AUTOMATED TRACK EXTRACTION FOR COHERENT CHANGE DETECTION IMAGES.

CM Stevenson Dstl Porton Down, Salisbury, Wiltshire. SP4 0JQ. UK
M Nottingham Dstl Porton Down, Salisbury, Wiltshire. SP4 0JQ. UK
DG Muff Dstl Porton Down, Salisbury, Wiltshire. SP4 0JQ. UK
D Blacknell Dstl Porton Down, Salisbury, Wiltshire. SP4 0JQ. UK

1 INTRODUCTION

Automatic track extraction is a field that has been extensively studied, particularly for cartography. If manually performed, track extraction is a time consuming operation and a variety of approaches to automate the process have been developed[1,2,3].

Synthetic Aperture Radar (SAR) Coherent Change Detection (CCD) assesses the similarity of scene reflectivity of two SAR images by computing the magnitude of their complex cross correlation, or coherence[4]. Small scene changes which substantially modify the SAR magnitude and or phase, such as ground imprints from vehicle tracks, will result in low coherence in the CCD image. The sample coherence is typically assessed over a sliding N-pixel window and has been shown to comprise five dominant contributors[5,6] as in equation 1.

$$\rho = \rho_{SNR}\rho_{base}\rho_{proc}\rho_{vol}\rho_{scene} \qquad (1)$$

ρ_{SNR} is determined by the relative backscatter signal to radar receiver noise ratio and ρ_{base} is the spatial baseline decorrelation attributed to differences in acquisition geometry for the image pair. ρ_{proc} is the decorrelation due to differences in image formation techniques and errors in registration between reference and match images. For the purposes of this work it is assumed that the data acquisition and image formation techniques are optimized for CCD production and the first three terms of equation 1 are therefore assumed to be ≈1. The sample coherence is taken to be dependent upon ρ_{vol}, the decorrelation due to volume scattering in vegetated areas and ρ_{scene}, decorrelation associated with in-scene disturbances.

Automatic track extraction techniques for CCD imagery typically struggle with false alarms due to low coherence in vegetated areas and radar shadow regions[7]. CCD images also suffer from a complex granular noise pattern and tracks often present large variations in along-track amplitude. This is problematic for feature extraction, as subtle brightness transitions may be masked out and brightness connectivity within common scattering features lost. As a result, tracks in CCD images do not typically possess the same predictable features as roads in urban areas in electro-optical imagery e.g. smooth straight edges, constant width, low variations in grayscale amplitude, consistent junction patterns, low degree of curvature etc. As such, many of the techniques which have seen moderate success in urban road extraction do not cope well with the ill-constrained nature of the off-road tracks in CCD imagery.

In recent years, a number of techniques have been developed to optimize CCD track detection whilst minimizing false alarms. These include; the application of steerable filters to improve track fidelity[8,9], selection of suitable track seed points for track tracing, constraining line detections to those which form part of a parallel pair as one might expect from vehicle tracks[9] or using multiple passes to alleviate false change in regions consistently showing low coherence in multiple CCD images[10].

In this work we implement an analysis of the Hessian matrix to derive second order structural information[11] from both the CCD image and the underlying reference and match images to mitigate false alarms. Analysis of the Hessian matrix avoids the need to apply banks of steerable filters tuned for specific track orientations. A pixelwise trackness index is produced, in which each pixel in the CCD image is assigned a value [0,1] denoting how track-like the pixel is based on characteristic track features and local structure of the image. A value of 1 indicates that the pixel is part of a structure that is very track-like and 0 not at all track-like. This index is computed without the need for user intervention such as seed choice, and no constraints are placed on the type of track to be detected. The trackness filter is integrated across multiple scales to allow detection of single tracks of a range of widths in a single CCD image pair.

2 METHODS

Automated track extraction techniques are applied to CCDs created using the Gotcha public release dataset[12] based on the sample coherence evaluated over a 5 x 5 pixel window. The track extraction algorithm comprises four main steps:

1. Shadow mitigation
2. Normalised ridge saliency computation
3. Direction filter computation
4. Pixel-wise trackness index assignment and vectorisation

Details of the above stages are given in sections 2.1 – 2.3 below. In order to assess the performance of the algorithm, a ground truth set of data was created by manual selection of the features identified as "true" track-like changes. These are used to create pixel-wise Receiver Operating Characteristic (ROC) curves.

2.1 Shadow mitigation

As described in section 1, a large number of the false changes in a CCD image occur in radar shadow regions common to both the reference and match images. This is mitigated by assessing the first order statistics of the reference and match image pair and setting the coherence of the areas that correspond to low-radar cross section areas in both images to unity in the CCD image[7].

As in work by Stojanovic[7], for each pixel *i* in the CCD image, an m x m window centred on the corresponding i^{th} pixel in the reference and match images is extracted. The average power is computed for the coherent sum and coherent difference of the reference and match windows. The pixel is defined as being a low-radar cross section area if the average power is below a selected threshold, *T*, as defined in equation 2.

$$\frac{1}{2}\left(\frac{1}{M}\sum_{i=1}^{M}\left|x_i + x_{i+M}\right|^2 + \frac{1}{M}\sum_{i=1}^{M}\left|x_i - x_{i+M}\right|^2\right) < T \qquad (2)$$

M is the number of pixels (m x m) in the local neighbourhood of the i-th pixel, $x_1, \ldots, x_M$ are pixels of the reference image and $x_{M+1}, \ldots, x_{2M}$ are the corresponding test pixels in the i-th pixel neighbourhood of the match image.

The performance of the track extraction tool was assessed over a range of T values (10-50) and was found to be insensitive to variations in T within this range, as such, throughout the paper a value of T=30 was used for demonstration purposes. The CCD image with shadow regions masked to unity is termed the *nsCCD* image. Prior to ridge detection, to reduce the effects of the granular noise pattern whilst maintaining edge structure, a 9 x 9 median filter is applied to the *nsCCD* image.

2.2 Normalised ridge saliency (nRS) computation

A modified ridge detector is applied to SAR CCD imagery and optimized to enhance dark linear track-like features.

The track enhancement filter adopts a multiscale derivative of Gaussians approach and eigenvalue analysis of the Hessian matrix to derive second order structural information about an image $\boldsymbol{L}$. Differential operators of $\boldsymbol{L}$ at point x and scale s are defined as a convolution with derivatives of Gaussians as shown in equation 3, with the D-dimensional Gaussian defined as in equation 4.

$$\frac{\partial}{\partial x}L(\mathbf{x},s)=s^{\gamma}L(\mathbf{x})*\frac{\partial}{\partial x}G(\mathbf{x},s) \qquad (3)$$

$$G(\mathbf{x},s)=\frac{1}{\sqrt{\left(2\pi s^{2}\right)^{D}}}e^{-\frac{\|\mathbf{x}\|^{2}}{2s^{2}}} \qquad (4)$$

The γ parameter in equation 3 defines a family of γ *normalised derivatives* and allows for a meaningful comparison of the response of differential operators across multiple scales. In order that the selected scale reflects the width of the ridge, for the ridge saliency metric the γ parameter is set to γ =3/4, as detailed by Lindeberg[13]. The Hessian is the matrix of second derivatives as defined in equation 5, where L_{xx} is the second derivative with respect to x, L_{yy} is the second derivative with respect to y and the mixed partial derivative is L_{xy}.

$$H=\begin{bmatrix} L_{xx} & L_{xy} \\ L_{yx} & L_{yy} \end{bmatrix} \qquad (5)$$

The eigenvalues and eigenvectors of the Hessian allow for geometric interpretation of the image. The eigenvector, whose eigenvalue, λ_2, has the largest absolute value, corresponds to the direction of greatest curvature i.e. rate of change of intensity. The eigenvector, whose eigenvalue, λ_1, has the smallest absolute value, corresponds to the direction of least curvature i.e. along track. The eigenvalues denote the respective amounts of curvature.

The eigenvalues of the Hessian are sorted in terms of absolute value such that:

$$|\lambda_1|\le|\lambda_2| \qquad (6)$$

In an ideal ridge structure λ_1 would approach zero, with λ_2 being large in comparison. Since tracks in CCD imagery are typically dark on a white background, a further constraint is placed on the polarity of λ_2 that it must be positive. In contrast to eigenvalue ratio methods[11], due to the along-track amplitude variance of CCD tracks (see figure 1), a Ridge Saliency (RS) metric is derived solely on γ-*normalised* λ_2, analogous to the method used by Steger[14]. The RS is defined as in (7) and is assessed across all scales.

$$RS=\max(\lambda_2>0) \qquad (7)$$

No consideration is paid to λ_1, meaning that the metric is sensitive to both elongated and blob-like structures. However, as can be seen in figure 1, tracks in CCD images present themselves as a series of rectangles in line with the path orientation rather than a single continuous ridge[15]. As such it is advantageous to have a metric which is sensitive to both structures. Figure 1 also demonstrates that λ_2 provides a consistent feature with which to identify track pixels.

A ridge saliency metric is computed for the *nsCCD* image to determine track-like features not attributed to shadows. Scales are chosen to match the size in pixels of features wishing to be extracted. The Gotcha dataset has a pixel spacing of 0.2m, hence the ridge saliency metric is computed for scales 1-10 in order to detect tracks of size 0.2 -2m in width. The ridge saliency metric is normalized prior to combination with the direction layer such that the values span the interval [0, 1]. The normalized ridge saliency metric is referred to as nRS.

The performance of the ridge saliency metric is assessed and compared against standard measures of ridge strength[13] such as; the γ-normalized maximum absolute principal curvature (equation 8), γ-normalized square principal curvature difference (equation 9) and square of the γ-normalized principal curvature difference (equation 10).

$$M_{\gamma-norm}L = \max(|\lambda_1|, |\lambda_2|) \qquad (8)$$

$$N_{\gamma-norm}L = (\lambda_2^2 - \lambda_1^2)^2 \qquad (9)$$

$$A_{\gamma-norm}L = (\lambda_2 - \lambda_1)^2 \qquad (10)$$

2.3 Direction filter

As a by-product of the ridge saliency computation, the eigenvector corresponding to the direction of greatest curvature is calculated. The eigenvector x, y components relating to the direction of greatest curvature are converted to an angular map of direction spanning [-π,π] using equation 11.

$$Dir = \arctan\left(\frac{v_{2y}}{v_{2x}}\right) \qquad (11)$$

An example angular direction map can be seen in figure 4A where it can be seen that this provides a direction measure for each pixel which can be used to filter out regions, such as vegetated areas, which present rapid changes in principal curvature direction over small pixel neighbourhoods. This is in contrast to tracks which show reasonably constant direction vectors along the track or slowly changing direction vectors in small pixel neighbourhoods.

To reduce clutter prior to vectorisation and remove contributions from vegetated areas a direction filter layer is constructed. The angular range of the direction map is split into 10 sub-regions e.g. ([–π <Dir < –9π/10], [–9π/10 <Dir < –8π/10],... ,[8π/10 <Dir < 9π/10], [9π/10 <Dir < π]) and for each sub-region, the direction map is thresholded to find clusters with small variations in direction. Morphological area opening is applied to the resultant 10 binary images to remove any noise attributed to connected components of fewer than 500 pixels (as shown in figures 4B and C). The resulting 10 binary masks are then summed to provide a direction filter layer. This is then used to mask the trackness index image prior to vectorisation.

A threshold is applied to the trackness index image to provide a binary image with the desired probability of detection and false alarm rate. The direction information for each of the remaining track components is then used to vectorise the final detected regions.

3 RESULTS AND DISCUSSION

Figure 1A illustrates the amplitude properties of a CCD image chip containing a track. The along-track variation in amplitude can clearly be seen and is consistent with the work carried out by Turner[15], where it is observed that tracks in CCD images present themselves as a series of rectangles in-line with the path orientation, rather than a continuous ridge feature.

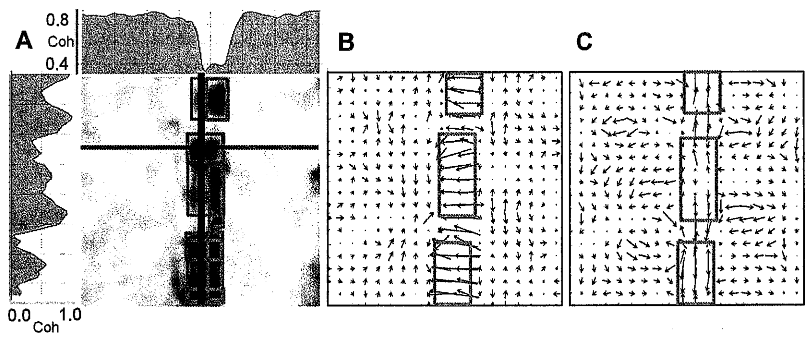

Figure 1: [A] Example chip from the CCD image, cross track (upper panel) and along track (left hand panel) amplitude variations for the red crosshair locations are shown. The typical rectangular features of the track are highlighted by blue boxes. The associated principal directions of curvature are shown for the same chip; [B] first eigenvector $\underline{v_2}$ scaled by λ_2, illustrating the direction of greatest curvature and [C] eigenvector $\underline{v_1}$ indicating the direction of minimal curvature, scaled by λ_1. The track portion shown in [A] is highlighted using blue boxes in [B] and [C].

Figures 1B and C show the impact of the CCD noise and along-track variation in amplitude on the extracted eigenvalues and eigenvectors. It is evident from the variation in amplitude and polarity of λ_1 seen in figure 1C that ridge saliency metrics that include λ_1 as a parameter are likely to introduce noise into the saliency metric and produce less robust, more disjointed track detections. Many ridge detection techniques take a ratio of λ_2 and λ_1, or assess differences in the two eigenvalues, under the assumption that in the along-track direction the degree of curvature will be minimal. However, figure 1A and C clearly demonstrate that this is not the case for tracks in CCD images. In comparison, figure 1B illustrates the suitability of the λ_2 parameter for characterising the track as a strong contrast can be seen with respect to the background patch with the feature being persistent along the length of the track portion.

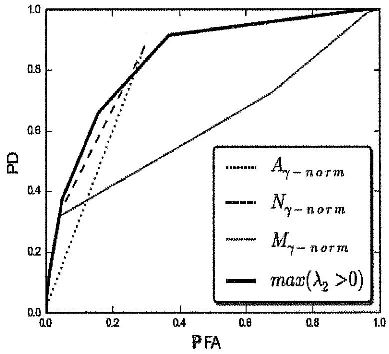

Figure 2: Pixel-wise ROC curve comparing 4 methods of ridge detection. It can be seen that, even with very lenient thresholds, the $A_{\gamma\text{-norm}}$ and the $N_{\gamma\text{-norm}}$ techniques fail to recover all of the track features.

Figure 2 illustrates the suitability of each of the ridge saliency metrics described above for CCD track extraction. Pixel-wise target detection receiver operating characteristic curves were produced for each of the measures and it can be seen that although the $A_{\gamma\text{-norm}}$ and $N_{\gamma\text{-norm}}$ methods perform reasonably well with a relatively low Probability of False Alarm (PFA), both techniques fail to recover all of the track features even at very low thresholds (a maximum of 88% of track pixels are detected). This may be due to the fact that they both rely on differences between the maximum and minimum eigenvalue. The granular nature of the CCD image means that large amounts of noise are present in the minimum eigenvalue, as seen in figure 1B, introducing noise and hence disjointed track detections into the final product. Furthermore, as a consequence of assessing differences in principal curvatures, both $A_{\gamma\text{-norm}}$ and $N_{\gamma\text{-norm}}$ are less sensitive to blob-like structures which are likely to be components of the CCD tracks at some scale. The $M_{\gamma\text{-norm}}$ metric, in contrast, having sensitivity to both blob and ridge like structures, detects all pixels present in the track. However, since no condition on the sign of the maximum eigenvalue has been applied for the plot shown in figure 2, the false alarm rate is poorer than the other methods as both bright and dark ridges are detected. The method adopted here for the track enhancement filter, shown by the solid black line, shows an improvement in performance over the other ridge detection techniques, although further modifications are required to ensure a useable performance with suitable probability of detection and minimal false alarms.

Figure 3: [A] CCD image chip, [B] normalized ridge saliency using the standard CCD product, [C] normalized ridge saliency using the *nsCCD* product to remove the effects of shadow regions.

Figures 3 A-C show an example portion of the Gotcha dataset at various stages of the track enhancement filter. Figure 3A shows a mixture of scene content including a track running down the right hand side, buildings, vehicles and vegetation. Notable areas of false change can be seen attributed to vegetation and building shadow. Figure 3B shows the normalised ridge saliency metric applied to the standard CCD product. High values of nRS indicate that a pixel is more likely to belong to a dark ridge. However, it can be seen that dark linear building shadows give strong nRS values comparable to that of the track. Furthermore, since the nRS metric is also sensitive to blob-like structures, some parts of the vegetated areas respond strongly. Figure 3C shows the normalised ridge saliency metric when applied to the *nsCCD*. The benefits in terms of false alarm reduction can clearly be seen with shadow regions no longer responding strongly in the nRS metric.

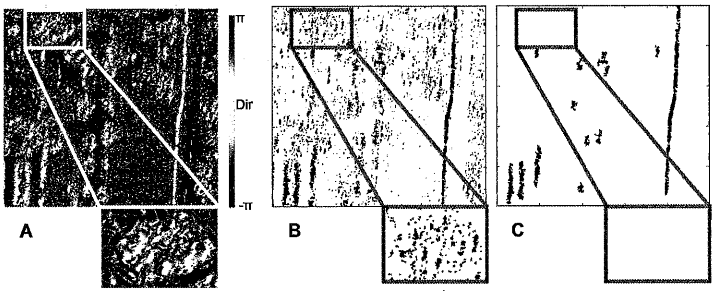

Figure 4. For the example image chip shown in figure 3 a stage of the direction filter is shown. [A] Angular representation of the eigenvector $\underline{v_2}$ illustrating the direction of greatest curvature, the zoom box highlights a vegetated area. [B] A filtered version of [A] is shown for $-7\pi/10 < \text{Dir} < -6\pi/10$, the zoom box highlights the corresponding vegetated area. [C] The binary filtered image in [B] is shown following morphological opening to remove small clusters. Contribution from the vegetated area is completely removed, as can be seen in the zoom box.

Figure 4A shows the angular map direction of greatest curvature. It can be seen, from the zoom box, that this presents a method for reducing false alarms introduced by vegetation as these areas show rapid local changes in direction when compared to the track portion. The intermediate stages of the direction filter layer are shown in figures 1B and C. The broad vegetated area shows fewer and smaller areas of constant direction than the tracks and morphological area opening removes the contributions from the region entirely. Some small responses remain from areas of the vegetated patches which appear as blobs of constant direction on a scale similar to that of the feature to be detected. This could potentially be improved by segmentation of the reference and match image prior to CCD creation and, based on local grayscale texture, setting those areas to unity in the CCD images, as was implemented for the shadow mitigation. This is an area for future work.

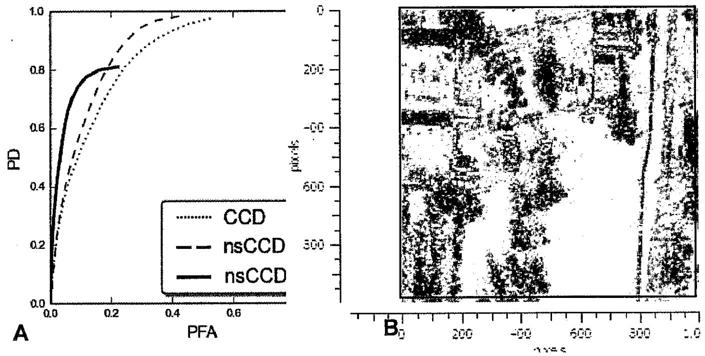

Figure 5: [A] Pixel-wise target detection ROC curve for standard ridge analysis and ridge analysis using basic metrics to remove contributions from regions of low return and vegetation. CCD denotes the basic normalised ridge saliency measure of the standard CCD product. *nsCCD* denotes the normalised ridge saliency measure of the *nsCCD* product with shadow regions masked out. *nsCCD* & Dir refers to shadow masking prior to ridge detection followed by direction filtering to remove contributions from vegetation. [B] Example CCD chip overlaid with the vector layer (red) extracted using the track enhancement filter with threshold set to 0.15 to give a PFA of 0.03.

It can be seen in figure 5A that inclusion of shadow masking of the CCD prior to ridge detection improves the performance of the track detector and that for low PFAs including the direction filter can further enhance the performance. For a PFA of 10% the combined use of the shadow masking and direction filter resulted in a probability of detection of 78% in comparison to 52% for the simple line detector. However, inclusion of the direction filter does mean that not all track pixels are detectable even at the most lenient threshold. It should be noted that since the ground truth was marked up manually at the judgement of the observer it will not be entirely accurate and a properly ground truthed dataset would be needed to produce more reliable ROC curves. Since the direction information is obtained during the analysis of the Hessian matrix and requires little extra processing time the results presented here suggest that it may be beneficial to include in track extraction algorithms.

Figure 5B shows an example of the extracted vector layer overlaid in red on the original CCD. It can be seen that the track pixels are fairly well detected, with some false alarms in vegetated areas as previously discussed. The remaining false alarms appear in areas demonstrating layover differences in bright returns. No attempt to correct for this was made in this work; however an extra layer could easily be included in the process to reduce the weighting of these areas on the trackness index. On the whole the performance of the algorithm is reasonable, but an area of future work would be to include a method, such as Dijkstra's technique[16] to connect the disjointed path segments detected using the trackness algorithm, this may also help to mitigate some of the small responses in vegetated areas spatially distinct from other track features.

The Gotcha dataset, although having a good mixture of scene content, has a limited number and type of tracks. A more meaningful assessment of the algorithm could be carried out on a dataset with a wider variety of track types in scene to further interrogate the strengths and weaknesses of the technique.

4 CONCLUSION

We have demonstrated an algorithm to extract a range of tracks from a CCD image pair. The algorithm has the flexibility to detect both single and in principle parallel tracks (this has not been tested as there were no visible parallel tracks in the dataset tested). The algorithm can tolerate tracks of variable widths and sensitivity to false change in vegetated or shadow regions has been reduced resulting in a moderate false alarm rate. Future work will comprise further testing of the algorithm on a more extensive set of data, including a wider variety of scene content and track types.

4 REFERENCES

1. M. Bicego, S. Dalfini, G. Vernazza, V. Murino. Automatic road extraction from aerial images by probabilistic contour tracking. Proc. of IEEE Int. Conf. on Image Processing. 3:585–588. (2003).
2. D. Klang. Automatic detection of changes in road databases using satellite imagery. Proc. Int. Arch. J. Photogrammetry and Remote Sensing. 32:293-298. (1998).
3. R. Bajcsy and M. Tavakoli. Computer recognition of roads from satellite pictures. IEEE Trans. Syst., Man, Cybern. SMC-6:623-637. (1976).
4. M. Preiss, D.A. Gray and N.J.S. Stacy. Detecting scene changes using synthetic aperture radar interferometry. IEEE Geosci. Remote Sens. 2451-2453. (2006).
5. H.A. Zebker and J. Villasenor. Decorrelation in Interferometric Radar Echoes. IEEE Geosci. Remote Sens. 30(5):950-959. (1992).
6. A. Bouraba, A. Younsi, A. Belhadj-Aissa, M. Acheroy, N. Milisavljevic and D. Closson. Robust techniques for coherent change detection using COSMO-SKYMED SAR images. Prog. Electromagn. Res. 22:219-232. (2012).
7. I. Stojanovic, L. Novak. Algorithms improve synthetic aperture radar coherent change detection performance. SPIE Newsroom, doi:10.1117/2.1201307.004889. (July 2013).
8. M. Cha and R. Phillips. Automatic track tracing in SAR CCD images using search cues. IEEE Conference Record of the Forty Sixth Asilomar Conference, Signals, Systems and Computers (ASILOMAR). 1825-1829. (2012).
9. R. Phillips. Activity Detection in SAR CCD. IEEE International Geoscience and Remote Sensing Symposium (IGARSS). 2998-3001. (2013).
10. M. Newey, J. Barber, G. Benitz and S. Kogon. False alarm mitigation techniques for SAR CCD. IEEE Radar Conference (RADAR). 1-6. (2013).
11. A.F. Frangi, W.J. Niessen, K.L. Vincken and M.A. Viergever. Multiscale vessel enhancement filtering. In Medical Image Computing and Computer-Assisted Intervention - MICCAI'98, Lecture Notes in Computer Science, Springer Verlag, Berlin, Germany, 1496 : 130-137. (1998).
12. S. Scarborough, L.Gorham, M.J. Minardi, U.K. Majumdar, M.G. Judge, L. Moore, L.Novak, S.Jaroszewski, L. Spoldi and A. Pieramico. A challenge problem for SAR change detection and data compression. Proc. SPIE 7699, p. 76990U. doi:10.1117/12.855378. (2010)
13. T. Lindeberg. Edge Detection and Ridge Detection with Automatic Scale Selection, J.Comp.Vis. 30(2): 117-156. (1998).
14. C. Steger. An unbiased detector of curvilinear structures. IEEE Trans. Pattern Anal. Mach. Intell. 20(2):113-125. (1998).
15. E. Turner, R.D. Phillips, C. Chiang and M. Cha. Inserting simulated tracks into SAR CCD imagery. Modeling & Simulation International (SCS) 2012 Autumn Simulation Multi Conference, San Diego, CA. (2012).
16. E.W. Dijkstra. A note on two problems in connection with graphs. *Numerische Mathematik.* 1:269-271. (1959).