may also contain structure (clutter) of high contrast.

Usually, the first step of point target detection is to suppress the clutter of the stationary background in the image. A clutter suppression step should remove the information of the static background while preserving the target signal energy.

One of the essential steps for background suppression is to determine the apparent motion between the frames, i.e., the registration step. The apparent motion of the background can, on a small scale, often be described by translational motion between two subsequent camera frames I_k and I_{k-1} .

Abstract. When bright moving objects are viewed with an electro-optical system at long range, they appear as small, slightly blurred moving points in the recorded image sequence. Typically, such point targets need to be detected in an early stage. However, in some scenarios the background of a scene may contain much structure, which makes it difficult to detect a point target. The novelty of this work is that superresolution reconstruction is used for suppression of the background. With superresolution reconstruction a high-resolution estimate of the background, without aliasing artifacts due to undersampling, is obtained. After applying a camera model and subtraction, this will result in difference images containing only the point target and temporal noise. In our experiments, based on realistic scenarios, the detection performance, after background suppression using superresolution reconstruction, is compared with the detection performance of a common background suppression method. It is shown that using the proposed method, for an equal detection-to-false-alarm ratio, the signal strength of a point target can be up to 4 times smaller. This implies that a point target can be detected at a longer range. © 2008 Society of Photo-Optical Instrumentation Engineers. [DOI: 10.1117/1.2977790]

Subject terms: superresolution reconstruction; detection; point target; background suppression; ROC.

Paper 080129RR received Feb. 14, 2008; revised manuscript received Jul. 1, 2008; accepted for publication Jul. 2, 2008; published online Sep. 18, 2008.

A standard way of performing background suppression is to shift, interpolate, and subtract (SIS). One of the frames is corrected for the shift (dx, dy) using interpolation (\tilde{I}_{k-1}) and is subtracted from the other frame. In the experiments in this paper we use B-spline interpolation. After subtraction a difference image ΔD_k^{SIS} results:

$$\Delta D_k^{\text{SIS}}(x, y) = I_k(x, y) - \tilde{I}_{k-1}(x + dx, y + dy). \quad (1)$$

Note that for point targets with a small apparent motion with respect to the background, the point target's signal energy in the difference image ΔD_k^{SIS} is almost lost. Another problem of SIS is that due to undersampling by the image sensor, aliasing artifacts in the recorded image sequence remain in the difference image. Both will hamper point target detection.

In this paper we propose to use superresolution (SR) reconstruction to improve the detection of moving point targets. The SR reconstruction algorithm is used in the background suppression step. In previous work^{1,2} we developed SR reconstruction techniques to improve the spatial

resolution of undersampled image sequences by exploiting the subpixel shift between the frames. Using SR for point target detection has the advantages that (1) the signal and aliasing contributions in the last frame can be predicted, which substantially reduces the aliasing-related clutter in the difference image, (2) the temporal noise is reduced, which improves the amplitude-to-noise ratio (ANR) of the point target in the difference image, and (3) the ANR in the background estimate is suppressed, which increases the ANR in the difference image. The latter is especially noticeable for point targets with a small apparent motion with respect to the background. Note that point targets that have *no* motion with respect to the background will be totally included in the background estimate. Therefore, they are not visible in the difference image.

After SIS or the SR background suppression step, standard detection algorithms such as thresholding or track-before-detect^{3,4} can be used. In this paper results are shown for a three-out-of-five tracking algorithm⁵ and for direct thresholding.

This paper is organized as follows. In the next section the advantages of SR reconstruction for point target detection are discussed from a theoretical perspective. In Sec. 3 the SR-based point target detection method is presented. In Sec. 4 the setup of these experiments is described. The experimental results are shown in Sec. 5. Finally, conclusions are presented in Sec. 6.

2 Theory

SR reconstruction is a well-known technique to increase the spatial resolution of a sequence of aliased low-resolution (LR) images using temporal information. The *zoom factor* of a SR reconstruction method is the ratio of the size of the resulting high-resolution (HR) image to the size of the LR images. Numerous SR reconstruction methods are described in the literature. Overviews are given by Park et al.,⁶ Farsiu et al.,⁷ and Van Eekeren et al.⁸ Generally, SR reconstruction can be split into three parts:⁹ (1) registration, (2) fusion, and (3) deblurring. The first part is necessary to align the content of all frames with subpixel accuracy. The two following steps will fuse the aligned data on a HR grid and deblur the result.

SR reconstruction can be used for point target detection to improve the background suppression step. With SR reconstruction it is possible to create a HR model of the background. This HR background model contains less or no aliasing and ideally does not contain the point target. With the HR background model \mathbf{Z} (reordered in a vector) and a transfer matrix \mathbf{H}_k , an estimate can be made of LR image \mathbf{I}_k (reordered in a vector). The transfer matrix \mathbf{H}_k describes (1) the model of the camera, (2) the estimated motion between \mathbf{Z} and \mathbf{I}_k , and (3) the zoom factor. A difference image of frame k is then created by applying

$$\Delta \mathbf{D}_k^{\text{SR}} = \mathbf{I}_k - \mathbf{H}_k \mathbf{Z}. \quad (2)$$

Here $\mathbf{H}_k \mathbf{Z}$ suffers from aliasing exactly the same way as \mathbf{I}_k , so the difference image is free from clutter due to aliasing. If the images were subtracted in the HR space, this would not be the case, for the background image \mathbf{Z} is aliasing-free and the HR version of \mathbf{I}_k contains the interpolated aliasing.

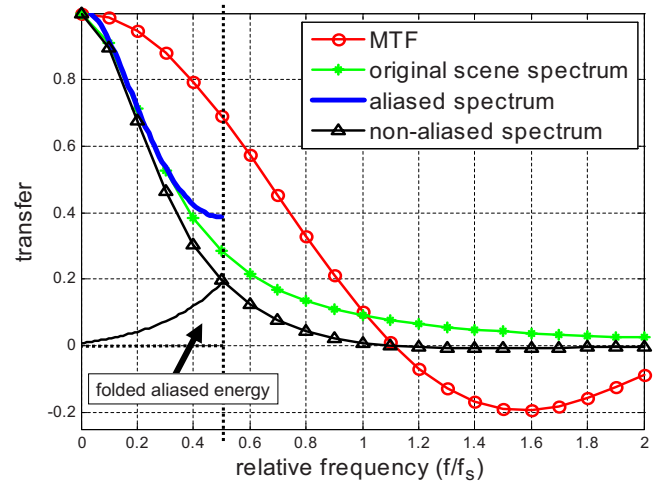


Fig. 1 The aliased spectrum (no marks) of a signal that is bandlimited and undersampled. The modulation transfer function (MTF) is modeled by Gaussian lens blur ($\sigma_{\text{psf}}=0.3$) and uniform sensor blur (area fill factor=81%). The nonaliased spectrum (marked with triangles) results after applying the MTF to the original spectrum. To obtain the aliased spectrum, the spectral energy above the half sampling frequency $f_s/2$ needs to be folded (marked with the arrow) and added to the nonaliased spectrum. Note that the aliased spectrum does not have any information above the half sampling frequency. All spectra are normalized so that the dc value equals one.

In the next subsections the advantages of SR reconstruction for point target detection are explained from a theoretical perspective.

2.1 Aliasing Noise Reduction

In a camera system the measured signal is limited by (1) the band limitation of the optics and (2) the sampling of the sensor. Aliasing is an effect due to undersampling. Both limitations are depicted in Fig. 1.

Here, the blurring of the lens is modeled with Gaussian blurring ($\sigma_{\text{psf}}=0.3$), and the sensor is modeled as a 2-D array of nonoverlapping square photosensitive elements with fill factor 81%. The scene spectrum is modeled with a quadratic decay, which is characteristic of natural images.¹⁰ Note that this may differ slightly from the real scene spectrum of the images used in the simulations of which no spectrum was determined.

Applying SR reconstruction increases the sampling rate so that the aliased frequency spectra are unfolded and part of the high-frequency spectrum is recovered. This implies that a better—i.e., (almost) aliasing-free—HR estimate of the background is obtained. Applying the camera model for frame k to this HR background image yields the same aliasing artifacts as in the recorded image I_k . Subtraction of these two images is very effective in suppressing the background, because by sampling the HR image at exactly the same grid positions (including subpixel shift) as the corresponding LR image, exactly the same aliasing artifacts for frame k are created. After subtraction, the difference image will contain only temporal noise and the point target signal. Note that the main aliasing effect on the point targets is that their maximum energy per frame is not constant.

2.2 Temporal-Noise Reduction

All cameras add temporal noise to the scene information. Let us assume that there are N recorded frames available, containing additive Gaussian-distributed noise with standard deviation σ_n . The resulting noise in a difference image after SIS [Eq. (1)] will be $\sqrt{2}\sigma_n$. Note that we assume that the images are corrected for nonuniformity. This can be done by the camera or based on the images. In this paper, we do not evaluate the effects of nonuniformity reduction by SR on the detection of point targets.

Now, assume that SR reconstruction is used to calculate a difference image as in Eq. (2). Here, the estimated LR frame $\mathbf{H}_k\mathbf{Z}$ is based on N recorded frames, which reduces the noise standard deviation by a factor \sqrt{N} . Therefore, the resulting noise in a difference image after SR reconstruction is

$$\sigma_n^{\text{ASR}} = \left(\frac{N+1}{N}\right)^{1/2} \sigma_n. \quad (3)$$

If many frames are used, the noise in the resulting difference image will be only slightly higher than σ_n , the noise in a single LR image. The ratio of noise in the SR difference image to that in the SIS difference image is

$$\frac{\sigma_n^{\text{ASR}}}{\sigma_n^{\text{ASIS}}} = \frac{\left(\frac{N+1}{N}\right)^{1/2} \sigma_n}{\sqrt{2}\sigma_n} \approx \frac{1}{\sqrt{2}}. \quad (4)$$

This means that for the same rate of false detections, the point target amplitude that can be detected with SR reconstruction in a temporal-noise-limited situation will be a factor $\sqrt{2}$ lower.

2.3 Point Target Amplitude Preservation

Another advantage of background suppression using SR reconstruction is that the point target intensity in the difference image is preserved for large and small apparent motion of the point target with respect to the background. Ideally, the point target is not present in the projection of the HR background image, i.e., the point target intensity in the difference image is the same as the amplitude of the point target. The difference image is calculated with Eq. (2). In the nonideal case, however, the point target can be visible in the projection of the HR image \mathbf{Z} . Note that the main aliasing effect on the point targets is that their maximum energy per frame is not constant.

To analyze the point target amplitude preservation, point targets with amplitude one are simulated. First, point targets are placed in a superscale image with a constant background, which is a factor 15 larger than that of the LR image \mathbf{I}_k . This superscale image is shifted and blurred with the MTF as described in Fig. 1 and afterwards subsampled with factor 15. The simulation method is fully explained in Sec. 4.3. To the resulting LR images a small amount of Gaussian-distributed noise ($\sigma_n=0.002$) is added.

The effect of point target amplitude preservation is measured as the maximum of the difference images ΔD_k^{SR} . If this maximum is around one, the point target amplitude is well preserved. The point target intensity in the difference image is simulated for two different SR reconstruction

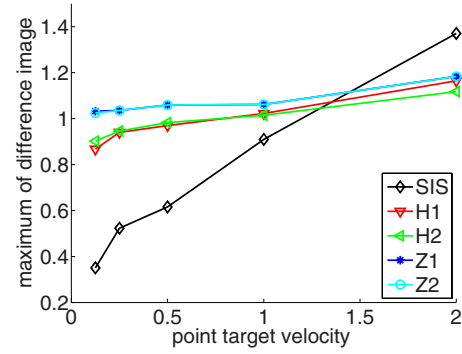


Fig. 2 Relative point target amplitude in a difference image as a function of the point target velocity. The relative point target amplitude for a difference image resulting from SIS (B-spline) is indicated with diamonds. Zomet (robust SR) is indicated with Z1 and Z2 (zoom factor 1 and 2 respectively), and Hardie (nonrobust SR) is indicated with H1 and H2. A total of 48 frames are used for the SR reconstruction, and Gaussian blur with $\sigma=\sigma_{\text{cam}}=0.41$ is used to describe the camera blur.

methods: Hardie (nonrobust) and Zomet (robust). Both methods model the camera blur with Gaussian blur ($\sigma_{\text{cam}}=0.41$) and use 48 frames for reconstruction. A more detailed explanation of both methods is given in Sec. 3.2.

As a comparison the point target amplitude preservation of SIS for varying point target motion is simulated as well. Here, the difference images ΔD_k^{SIS} are calculated with Eq. (1). The results are shown in Fig. 2. As expected, the robust Zomet method (Z1 and Z2) performs best for point targets with small apparent motion with respect to the background. There is no significant effect between the different zoom factors. Note that the robust Zomet method preserves the point target amplitude better than the nonrobust Hardie method. This can be explained by the fact that the point target is treated as an outlier in Zomet's SR reconstruction. The difference between robust and nonrobust SR is explained in more detail in Sec. 3.2. If the point target velocity is large (>1.5 pixels/frame), the point target profile of the previous recorded frame will hardly influence the point target profile in the current frame. For those cases the point target amplitude in the difference image is maximal. The high maximum value in the difference image after SIS for a high point target velocity can be explained by the B-spline interpolation that is used for the shift. This interpolation can cause lobes that are below the background and add to the point target in the noninterpolated frame.

Summarizing, using SR reconstruction for background suppression has the following advantages from a theoretical perspective: (1) aliasing artifacts are reduced, (2) temporal noise is reduced, and (3) the point target is better preserved in the difference image for small apparent motion with respect to the background. Therefore, the largest gain from using SR reconstruction for background suppression is expected for recorded sequences with much structure in the background (causing significant aliasing artifacts) and a small apparent point target motion.

3 Point Target Detection Using Superresolution Reconstruction

This section describes the point target detection method based on background suppression using SR reconstruction.

Although this method is based on existing, well-known techniques, the combination and use of those techniques is innovative. First, the registration method is described, followed by the SR reconstruction method and finally the detection and tracking methods.

3.1 Registration

Registration aligns the content of all LR frames prior to SR reconstruction. This registration step is also needed for background subtraction with SIS. There are a variety of image registration techniques described in the literature.¹¹ We perform registration with a very precise iterative gradient-based shift estimator.¹² This estimator¹³ finds the displacement (dx_{k1}, dy_{k1}) between two shifted images, $\mathbf{I}_{k-1}(x, y)$ and $\mathbf{I}_k(x, y)$, as a least-squares solution:

$$\text{it: 1: } \min_{dx_{k1}, dy_{k1}} \frac{1}{P} \sum_{x, y} \left(I_k - I_{k-1} - dx_{k1} \frac{\partial I_{k-1}}{\partial x} - dy_{k1} \frac{\partial I_{k-1}}{\partial y} \right)^2. \quad (5)$$

Here, the image \mathbf{I}_k is approximated with a Taylor expansion of \mathbf{I}_{k-1} , (x, y) are the pixel positions, and P is the number of pixels in \mathbf{I}_k . The partial derivatives are calculated with a Gaussian gradient filter (see p. 64 in Ref. 14).

The solution of Eq. (5), (dx_{k1}, dy_{k1}) , is biased; the bias is corrected in an iterative way:

$$\text{it: } n: \min_{dx_{kn}, dy_{kn}} \frac{1}{P} \sum_{x, y} \left(\tilde{I}_k(x + dx_{k(n-1)}, y + dy_{k(n-1)}) - I_{k-1} - dx_{kn} \frac{\partial I_{k-1}}{\partial x} - dy_{kn} \frac{\partial I_{k-1}}{\partial y} \right)^2. \quad (6)$$

In iteration n ($n > 1$), \mathbf{I}_k is translated by interpolation (indicated by the tilde) with the estimated subpixel displacement $(dx_{k(n-1)}, dy_{k(n-1)})$ from the previous iteration. Now, the displacement (dx_{kn}, dy_{kn}) between the shifted \mathbf{I}_k and \mathbf{I}_{k-1} is estimated. This displacement is accumulated with the displacement obtained in the previous iteration. This schema is iterated until convergence and results in a very precise ($\sigma_{\text{disp}} \approx 0.01$ pixel for noise-free data) unbiased registration.¹² The total estimated displacement with the iterative gradient-based shift estimator after M iterations is

$$(dx_k, dy_k) = (dx_{k1}, dy_{k1}) + \dots + (dx_{kM}, dy_{kM}). \quad (7)$$

Note that this registration method, due to its iterative character, can also cope with multiple-pixel image shifts. In such a case, the registration will not be accurate after the first iteration, because the Taylor expansion is not accurate for large shifts. However, after a few iterations the remaining shift will be small and hence the Taylor expansion becomes accurate.

3.2 Robust Superresolution Fusion and Deblurring

The second and third step of superresolution reconstruction are fusion and deblurring. Numerous SR reconstruction methods can be found in the literature; some methods work in the Fourier domain,^{15,16} there exist both robust methods¹⁷ and nonrobust methods,¹⁸ and some methods² are adaptive. Van Eekeren et al.⁸ made a quantitative per-

formance comparison between a selection of different SR reconstruction methods. One of the best-performing methods is the one proposed by Hardie et al.¹⁸ Like many other SR reconstruction methods, it models the image formation process in the following way:

$$\mathbf{I}_k = \mathbf{D}_k \mathbf{C}_k \mathbf{F}_k \mathbf{Z} + \boldsymbol{\theta}_k = \mathbf{H}_k \mathbf{Z} + \boldsymbol{\theta}_k, \quad (8)$$

where \mathbf{I}_k is the k 'th LR frame, \mathbf{Z} is the HR image scene, and $\boldsymbol{\theta}_k$ is normally distributed additive noise, all reordered in vectors; \mathbf{F}_k is the geometric warp matrix based on the results of the registration; \mathbf{C}_k is the blurring matrix of the camera; and \mathbf{D}_k is the decimation matrix, which resamples the image to low resolution. For simplification all matrices are combined in \mathbf{H}_k . The blurring of the camera is modeled by Gaussian blurring. Note that it is allowed to represent basic operations such as warping and blurring in a matrix, because they are linear in the image intensities.

As already stated in Sec. 2.3, a robust SR algorithm for background suppression is proposed because the point target is better preserved in the difference image. A robust algorithm is less sensitive to outliers in the background data, such as moving point targets. With enough frames available and sufficient apparent motion of the point target with respect to the background, a robust SR algorithm will treat the point target as an outlier. For this reason we use a robust method, proposed by Zomet et al.,¹⁹ in the experiments. This method is similar to Hardie's but uses robustness in the minimization procedure.

This is best explained by comparing the minimization procedure used by Hardie with that used by Zomet. We start by giving a short derivation of Hardie's method, and then stress the differences from Zomet's.

The total squared error of resampling the HR image \mathbf{Z} is given by

$$L(\mathbf{Z}) = \frac{1}{2} \sum_{k=1}^N \sum_i [\mathbf{I}_k(i) - (\mathbf{H}_k \mathbf{Z})(i)]^2 \quad (9)$$

with N the total number of LR frames and i the LR pixels. Note that \mathbf{Z} is based on all the \mathbf{I}_k 's. Taking the derivative of L with respect to \mathbf{Z} results in

$$\nabla L(\mathbf{Z}) = \sum_{k=1}^N \mathbf{H}_k^T (\mathbf{H}_k \mathbf{Z} - \mathbf{I}_k) = \sum_{k=1}^N \mathbf{G}_k \quad (10)$$

with \mathbf{H}_k^T the transpose of \mathbf{H}_k . A gradient-based iterative minimization method updates the estimation in each iteration n by

$$\mathbf{Z}^{n+1} = \mathbf{Z}^n + \epsilon \nabla L(\mathbf{Z}) \quad (11)$$

with ϵ the step size in the direction of the gradient. The procedure can be seen as a version of the iterated back-projection method.²⁰ In each iteration the difference between the resampled HR image $\mathbf{H}_k \mathbf{Z}$ and the LR image \mathbf{I}_k is projected back to the HR grid.

A replacement of the sum of backprojected images \mathbf{G}_k in Eq. (10) with a scaled pixelwise median introduces the robustness of Zomet's method:

$$\nabla L(\mathbf{Z}) \approx N \cdot \text{median}(\mathbf{G}_k)_{k=1}^N. \quad (12)$$

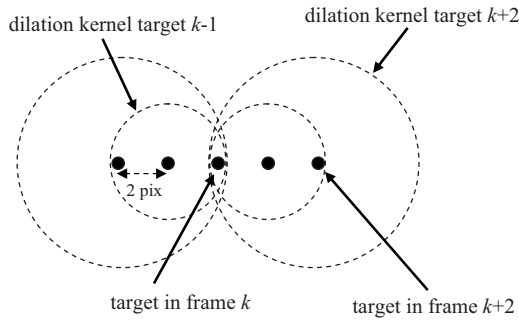


Fig. 3 The different kernel sizes used for the special dilation in the three-out-of-five tracking algorithm. The smallest kernel size is 2 pixels, which supports tracking of point targets with an apparent motion up to 2 pixels/frame.

3.3 Detection and Tracking of Point Targets

After background subtraction, the point targets need to be detected in the difference image. The difference images show the amplitude difference between a moving target and its local background, noise, and aliasing artifacts (SIS method). On such a difference image the detection of the objects is done. The simplest detection technique is to threshold the magnitude of the difference image. All pixels with a value above a certain threshold are detected as targets:

$$T_k = \begin{cases} 0 & \text{for } \Delta D_k(x,y) \leq \text{threshold}, \\ 1 & \text{for } \Delta D_k(x,y) > \text{threshold}. \end{cases} \quad (13)$$

This detection method works well for targets that have a sufficiently high ANR, so that the moving object can be detected in every frame. However, objects having an amplitude close to the local background may not be detected this way.

The specificity of a detection algorithm can be increased if it is performed on a series of subsequent difference images instead of a single difference image. Tracking can be used to associate the detections in the images. It is assumed that a target path is a continuous path over time, which means that the positions of correct detections are highly correlated over the frames. Uncorrelated detections are unlikely to be correct detections.

In this paper a three-out-of-five tracking algorithm⁵ is used to increase the specificity. This tracking algorithm performs first a special dilation on five successive frames after thresholding. This allows a limited displacement of the point target in the next frames. In order to keep the moving point target in track, a dilation with different kernel sizes is performed on frames 1, 2, 4, and 5. The center frame, 3, is not dilated. Kernel sizes are chosen such that point targets that have an apparent motion with respect to the background of maximum 2 pixels/frame can be tracked (see Fig. 3). Afterwards a pixelwise summation is performed. Pixels with a sum larger than or equal to 3 are marked as detections after tracking. This means that the targets are present in at least three of the five frames. Note that this tracking algorithm can be improved (using velocity and heading of the point target), which will result in a further reduction of the number of false alarms without losing sensitivity.

4 Experimental Setup

The performance of the different algorithms is tested on images containing point targets. These images are constructed by inserting simulated point targets in a real image sequence. To simulate a realistic scenario, first a few real-world scenarios are analyzed.

4.1 Real-World Scenario

An incoming missile at long distance is observed as a point in a recorded sequence. Such a missile must be detected as early as possible. In this analysis two missiles, a Stinger²¹ and an AA-10,²² are chosen, because most of these missiles are radar-silent. This means that they must be detected using an electro-optical sensor. First, let us analyze the observed velocity of missile. The apparent velocity (expressed in radians per second) of a missile with respect to the background from the observer's point of view can be described with

$$v_t = \frac{v}{d} \sin \alpha. \quad (14)$$

Here, v is the velocity of the missile, d is the distance to the observer, and α is the angle between the missile's path and the shortest path to the observer. The apparent motion of a missile in camera coordinates depends on the instantaneous field of view (IFOV) and the frame rate of the camera.

Realistic specifications of an infrared camera for the task of missile detection are: a center wavelength of $4 \mu\text{m}$, IFOV=1.5 mrad, a frame rate of 15 Hz, and a sensitivity of 0.025 K. The last determines the amount of noise radiance. With these specifications a missile, such as a Stinger, flying at Mach 2 at 2 km with $\alpha=10$ deg, has an apparent motion of approximately 2.6 pixels/frame. For a scenario with a larger range (e.g., an AA-10 flying at Mach 4 at 60 km with $\alpha=10$ deg) the apparent motion is approximately 0.17 pixels/frame.

The observed missile intensity depends on its radiated energy at a certain wavelength. Propagation losses are ignored in our analysis. The observed missile and background are regarded as blackbodies from which the radiated energy per unit time per unit surface area per unit wavelength can be calculated with Planck's law.²³ The total observed radiance E_t on one sensor element is defined as

$$E_t = \beta E_m + (1 - \beta) E_{bg}, \quad (15)$$

with E_m the radiance of the missile, E_{bg} the radiance of the background, and β the area fraction of the missile. The difference in radiance between an observed missile and its background is defined as $\Delta E_m = \beta(E_m - E_{bg})$. For two different real-world scenarios ΔE_m is calculated and compared with the clutter radiance and the noise radiance:

- *High-clutter scenario:* In this scenario a Stinger is fired from the ground to an air target at 3 km. The temperature of the background is chosen to be 290 K, and the clutter is chosen to be $\Delta T_{cl}=1$ K. The Stinger has a velocity of Mach 2 and has a diameter of 7 cm. Its velocity determines its aerodynamic temperature²⁴ to be 480 K. The diameter of the missile, the distance to the target, and the IFOV and MTF of the camera

determine the area fraction β to be 1.9×10^{-4} . This results in the following ratio between the missile's radiance and the noise radiance: $\Delta E_m/E_n \approx 24$ with E_n the noise radiance. The ratio between clutter radiance and noise radiance is $\Delta E_{cl}/E_n \approx 41$. This indicates a clutter-dominated scenario.

- *Low-clutter scenario*: In this scenario an AA-10 is fired from an aircraft at 9-km altitude to an air target at a distance of 60 km. The air temperature is estimated (with a vertical gradient of -10 K/km) to be 203 K, and the clutter is estimated to be $\Delta T_{cl}=0.1$ K. The AA-10 flies at Mach 4 and has a diameter of 23 cm. Its velocity determines its aerodynamic temperature²⁴ to be 736 K. The diameter of the missile, the distance to the target, and the IFOV, and MTF of the camera determine the area fraction β to be 5.1×10^{-6} . This results in the following ratio between the missile's radiance and the noise radiance: $\Delta E_m/E_n \approx 9$. The ratio between clutter radiance and noise radiance is $\Delta E_{cl}/E_n \approx 0.04$. This indicates a noise-dominated scenario.

4.2 Simulated Scenario

The point target images used for the experiments are constructed from a real image sequence, which was recorded with an infrared camera (Radiance HS, 3 to 5 μm , 256×256 , 15 frames/s). The recorded images, which have an intensity range of [1000, 1172] gray values, contain noise with an estimated standard deviation of 1 gray value. Furthermore, they contain artifacts such as bad pixels and non-uniformity. Before inserting the point targets, the recorded image sequence is corrected for those two types of artifacts.¹ This will improve the detection results and will make it easier to compare the results of our experiments. The camera movement of the recorded sequence is approximated by a frame-to-frame translation, which is on average over the frames $v_x=3.10$ pixels/frame and $v_y=0.64$ pixels/frame.

4.3 Simulated Point Targets

The point targets are simulated and added to the LR camera images. First the point targets are placed in a superscale image, which is a factor 15 larger than the LR camera image. Here the position of the point target is integer-based. To obtain the LR image with the point target, a camera model is applied to the superscale image. In this camera model the MTF of the camera is modeled by a lens blur ($\sigma_{psf}=0.3$ LR pixel) and a fill factor (81% area). The camera MTF is plotted in Fig. 1. The camera model also subsamples the superscale image with a factor of 15. This subsampling is done by taking each 15th pixel. The resulting LR image contains the point target with aliasing. The maximum point target energy depends on the LR subpixel position of the LR image. We define the amplitude of a point target as the average maximum intensity of the point target in all available LR images. The point target simulation is visualized in Fig. 4. Here can also be seen that due to aliasing the maximum energy per frame of the point targets is not constant.

Adding the point target to the background instead of replacing the background introduces an error. In this simu-




Fig. 4 Two simulated point targets with different locations. In (a), the input point target in superscale is visualized. Note that the point is slightly larger than 1 pixel for visualization purposes. In (b), the camera model has been applied to the point target. It can be seen that the point target energy is spread over a large number of pixels. When the image is downsampled (c), the point target energy is still in more than one pixel.

lation the error is small, for reasons. First, the point target is placed in a superscale image and downsampled as described, instead of placing it directly into the LR image. In this way, the point target will suffer from aliasing in a similar way to the background. Second, the target is a point target and has therefore a small footprint. The error that is made by adding instead of replacing in the superscale image is $\Delta E_{add} = \beta \Delta E_{cl}$, which is the clutter radiance that is not replaced. In our scenarios, the worst case of ΔE_{add} is $E_n \times 1.9 \times 10^{-4} \times 41 \approx E_n \times 10^{-2}$. This means that the maximum error is 100 times smaller than the temporal noise.

For the experiments, the amplitude and apparent motion of the point target with respect to the background are varied according to the calculations of the different real-world scenarios. The point target amplitude is varied between 4 and 56 gray values, and the apparent motion of the point target is varied between 0.125 LR pixels per frame (almost no movement with respect to the background) and 2 LR pixels per frame. For each velocity eight different subpixel start locations of the point target are chosen. To simulate the different clutter scenarios, two different kind of sequences are constructed: one with the point target in a low-clutter region and one with it in a high-clutter region of the real image sequence. The upper parts (256×128) of a few frames of a constructed point target sequence are shown in Fig. 5. Here, the point target is placed in a low-clutter region with an apparent motion with respect to the background of 2 pixels/frame.

4.4 Processing Details

The constructed LR images are registered using the techniques presented in Sec. 3.1. The number of iterations used is 5. The σ of the Gaussian derivative filters is 1. Then the constructed LR images are processed by three different background suppression methods: (1) SIS with B-spline interpolation, (2) Zomet's robust SR reconstruction method with zoom factor 1, and (3) Zomet's robust SR reconstruction method with zoom factor 2. The camera model used in Zomet's method consists of a Gaussian blurring only. In our experiments this blurring has been set to $\tilde{\sigma}_{cam}=0.41$, which is the best Gaussian fit to the real camera model ($\sigma_{psf}=0.3$ and fill factor 81%). In the Zomet reconstruction 10 iterations are used. No regularization is used in the Zomet algorithm.




Fig. 5 Three frames (256×128) of a constructed point target sequence. The position of the point target can be seen in Fig. 6. The amplitude of the point target is 56 gray values. The point target is moving with an apparent velocity of 2 LR pixels per frame with respect to the background.

5 Results

Figure 6 shows the difference images for the three different background suppression methods. It can be seen that the difference image resulting after background suppression with Zomet's SR reconstruction method with zoom factor 2 contains much less background contributions than the other methods. This effect is best seen in the middle part of the image, where the structure of the buildings is hardly visible in comparison with the other two difference images. Furthermore it can be seen that both difference images based on Zomet's method contain less noise than the difference image based on SIS.

To evaluate the performance of the different methods under different scenarios, sequences of 48 frames are used. The different scenarios are created using: (1) different clutter levels (indicated with $CNR = \text{clutter noise ratio}$), (2) different point target amplitudes (indicated with $ANR = \text{amplitude noise ratio}$), and (3) different point target velocities (indicated with $PTV = \text{point target velocity}$). The last two can be controlled, because they are simulated, but the clutter level cannot. Therefore, as low-clutter region the sky area in the image is selected, and as high-clutter region the building area in the middle. The clutter is defined as the maximum gradient magnitude present in the region of interest (ROI) used for a specific scenario. In Fig. 7 the ROIs (one for high and one for low clutter) are visualized for frame 24. Note that in our analysis the point target is always present in the ROI. Furthermore, an external mask is used to mask the bad pixels in the original camera scene.

The detection results are presented in two different ways. First, the results are presented by means of receiver operating characteristic (ROC) curves. These curves represent the relation between the true-positive rate (sensitivity)

and the false-positive rate (1-specificity) for different threshold values. Next, the performance for the different algorithms is compared for a representative operating point.

5.1 ROC Curves

An ROC curve relates the sensitivity to the specificity of an algorithm. For a detection method an ROC curve can be determined by varying the threshold and counting the true and the false detections (knowing the ground truth). In our analysis, first the fraction of true detections is determined. A true detection occurs when in the point target ROI—a 5×5 neighborhood around the point target—a detection is present. Here, a detection is defined as one or more connected pixels after thresholding or tracking. Because every frame contains one point target, the number of true detections divided by the total number of frames equals the fraction of true detections. This is indicated on the vertical axis of the ROC curve.

The next part of the analysis is to determine the number of false detections. First, the true detections are removed if they are smaller than twice the size of the point target ROI. This is done to make sure that true detections are not counted as false detection as well, except when a true detection is large too. The latter situation can occur for small threshold values. After this removal, the number of false detections in each frame is determined after labeling (with 4-connected neighbors) all detections. In the ROC curves that are presented here, the number of false detections per second is plotted on the horizontal axis. These numbers correspond with a frame rate of 15 frames/s and a frame size of 256×256 pixels. Because the evaluation is done on




Fig. 6 Difference images for the different background suppression methods displayed with intensity range $[-6, 6]$. The positions of the point target are indicated with a circle. The difference image is shown for the 24th frame in the sequence. The amplitude of the point target is 12 gray values, and its apparent motion with respect to the background is 2 LR pixels per frame.




Fig. 7 Regions of interest in frame 24 that are used for analyzing the point target detection in a low- and high-clutter scenarios. All images are visualized in $[1000, 1120]$. Note that no simulated point target is present in these images.

a smaller, defined ROI (see Fig. 7), the measured numbers of false detections are scaled so that they correspond to a 256×256 -pixel frame.

In Fig. 8 the improvement in detection performance that can be obtained by using tracking is shown. The high-clutter scenario is used with a low point target velocity (0.125 pixel/frame) and an ANR of 12. It can be seen that tracking improves the detection results for background subtraction with both Zomet and the SIS method.

Figure 9 shows ROC curves of a low-clutter scenario. These curves show that the background subtraction methods using SR reconstruction outperform the SIS method. For fast-moving point targets in a low-clutter scenario (lower row of Fig. 9), the Zomet method performs better due to its noise reduction capabilities. For slow-moving point targets (upper row of Fig. 9) the Zomet method also performs better because the point target is efficiently suppressed in the background estimation, resulting in more point target energy in the difference image. Therefore, the improvement using Zomet is much larger for point targets with a low apparent velocity than for point targets with a high apparent velocity. As expected, the difference in performance between the two zoom factors of Zomet is not significant, because there is not much residual aliasing in the low-clutter scenario.

Figure 10 shows ROC curves of the high-clutter scenario. These results show the excellent performance of Zomet's SR method with zoom factor 2. This method reduces the aliasing artifacts, which are specifically present in the high-clutter region, much better than both other methods. This can also be seen from the difference images depicted in Fig. 6 in the building ROI. In the upper row of Fig. 10 (small apparent point target motion with respect to

the background), Zomet's SR method (both zoom factors) outperforms the SIS method. This can be explained by a higher point target amplitude in the difference image of Zomet's method. The difference image of Zomet with zoom factor 2 also contains less aliasing artifacts. Note that in Fig. 10(d) the SIS method performs better than Zomet's method with zoom factor 1 for small ANR. This is explained by the fact that the background estimation generated with the Zomet 1 method will have aliasing errors. Because this background estimation is used for the detection in all frames, these aliasing errors will result in correlated false detections. In the SIS method, there are also aliasing errors in the frames, but these errors will be uncorrelated over subsequent frames, and will therefore not lead to correlated detections. Because correlated false detections are assumed to be correct detections by the tracking algorithm, the Zomet 1 method will lead to more false detections. This result shows that it is useful to apply SR reconstruction with a zoom factor larger than 1, because this reduces the aliasing noise.

5.2 Performance Comparison

The area under the ROC curve is often used as a performance measure.^{25,26} In our case the interesting part of the ROC curve is where the number of false detections is small.²⁷ Therefore, the area under the ROC curve is determined up to a value of 20 false detections per second (for a 256×256 -pixel frame and 15 frames/s).

In this paper the performance of a specific scenario is determined by ANR_{80} , the ANR of the point target that corresponds with a 80% area under the ROC curve. This is the ANR for which 80% of the point targets are detected with 20 false detections per second. Note that a smaller ANR_{80} indicates better performance. An operating point of 80% area under the ROC curve up to 20 false detections per second seems useful: The fraction of true detections is high enough to perform more advanced tracking while reducing the number of false positives even further.

For each scenario, ROC curves are measured for varying ANRs. Under each ROC curve the area under the ROC curve up to 20 false detections per second is calculated. By linear interpolation between those areas the ANR_{80} is determined. The resulting ANR_{80} 's are shown in Table 1, including the precision. The precision of the ANR_{80} is deter-

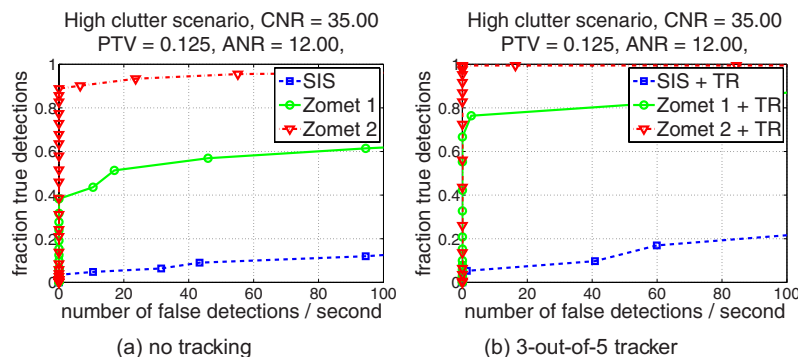


Fig. 8 ROC curves of the three different background suppression methods for a high-clutter scenario (buildings), a point target velocity of 0.125 pixel/frame, and an ANR of 12. Part (a) shows ROC curves obtained without using tracking, and (b) shows ROC curves after using a three-out-of-five tracker.

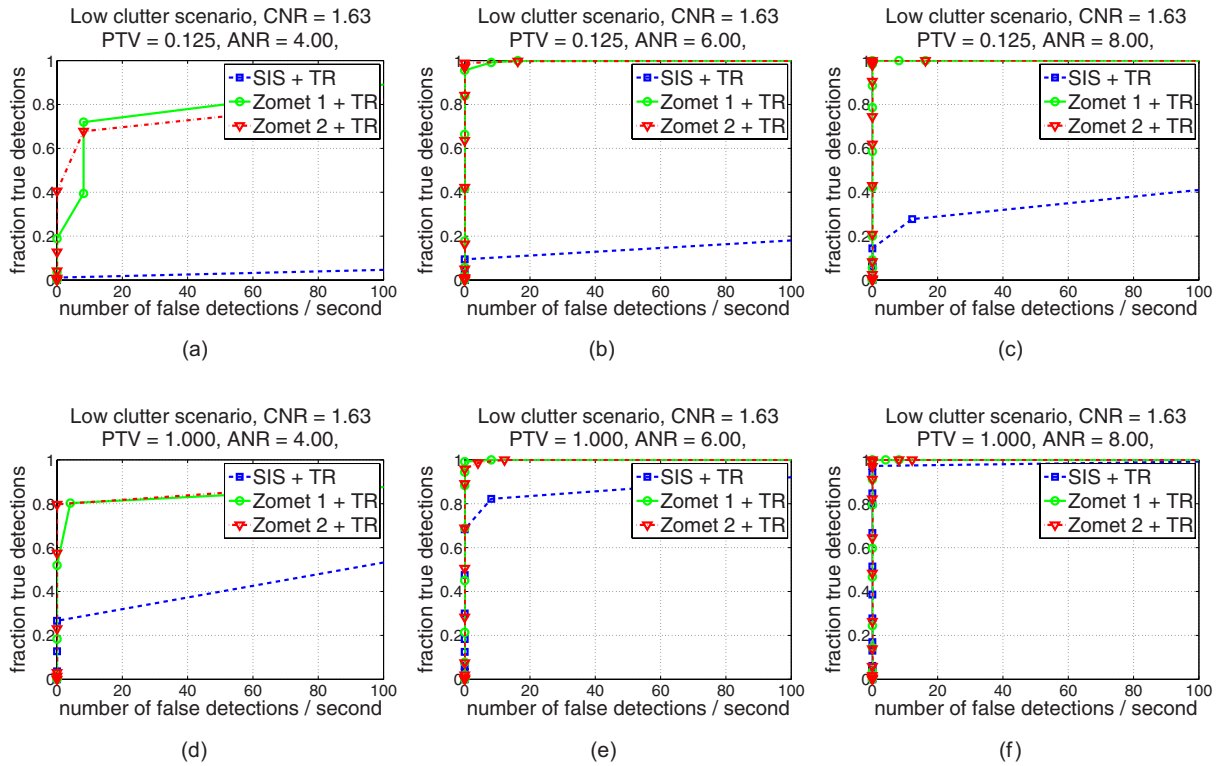


Fig. 9 ROC curves of the three different background suppression methods plus tracking for a low-clutter scenario (sky). SIS, shift, interpolate, and subtract; Zomet 1 and 2, Zomet's robust SR method with zoom factor 1 and 2, respectively; TR, tracking. The upper row [(a), (b), (c)] shows results of data with low apparent point target velocity (0.125 pixel/frame), and the lower row [(d), (e), (f)] with high apparent point target velocity (1 pixel/frame). Each column shows a specific point target amplitude-noise ratio.

mined from the standard deviation of each point (eight measurements) on the ROC curves. The inaccuracy of the interpolation is not taken into account in this precision. As can be seen in Table 1, the improvement of tracking is on the order of 1.5, except for point targets with the highest apparent target velocity (PTV) of 2 pixels/frame, where the effect of tracking is negligible. This might be explained by the limited association window used in the tracking algorithm.

The relative performance of the proposed detection method using Zomet's robust SR reconstruction for background suppression, compared with the detection method based on SIS, is presented in Table 2. Each number indicates the ratio of the ANR_{80} of SIS+TR (baseline) to the ANR_{80} of Zomet 1+TR or Zomet 2+TR. Larger numbers indicate better performance.

For the smallest apparent target velocity the improvement is at least 2.3 for the Zomet method compared to SIS. For high-clutter scenarios, Zomet 2 is significantly better than Zomet 1, while for low-clutter scenarios the difference between Zomet 1 and Zomet 2 is not significant. The largest improvement compared to the baseline is obtained with Zomet 2 for the scenario of high clutter and a small apparent point target velocity. Here, the improvement is almost a factor 4.

The measured improvement in performance of a detection method using SR reconstruction over the expected theoretical improvement is also given in Table 2. The theoret-

ical values are based on the analysis in Secs. 2.2 [Eq. (4)] and 2.3 (Fig. 2) and are given by $\text{TI} = I_{\text{noise}} \cdot I_{\text{ampl}}$. Here, I_{noise} is the ANR improvement due to the temporal noise reduction, which is $\sqrt{2}$, and I_{ampl} is the ANR improvement due to point target suppression in the background by using robust SR reconstruction.

Because we cannot quantitatively estimate the reduction of the aliasing noise, the values for the high-clutter scenario are only indicative. As can be seen in this table, mostly the real performance is smaller than our theoretical expectations. For point targets with a high apparent motion with respect to the background, the improvement is primarily due to temporal noise reduction. Here, our measurements are close to what is expected. For small apparent motion of the point targets, the difference between theory and measurements is somewhat larger. This may be explained by the fact that the simulated point targets are placed in real recorded data, which introduces an error in the position and therefore in the apparent motion of the point target. Relatively, this error is larger for smaller apparent motions.

6 Conclusions and Discussion

In this paper we have presented a new method for point target detection based on superresolution (SR) reconstruction of the background. With a simulation based on a real-world sequence we show that the specificity and sensitivity of a point target detection method are improved. The im-

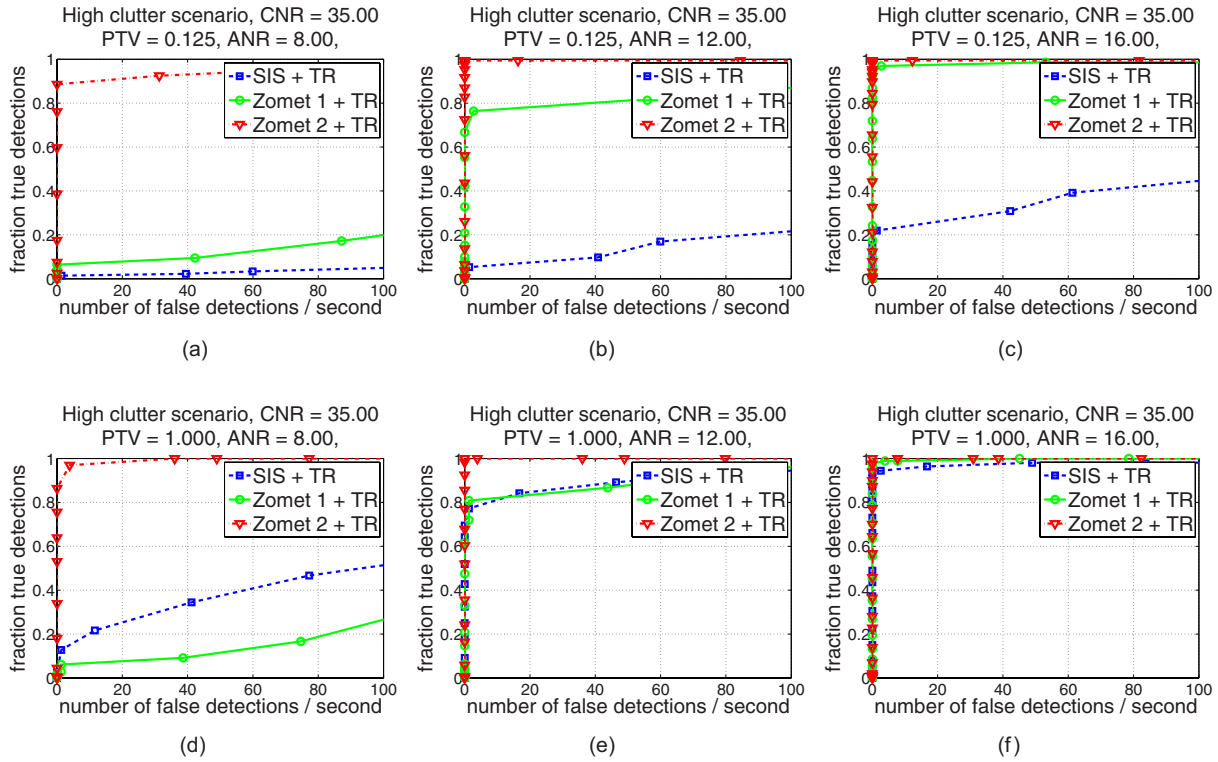


Fig. 10 ROC curves of the three different background suppression methods plus tracking for a high-clutter scenario (buildings). SIS, shift, interpolate, and subtract; Zomet 1 and 2, Zomet’s robust SR method with zoom factor 1 and 2, respectively; TR, tracking. The upper row [(a), (b), (c)] shows results of data with low apparent point target velocity (0.125 pixel/frame), and the lower row [(d), (e), (f)] with high apparent point target velocity (1 pixel/frame). Each column shows a specific point target amplitude-noise ratio.

Table 1 Point target amplitude noise ratios (ANR_{80}) that correspond with 80% area under the ROC curve up to 20 false detections per second for the shift, interpolate, and subtract (SIS) method and the Zomet method with zoom factors 1 and 2. For all methods the results with and without tracking (TR) are given. Note that a smaller ANR_{80} indicates better performance.

Clutter	PTV	ANR_{80}					
		SIS	SIS +TR	Zomet 1	Zomet 1 +TR	Zomet 2	Zomet 2 +TR
Low	0.125	23.3±3.9	15.6±3.7	8.0±0.7	5.3±1.3	9.4±0.6	5.1±0.8
	0.25	24.1±2.0	15.4±2.7	7.8±0.7	5.1±0.8	9.3±0.9	5.3±1.3
	0.5	15.5±3.9	12.0±2.9	7.7±0.9	4.2±1.1	8.7±1.2	5.4±1.2
	1	9.5±0.8	6.2±0.7	7.8±1.4	4.4±1.8	8.6±1.8	4.0±1.3
	2	7.4±0.6	7.2±1.2	6.0±0.3	5.6±0.7	6.7±0.7	5.7±0.9
High	0.125	57.5±12.9	29.6±4.7	16.5±1.0	12.8±1.9	10.6±0.7	7.7±0.2
	0.25	44.8±4.5	29.9±7.3	16.2±0.6	12.7±1.2	10.4±0.5	8.1±0.7
	0.5	31.8±3.6	21.7±2.7	16.0±0.6	11.9±1.6	10.0±0.6	7.6±0.2
	1	20.4±1.8	12.3±1.5	15.7±0.8	12.0±1.1	9.5±1.1	7.3±1.0
	2	15.0±0.9	11.8±1.9	15.7±0.5	14.8±1.2	7.9±1.1	5.8±1.1

Table 2 Relative performance of the proposed detection method using Zomet for background suppression and the detection method based on shift, interpolate, and subtract (SIS). For all methods only the results with tracking (TR) are given.

Clutter	PTV	ANP ₈₀ ratio		
		Zomet 1+TR versus SIS+TR	Zomet 2+TR versus SIS+TR	Theoretical improvement
Low	0.125	2.9±1.0	3.1±0.9	4.1
	0.25	3.0±0.7	2.9±0.8	2.8
	0.5	2.9±1.0	2.2±0.8	2.4
	1	1.4±0.6	1.6±0.8	1.7
	2	1.3±0.3	1.3±0.3	1.2
High	0.125	2.3±0.5	3.8±0.6	>4.1
	0.25	2.4±0.6	3.7±0.5	>2.8
	0.5	1.8±0.3	2.9±0.4	>2.4
	1	1.0±0.2	1.7±0.3	>1.7
	2	0.8±0.1	2.0±0.4	>1.2

improvement in specificity is based on two properties of the SR reconstruction algorithm: temporal noise reduction and anti-aliasing. Both these effects reduce the number of false alarms, because there is less noise in the background estimation and therefore also less noise in the difference image on which the detection is based.

The sensitivity of point target detection is increased by the point target suppression capabilities of SR reconstruction in the background estimate. Therefore, the amplitude of the point target is preserved in the difference image. This effect is larger for point targets with lower apparent target velocity. Robust SR reconstruction is used, because this suppresses outliers and therefore has hardly any contribution of the point target in its background estimation, whereas for nonrobust SR reconstruction methods a small portion of the point target energy will still be seen in the background estimation.

It can be seen that background suppression with SR reconstruction performs better than a standard shift, interpolate, and subtract (SIS) algorithm in almost all tested scenarios. As expected, SR reconstruction with zoom factor 2 performs better than SR reconstruction with zoom factor 1 in high-clutter scenarios. This effect is due to the fact that a better estimation of the background by using antialiasing, as is done with zoom factor 2, will decrease the number of false detections. In low-clutter scenarios a higher zoom factor does not improve the performance.

The improvement using SR reconstruction is only demonstrated for a limited data set. However, these results provide indicators for the performance of these techniques using other imaging systems and for other scenes. The performance depends on the properties of the imaging system, such as its sharpness and sampling frequency. These

properties affect the aliasing properties. According to the theory, the performance gain will be lower for systems with less aliasing. On the other hand, the performance gain will increase for systems with more aliasing. The results also depend on the amount of clutter in the scene. In our simulations we tested two scenarios: a high-clutter scenario, which was clutter-dominated, and a low-clutter scenario, which was temporal-noise-dominated. This provides two measuring points at the extremes of the clutter-to-noise ratio. The performance gain of scenes with another clutter-to-noise ratio will therefore be in between the low- and high-clutter improvements shown in this paper.

Summarizing, we have shown that point target detection after background suppression with SR reconstruction is significantly better than detection with the SIS method, especially in high-clutter scenarios and for low apparent target motion with respect to the background. While maintaining equal detection performance, the proposed method using SR reconstruction can detect point targets that have an up to 4 times smaller amplitude-to-noise ratio in the scenarios studied. In practice this implies that a point target can be detected at longer range.

Acknowledgments

Portions of this work are described in European patent application No. 06077053.4: Point Target Detection with Super-Resolution.

References

1. K. Schutte, D. J. J. de Lange, and S. P. van den Broek, "Signal conditioning algorithms for enhanced tactical sensor imagery," in *Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XIV*, Proc. SPIE 5076, 92–100 (2003).
2. T. Q. Pham, L. J. van Vliet, and K. Schutte, "Robust fusion of irregularly sampled data using adaptive normalized convolution," *EURASIP J. Appl. Signal Process.* 1010, 1–12 (2006).
3. Y. Xiong, J. X. Peng, M. Y. Ding, and D. H. Xue, "An extended track-before-detect algorithm for infrared target detection," *IEEE Trans. Aerosp. Electron. Syst.* 33(3), 1087–1092 (1997).
4. S. M. Tonissen and Y. Bar-Shalom, "Maximum likelihood track-before-detect with fluctuating target amplitude," *IEEE Trans. Aerosp. Electron. Syst.* 34(3), 796–809 (1998).
5. S. S. Blackman, *Multiple-Target Tracking with Radar Application*, p. 10, Artech House (1986).
6. S. C. Park, M. K. Park, and M. G. Kang, "Super-resolution image reconstruction: a technical overview," *IEEE Signal Process. Mag.* 20(3), 21–36 (2003).
7. S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar, "Advances and challenges in superresolution," *Int. J. Imaging Syst. Technol.* 14(2), 47–57 (2004).
8. A. W. M. van Eekeren, K. Schutte, O. R. Oudegeest, and L. J. van Vliet, "Performance evaluation of super-resolution reconstruction methods on real-world data," *EURASIP J. Adv. Signal Process.* 2007, 1–11 (2007).
9. T. Q. Pham, L. J. van Vliet, and K. Schutte, "Influence of signal-to-noise ratio and point spread function on limits of super-resolution," in *Image Processing: Algorithms and Systems IV*, Proc. SPIE 5672, 169–180 (2005).
10. A. van der Schaaf, "Natural image statistics and visual processing," Ph.D. Thesis, Rijksuniv. Groningen (1998).
11. B. Zitová and J. Flusser, "Image registration methods: a survey," *Image Vis. Comput.* 21(11), 977–1000 (2005).
12. T. Q. Pham, M. Bezuijen, L. J. van Vliet, K. Schutte, and C. L. Luengo Hendriks, "Performance of optimal registration estimators," in *Visual Information Processing XIV. Defense and Security*, Proc. SPIE 5817, 133–144 (2005).
13. B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proc. 1981 DARPA Imaging Understanding Workshop*, pp. 121–130 (1981).
14. I. T. Young, J. J. Gerbrands, and L. J. van Vliet, *Fundamentals of Image Processing*, Delft Univ. of Technology (1995).
15. R. Y. Tsai and T. S. Huang, "Multiframe image restoration and registration," in *Advances in Computer Vision and Image Processing*, Vol. 1, pp. 317–339, JAI Press, 1984.

16. E. Kaltenbacher and R. C. Hardie, "High resolution infrared image reconstruction using multiple, low resolution, aliased frames," in *Proc. IEEE National Aerospace and Electronics Conf.*, Vol. 2, pp. 702–709 (1996).
17. S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar, "Fast and robust multi-frame superresolution," *IEEE Trans. Image Process.* 13(10), 1327–1344 (2004).
18. R. C. Hardie, K. J. Barnard, J. G. Bognar, E. Armstrong, and E. A. Watson, "High-resolution image reconstruction from a sequence of rotated and translated frames and its application to an infrared imaging system," *Opt. Eng.* 37(1), 247–260 (1998).
19. A. Zomet, A. Rav-Acha, and S. Peleg, "Robust super-resolution," in *IEEE Conf. on Computer Vision and Pattern Recognition*, Vol. 1, pp. 645–650 (2001).
20. M. Irani and S. Peleg, "Improving resolution by image registration," *Graph. Models Image Process.* 53, 231–239 (1991).
21. http://en.wikipedia.org/wiki/FIM-92_Stinger.
22. http://en.wikipedia.org/wiki/Vympel_R-27.
23. G. B. Rybicki and A. P. Lightman, *Radiative Processes in Astrophysics*, John Wiley & Sons, New York, (1979).
24. R. D. Hudson, *Infrared System Engineering*, pp. 100–103, John Wiley & Sons (1969).
25. A. P. Bradley, "The use of the area under the ROC curve in the evaluation of machine learning algorithms," *Pattern Recogn.* 30(7), 1145–1159 (1997).
26. S. A. Rizvi, N. M. Nasrabadi, and S. Z. Der, "A clutter rejection technique for FLIR imagery using region-based principal component analysis," in *Proc. IEEE Int. Conf. on Image Processing*, Vol. 4, pp. 415–419 (1999).
27. A. P. Tzannes and D. H. Brooks, "Detection of point targets in image sequences by hypothesis testing: a temporal test first approach," in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, Vol. 6, pp. 3377–3380 (1999).



Judith Dijk received her MSc in 1998 from the Pattern Recognition Group (currently the Quantitative Imaging Group) at Delft University of Technology. In January 1998 she started a PhD in an IOP Beeldbewerking project called "Image quality measures for the perception of printing." This project was a collaboration between the Applied Physics Department in Delft and TNO Human Factors in Soesterberg. In 2004 Judith received her PhD with the thesis "In search

of an objective measure for the perceptual quality of printed images." In 2003 she started working as a research scientist in the Electro-Optics Group within TNO. Her research interests include color image processing, superresolution, and image quality assessment.



Adam W. M. van Eekeren received his MSc in 2002 from the Department of Electrical Engineering at the Eindhoven University of Technology. He did his graduation project within Philips Medical Systems on the topic of image enhancement. Subsequently he worked one year at the Philips Research Laboratory on image segmentation using level sets. In 2004 he started his PhD project entitled "Super-resolution on small moving objects" in the Electro-Optics

Group at TNO in collaboration with the Quantitative Imaging Group

at Delft University of Technology. His research interests include image restoration, superresolution, image quality assessment, and object detection.



Klammer Schutte (MSc, physics, 1989, University of Amsterdam) received his PhD in 1994 from the University of Twente for his thesis "Knowledge based recognition of man-made objects." Subsequently he had a postdoctoral position with the Delft University of Technology's Pattern Recognition (now Quantitative Imaging) Group. Since 1996 he has been employed by TNO, currently as a senior research scientist (electro-optics) within the Business Unit Observation System. Within TNO he has actively led multiple projects in areas of signal and image processing. Recently many of his projects have included superresolution reconstruction for both international industries and governments, resulting in superresolution-reconstruction-based products in active service. His research interests include pattern recognition, sensor fusion, image analysis, and image restoration. He is secretary of the NVBHPV, the Netherlands branch of the IAPR.



Dirk-Jan de Lange received his PhD in physics from the University of Utrecht in the Netherlands in 1998 for his work on nuclear physics. Since 1998 he is working at TNO Defence, Security and Safety in The Hague as a research scientist. His main interest is on real-time image processing, image enhancement, super-resolution, noise reduction, automatic target detection, turbulence compensation and sensor performance prediction.



Lucas J. van Vliet is a full professor in multidimensional image processing and analysis at the TU Delft. He studied applied physics at the TU Delft and received his PhD degree cum laude in 1993. His thesis entitled "Grey-scale measurements in multidimensional digitized images" presents novel methods for sampling-error-free measurements of geometric object features. He has worked on various sensor, image restoration, and image measurement problems in

quantitative microscopy and medical imaging. In 1996 he was awarded a fellowship of the Royal Netherlands Academy of Arts and Sciences (KNAW). He was a visiting scientist at LLNL (1987), UCSF (1988), Amoco ATC (1989–1990), Monash University (1996), and LBNL (1996).