Laboratory Experimentation for Dim Signal Detection in Cluttered Optical Data

Christopher T. Agh*, Matthew Buoni, Andrew P. Brown
Toyon Research Corporation, 6800 Cortona Drive, Goleta, CA, 93117

ABSTRACT

We discuss an algorithmic approach for detecting spatially stationary, dim signals in cluttered optical data. In the problem considered here, cluttered scene backgrounds are substantially more intense than sensor noise and signal variations from scene anomalies of interest. As a result, clutter estimation and rejection algorithms are performed prior to implementing signal detection schemes. Even then, stationary residual clutter may be spatially similar to, and have intensities much greater than, those of the signals of interest. This poses an extreme challenge for the automated detection of low-contrast scene anomalies, and detectors based solely on spatial properties of the optical scene generally fail. In our newly developed signal detection algorithm, we exploit not only the structure of the dim signals of interest, but also the time-lapsed residual clutter. By examining the properties and statistics of both the signals of interest and the signals we wish to reject, Toyon has developed an algorithm for the automated detection of low-contrast signals in the presence of high-intensity clutter. We discuss here the developed signal detection algorithm and results for overcoming the challenges inherent to heavily cluttered optical data.

Keywords: automated dim signal detection, signal classification, optical data, clutter suppression, residual clutter, sensor motion

1. INTRODUCTION

Many algorithms exist for detecting and tracking a variety of signals in cluttered optical data. Removal of the clutter is typically the first stage of such algorithms. Clutter modeling can be performed spatially and/or temporally[1], and involves building a clutter model based on previous optical data and subtracting the model from current data. Any inaccuracies in the clutter model, due to model/scene mismatch or errors in sensor motion jitter estimation, will lead to clutter leakage. In this situation, sharp gradients in image scene backgrounds tend toward more pronounced, high-intensity residual clutter. Depending on how a clutter model is built and applied, clutter leakage is typically stationary in pixel coordinates, but can vary in time, in both magnitude and sign, within a clutter-suppressed sequence of optical data. These clutter-based intensity changes may cause little or no issues for state estimation algorithms employed to detect and track moving objects, such as Kalman filters[2] or particle filters[3], since one can place motion requirements on the dynamics of the objects of interest. However, due to its spatially static and high intensity nature, clutter leakage is a challenging component for the automated detection of dim, stationary signals. Figure 1 shows two clutter-suppressed image sub-chips from the same processed image. These figures aim to show the spatial similarity of, and intensity differences between, example clutter leakage in Figure 1(a) and a simulated, low-contrast signal of interest in Figure 1(b). Both features occupy an approximate 3x3 to 5x5-pixel region; however the clutter leakage is clearly brighter.

*cagh@toyon.com; phone: 805-968-6787; fax: 805-685-8089; toyon.com

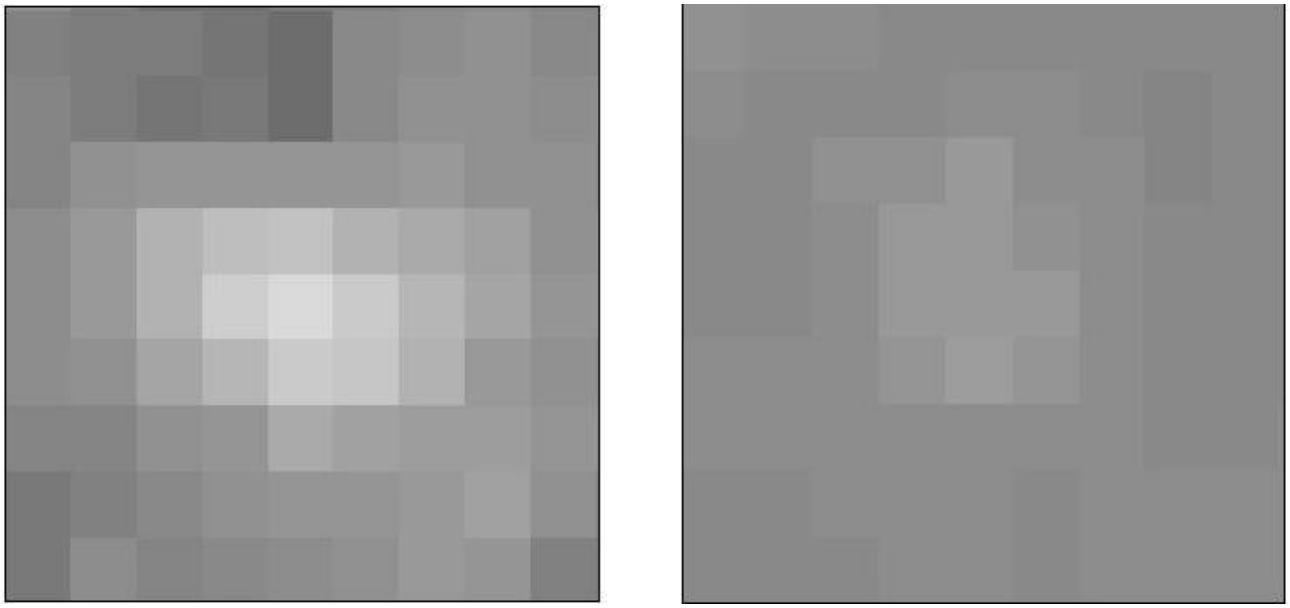

(a) (b)

Figure 1: Clutter-suppressed image sub-chips. (a) Example of clutter leakage. (b) Signal of interest.

The algorithm discussed in this paper jointly exploits the dim signals of interest and the signals produced by residual clutter. The amplitudes of the residual clutter may be up to an order of magnitude greater than the amplitudes of the dim signals, but our algorithm takes advantage of the statistical differences in the two signal types in order to track low-contrast anomalies, while simultaneously rejecting regions of high pixel intensities due to background clutter leakage. The algorithm may be implemented as a particle filter in which particles remain stationary and only track time-lapsed signal variations, with particle likelihoods calculated and used in the resampling stage, or it may be implemented as a detector in which each pixel in the image plane is tracked over time. The choice here is a matter of sufficient resources for adequate run-time. The specifics of the algorithm will be discussed in detail while keeping the implementation generic.

2. DATA

A simulation tool, developed by Toyon Research Corp., was used to generate optical data in order to test the stationary signal detector in a lab environment. The generated image sequences contain background clutter, and the simulator was configured to insert optical anomalies, as well as sensor artifacts, sensor noise, and sensor motion jitter. Each dim signal of interest was modeled as a point source and inserted into its corresponding image sequence at a random, sub-pixel location with properly applied optical point-spread-function (psf) smoothing. A psf standard deviation of 1.0 was used during the simulation. The time-domain signal of interest, $s(t)$, prior to insertion into the image frames, was comprised of an A/C fluctuating amplitude with a positive mean value. The mean value can be thought of as the DC component of the fluctuating signal. The determination of the signal mean, or DC-offset, was determined by the desired signal-to-clutter ratio (SCR). Here, SCR is defined to be

$$\mathrm{SCR} = \overline{s(t)} \Big/ \overline{x(t)}, \tag{1}$$

where again, $s(t)$ is the time-domain signal prior to insertion in the image frames, $x(t)$ represents the background intensities at the signal location prior to insertion, and the "overbar" in both the numerator and denominator denotes the mean value of the respective signals $s(t)$ and $x(t)$. The simulated dim signals, for which detection results are shown in Section 4, were generated with SCR = 0.001. The signal-to-noise ratio (SNR) is defined to be the root-mean-square (RMS) of $s(t)$ divided by the standard deviation of the noise. Noise was drawn randomly from a Gaussian distribution so that the signals of interest in the simulated optical data have an SNR = 8. Sensor motion jitter was also generated from a random Gaussian distribution with a frame motion standard deviation of 0.25 pixels.

A dim signal example prior to being inserted into the image frames is shown in Figure 2. The amplitude of this signal at each frame is then added to the simulated frame at the pre-determined signal location, with spreading of the signal due to the sensor psf. The pristine signal in Figure 2 has not yet been corrupted by background clutter or sensor noise, and the extremely low positive mean value (DC offset) of approximately 0.1 is clearly seen within the small amplitude range along the y-axis. To gain a better understanding of the difficulty of detecting these signals in cluttered optical data, Figure 3 shows the same dim signal (blue) after insertion into cluttered background images with additive sensor noise and after the clutter suppression process has been applied. An example of residual clutter observed in the same clutter-suppressed image sequence is overlaid (green) to show how the dim signal of interest is buried in undesirable clutter.

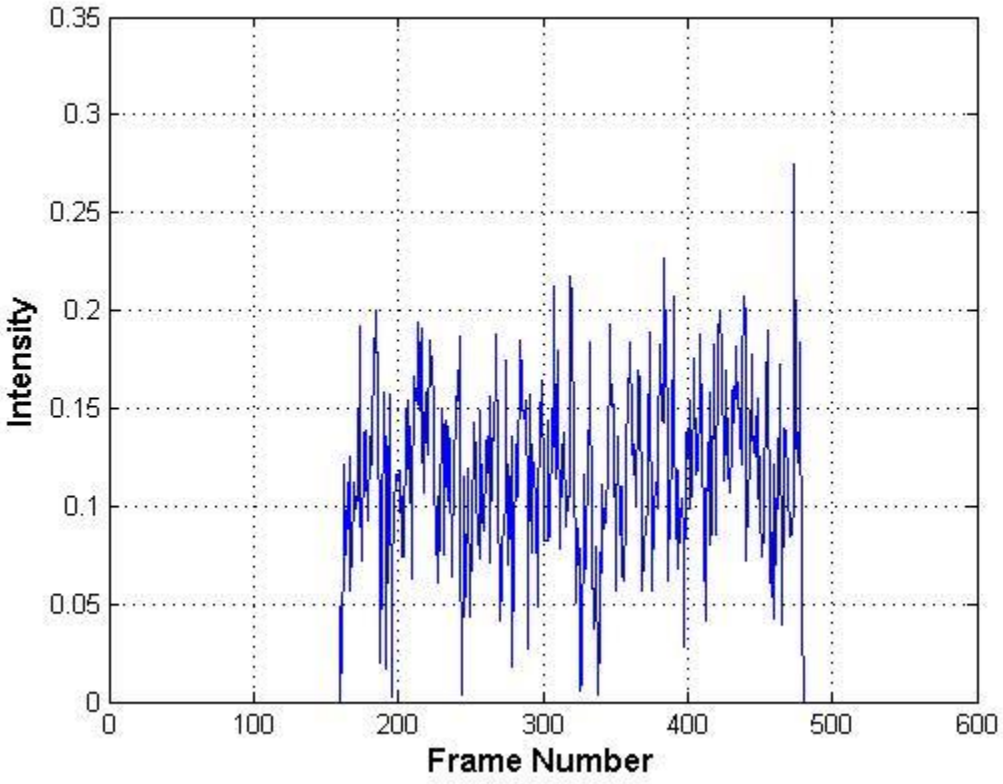

Figure 2: Uncorrupted dim signal prior to insertion into the simulated image frames. This signal has an SCR of 0.001.

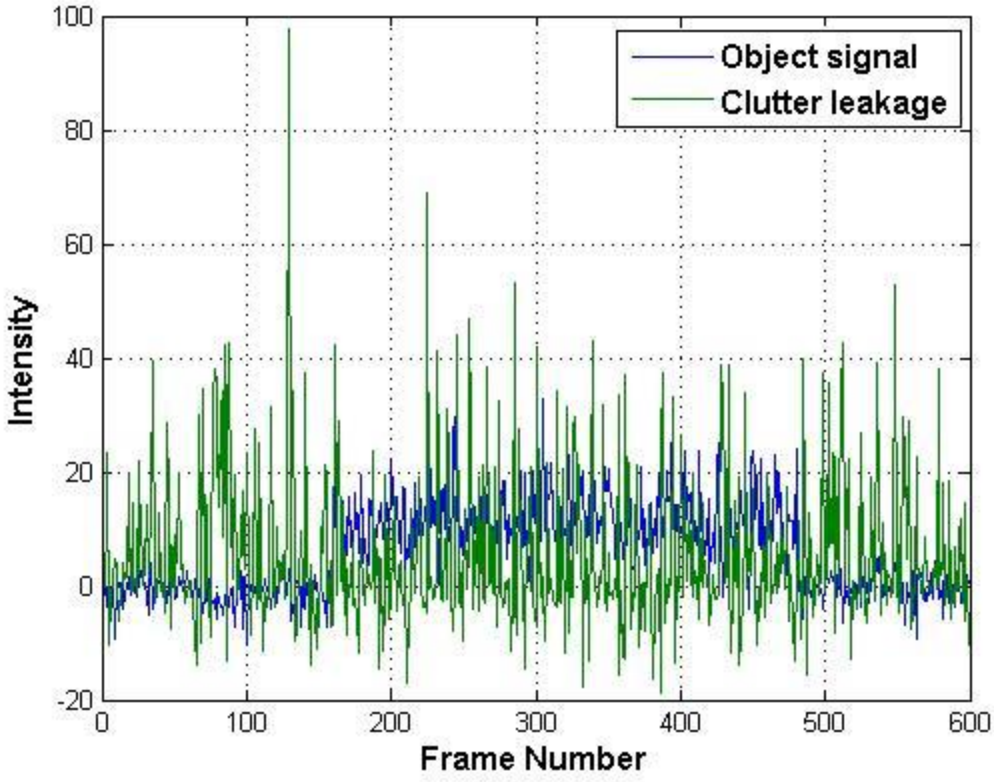

Figure 3: Dim signal of interest (blue) from Figure 2 along with clutter leakage (green) observed in simulated clutter-suppressed optical data.

3. DIM SIGNAL DETECTION

Spatial and temporal clutter estimation and rejection algorithms process raw sensor data and result in images that contain only uncorrelated noise, clutter leakage, and signals of interest. High-intensity clutter leakage can often cause false detects, especially in the case of stationary signals. Clutter leakage is due to slight misalignment during the clutter estimation and rejection process and is more prevalent with greater frame-to-frame jitter due to sensor motion. This causes dim, optical anomalies (especially those stationary in nature) to be lost amongst bright intensity areas where background clutter leaks through. However, the time-domain structure of the signals of interest is in general different than the structure of clutter leakage. Our algorithmic approach uses statistical methods and timing thresholds to make use of these differences as clutter-suppressed images are integrated over time. We continue this section with a description of the algorithm and a comparison of the low-contrast signals of interest versus high-intensity clutter leakage.

Signal detector

Initially, whether implemented in a particle-filter framework or not, signal observation/signal filter locations are uniformly distributed throughout the image plane. This could easily be implemented in a per-pixel fashion. Regardless of the distribution, at each frame each observed signal (interpolated pixel intensities if using sub-pixel locations) is stored in a buffer of length L, and an online-recursive mean and standard deviation are computed for each length-L signal history. The mean value of each current signal is then compared to a delayed signal mean in order to find locations of potential rises from the current background level, i.e. "bursts".

Let n be the current frame number and $\{x_i\}_{i=n-L+1}^{n}$ be the most current L image intensities observed at a particular observation location (i.e. at its location in the image plane over the past L frames). We then compute the mean and standard deviation of these observations.

$$y_n = \frac{1}{L}\left(\sum_{i=n-L+1}^{n} x_i \right) \tag{2}$$

$$\sigma_n^y = \sqrt{\frac{\sum_{i=n-L+1}^{n-D} \left(y_i - \bar{y} \right)}{n}} \tag{3}$$

for some delay $0 < D < L$, and where $\bar{y} = \text{mean}\left(\{ y_n \} \right)$. Each mean value y_n from equation (2) is over the set of all past L observations, and the new observed signal $\{y_n\}$ is a low-pass filtered version of the raw signal $\{x_n\}$, so much of the high frequency variation is removed. The standard deviation σ_n^y from equation (3) is only over the first $L - D$ of the past L observations. This is to ensure that the beginning of a "burst" in the most current observations do not increase the standard deviation too rapidly.

Next, we compute

$$z_n = \frac{y_n - y_{n-L}}{\sigma_n^y} \tag{4}$$

This sequence of "z-factors" may aid in the likelihood computation of each particle if implementing a particle filter and/or thresholding the observations of potential stationary bursts. A *preliminary detection* is one in which the observed z-factor is greater than some threshold θ_{on}. In this case, before iterating to the next image frame, the online recursive delayed mean and standard deviation in (4) is stopped for that particular signal observation and we set $y_{on} = y_{n-L}$ and $\sigma_{on}^y = \sigma_n^y$. The current observed signal mean comparisons are now performed with the current mean value and the delayed mean value and standard deviation just prior to the "burst", i.e. the z-factor is computed as

$$z_n = \frac{y_n - y_{on}}{\sigma_{on}^y} \tag{5}$$

An observed signal with z-factor greater than θ_{on} necessarily has a mean value higher than previously observed. If the z-factor continues to remain larger than θ_{on} for some preset number of frames, then *detection* is declared at that signal location and a "detect flag" is turned on for that observation. The signal mean comparisons, i.e. z-factors, continue to be performed at all signal locations with their "detect flag" in the "on" position, but the algorithm now checks for the z-factor to drop *below* some threshold θ_{off}. If this occurs, and continues to occur, for some preset number of contiguous frames, then the "detect flag" for that signal is turned off and we resume updating the online-recursive delayed mean and standard deviation for that observed signal. In this case, the mean of the signal is necessarily close to its original value before the preliminary detection occurred.

Signal comparisons

By tuning the delay parameter *L*, exclusion parameter *D*, z-factor thresholds θ_{on} and θ_{off}, as well as the frame counting thresholds that keep track of how long $z_n > \theta_{on}$ or $z_n < \theta_{off}$, we are able to detect and track signals whose amplitudes are much dimmer than the surrounding residual clutter. Also, the algorithm does not require *a priori* knowledge of the duration of the signals of interest. The figures below show the simulated signals (A/C signal with small D/C component) with SCR = 0.001, and different examples of residual clutter.

We show the raw observed signal $\{x_n\}$ from the sequence of clutter-suppressed images, and also the low-pass filtered sequence $\{y_n\}$ which is used by the algorithm to perform the comparisons and thresholding. The time-domain series are displayed in order to better understand the differences in signal structure between the signals of interest and the residual clutter, and also to help visualize how the record-keeping aspect of the algorithm might behave on such signals. Figure 4(a) shows the raw observed signal from a clutter-suppressed image sequence. Figure 4(b) shows the filtered version of

the same signal which has been greatly smoothed to reduce high-frequency content. In both figures we visually see the "burst", and that it remains elevated for approximately 300 frames. Similarly, Figure 5 shows the raw and filtered versions, respectively, of observed clutter leakage from the same clutter-suppressed image sequence. Amplitudes in the residual clutter reach over an order of magnitude higher than those of the signal of interest in Figure 4, however due to the statistics of the filtered signal it does not produce a false alarm in the algorithm described above. Figure 6 shows two more examples of filtered clutter leakage from different locations in the clutter-suppressed image plane overlaid with the filtered signal of interest from Figure 4(b). Again, although filtered clutter leakage may have much higher amplitudes than the signals of interest, the inherent time-domain structures of both the residual clutter and desired dim signal are utilized to distinguish between the two.

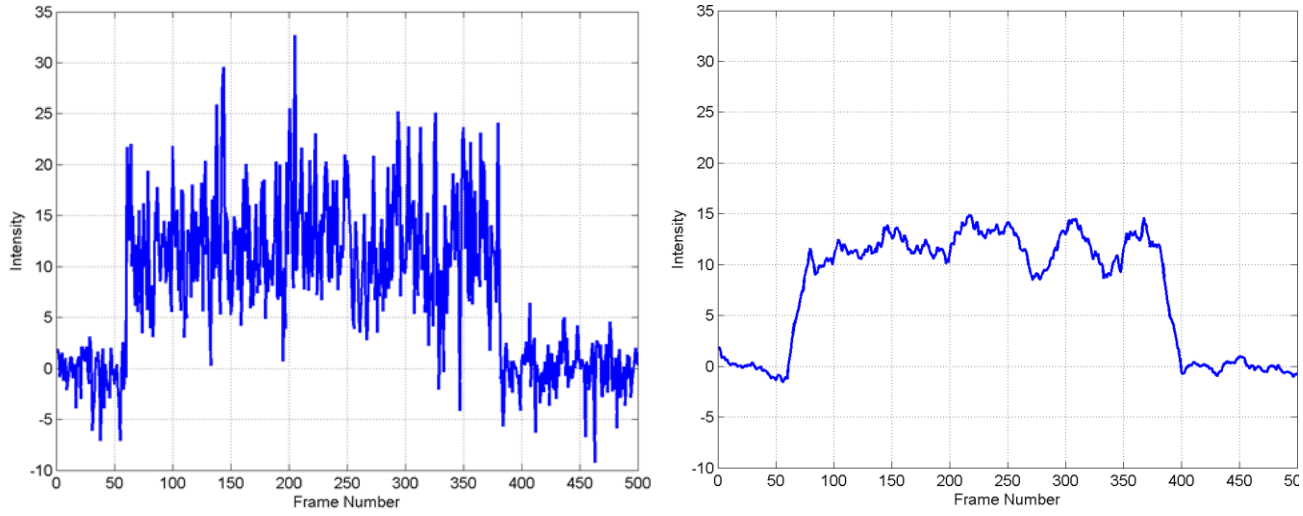

(a) (b)

Figure 4: Simulated, dim, signal observed in a clutter-suppressed image sequence. The raw signal is seen in (a), and the low-pass filtered signal (sequence of online-recursive means) is seen in (b).

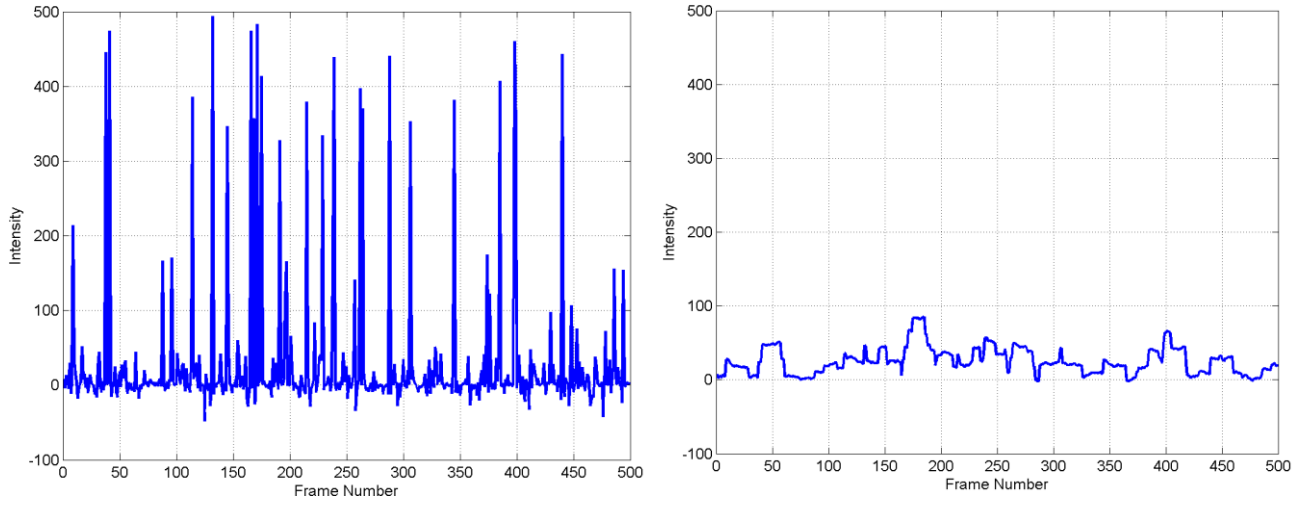

(a) (b)

Figure 5: Example of high-intensity clutter leakage observed in a clutter-suppressed image sequence. The raw signal is seen in (a), and the low-pass filtered signal (sequence of online-recursive means) is seen in (b).

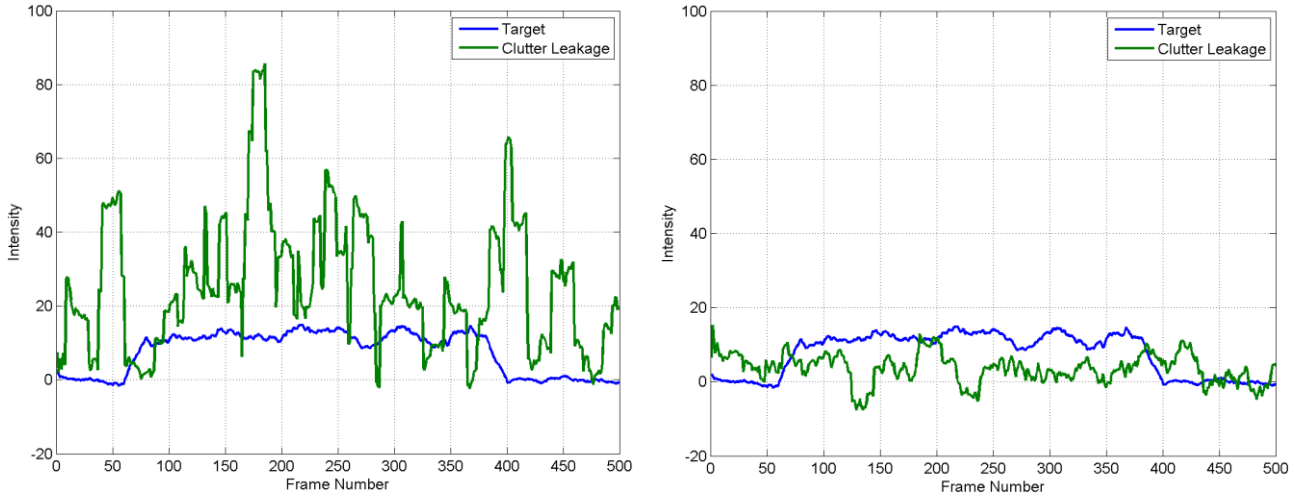

(a) (b)

Figure 6: Low-pass filtered versions of the simulated dim signal of interest (blue), and two other examples of clutter leakage (green).

4. RESULTS

Figure 7 shows two images from simulated optical data post clutter suppression and low-pass filtering. The original simulated data was generated with random sensor motion jitter having a standard deviation of 0.25 pixels in both the x- and y-directions. The image sequence also contains a low-contrast signal of interest as described above with SNR = 8 and SCR = 0.001. These images, post low-pass filtering, look similar to the corresponding images prior to low-pass filtering, but when viewed in succession at a quick rate the reduction in pixel intensity variation is noticeable. Figure 7(a) shows a filtered image without the presence of the desired signal, while Figure 7(b) contains the dim signal we are attempting to detect (located in the red box as a group of pixels with slightly higher intensity than the background). Both images contain a fair amount of high-intensity residual clutter which varies in time throughout the data, and is much brighter than the dim signal of interest.

Figure 8 shows two output frames from the detection algorithm. These can be constructed as particle weight images (2D probability mass functions for desired signal location), or as detection frames in which each pixel is proportional to its corresponding z-factor. Figure 8(a) is a frame when no target is present. The high-intensity regions due to residual clutter come in and out of view for various durations as each frame in the sequence is processed. Figure 8(b) contains the brightest region at the location of the simulated dim stationary signal, while bright spots due to clutter leakage are minimal in this image. During testing of the algorithm using various image sequences, high-intensity clutter leakage appears intermittently in the detection frames, however when the dim signal of interest is present the high-intensity regions converge to the correct simulated signal location and clutter leakage visibly falls to the background.

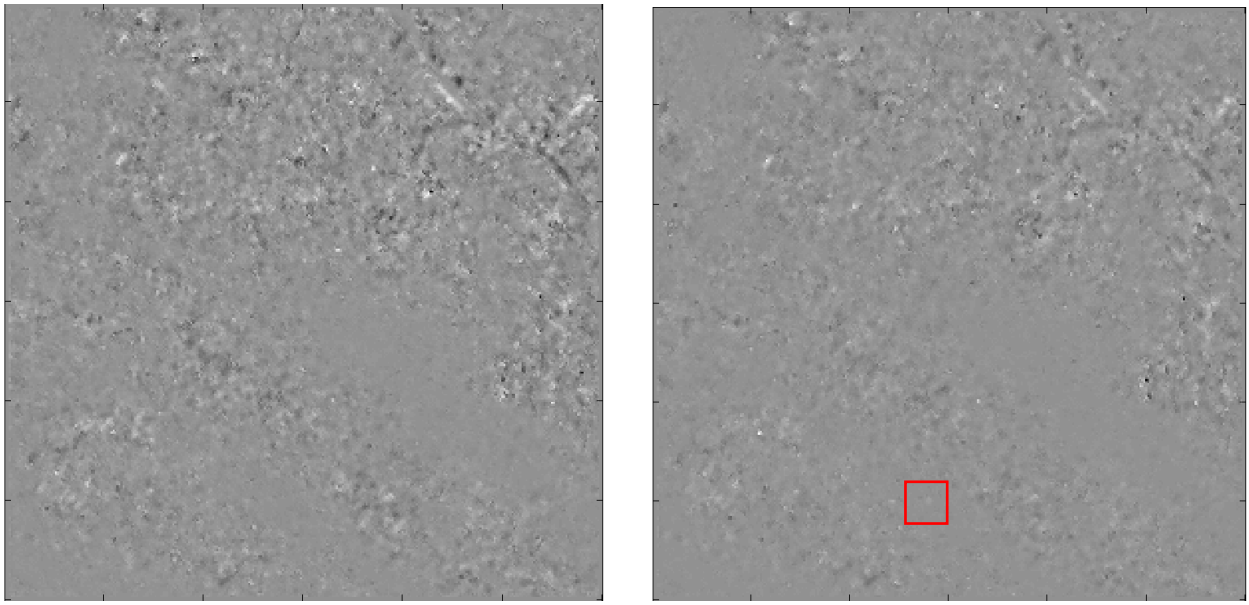

(a) (b)

Figure 7: Lowpass-filtered, clutter-suppressed image sequence frames. (a) No stationary signal of interest present. (b) Simulated signal of interest appears in the red box as a group of pixels slightly brighter than the background.

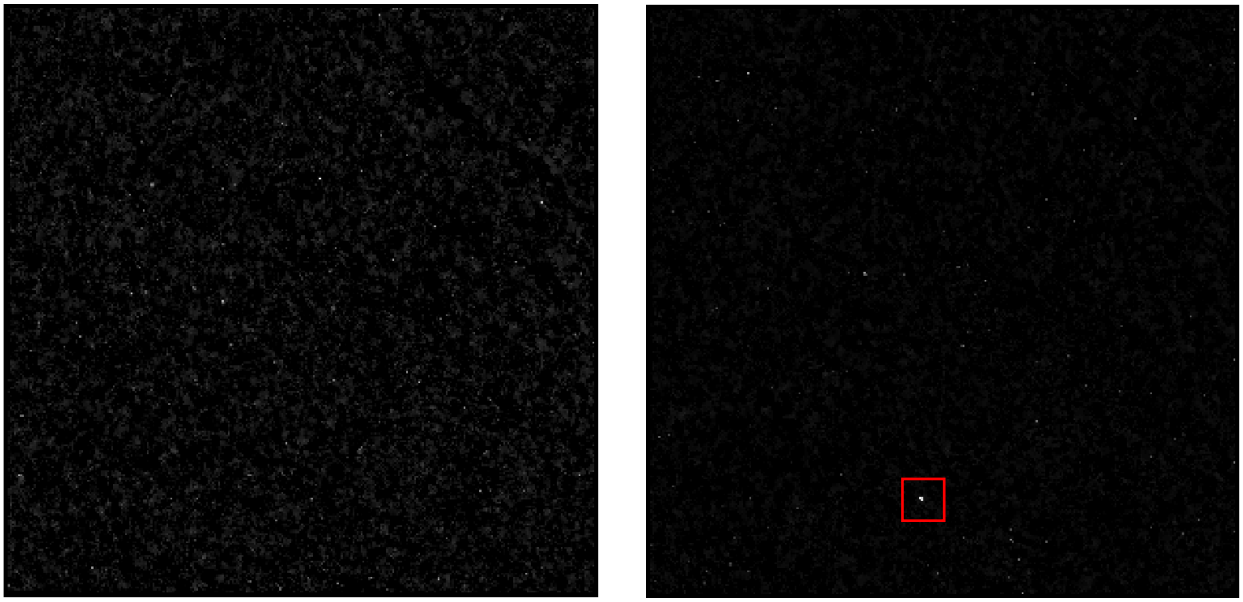

(a) (b)

Figure 8: Detection frame outputs from the stationary dim signal detection algorithm. (a) No signal presence. (b) Detection occurs at the location of the simulated dim signal of interest.

5. CONCLUSIONS

In order to detect dim signals in cluttered optical data it is imperative to model and reject clutter accurately. However, this is a difficult task and clutter leakage will occur with leakage amplitudes being a function of the sharpness of clutter gradients, as well as the amount of sensor motion while collecting data. These artifacts of clutter suppression are what cause the majority of false detects when searching for dim, stationary signals. In this paper, we discussed an algorithmic

approach for differentiating between dim signals of interest and signals typical of residual clutter. Through a study of the statistics of this leakage in simulated optical data, and laboratory experimentation, we have shown that our algorithm can detect extremely dim signals with SCRs as low as 0.001, while rejecting detections on bright spatial regions of clutter leakage that may have intensities up to an order of magnitude higher than the dim signals of interest. Natural extensions of this work include clutter modeling and rejection aimed at further reducing clutter leakage, and the detection of a variety of dim signal types in cluttered optical data.

ACKNOWLEDGEMENTS

The authors would like to thank Wellesley Pereira and Reed Weber for their support of this work through several technical discussions and helpful suggestions. We would also like to thank the Small Business Innovation Research program which provided partial funding for this effort.

REFERENCES

[1] Tartakovsky, A., Brown A., and Brown J., "Enhanced Algorithms for EO/IR Electronic Stabilization, Clutter Suppression, and Track-Before-Detect for Multiple Low Observable Targets," AMOS Conference, (2009).

[2] Bar-Shalom, Y., Li, X. R., and Kirubarajan, T., [Estimation with Applications to Tracking and Navigation], McGraw-Hill, New York, NY, (2001).

[3] Arulampalam, S., Maskell, S., Gordon, N., and Clapp, T., "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," IEEE Transactions on Signal Processing SP-50, 174–188, (Feb. 2002).