# Region Growing Shadow Segmentation in Synthetic Aperture Radar Images 

Kelce S. Wilson and Gregory J. Power<br>Air Force Research Laboratory, AFRL/SN<br>Wright-Patterson AFB, OH 45433


#### Abstract

This effort introduces a noise-tolerant algorithm for segmenting shadow regions in Synthetic Aperture Radar (SAR) images, using both edge detection and a region-grow method for filling between detected edges. The technique introduced here recognizes that real shadow boundaries are neither precisely known, nor are infinitesimally thin, due to SAR phenomenology such as motion and layover. As the identified shadow region transitions from actual shadow to target and clutter regions, a metric based on the rate-of-change of the region size allows automatic detection of the shadow edge. This works in concert with traditional edge detection methods to identify boundaries more reliably. An evaluation of the quality of the shadow segmentation is presented using multiple quality measures, including a percent-pixels-same measurement that is based on manual segmentations of the same image. The algorithms and evaluation is accomplished on SAR target image chips obtained from the Moving and Stationary Target Acquisition and Recognition (MSTAR) program sponsored by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL).


## 1 Automatic Segmentation

An automatic segmentation algorithm was developed earlier to identify target shadow regions for phenomenology studies and comparisons between artificially generated SAR images and their equivalent measurements [1]. It assumed an infinitesimal boundary between shadow and non-shadow regions, and used a region grow algorithm to find the border. One of the main discoveries of this initial work was that the edges of the artificial shadows and measured shadows differed significantly in terms of abruptness. This is because the
artificial shadows were calculated using the geometry of a CAD model and simulated illumination from a single point placed at the phase center of the SAR, while the actual SAR image was formed from integrating a series of illuminating angles. In each of the different real illuminating angles, the shadow would shift slightly, giving three classifications, rather than just two, for the clutter pixels surrounding a target: always illuminated, always shadowed, and sometimes illuminated but otherwise shadowed. This integration of shifting shadows produces a shadow region with blurred edges. As a result, while the region-grow metric would easily find the boundary of an artificial shadow, the boundary of a real shadow could often be missed. Figure 1 shows a comparison of the metric used for automatic boundary detection when applied to a representative artificial and measured image pair. The metric is the size of the identified shadow region with various trial values in a key comparison threshold; the algorithm searched for the largest jump in the metric for deciding on the proper value. As can be seen clearly, the jump for the measured image is much smaller.


Figure 1: Metric comparison for representative artificial and measured image pair

Unfortunately, the response of measured images to the metric makes reliance upon the single input impractical. When edge detection is
applied to an image, however, a relative maximum value of the convolution results, in the direction of an outwardly growing region, can be used as a second input to the automatic edge condition identifier. Such an additional input makes the algorithm considerably more robust. In this scheme, the region can "grow into" the shape determined by the edge detection process, with use of the earlier metric reduced mainly to identifying "leaks" in the detected edge. A leak, as used here means a break in an otherwise fully connected series of pixels identified as an edge. That is, if a group of pixels identified as belonging to an edge, survive minimum-cluster size and threshold value tests, they should ideally form a continuous bounding region surrounding the shadow region. This is rarely the case for measured images, though. The pixels that are missing and cause holes in the region are then named leaks. With simple neighborhood comparisons for the edge detection results, however, most small leaks can be patched before the region-grow begins.

After the initial edge detection kernels have been convolved with the image, and the result saved in a reference matrix, the algorithm identifies a point expected to be inside the shadow. It traverses the image in all four directions, up, down, left, and right, searching for either a series of consecutive unbroken threshold violations, or a maximum edge convolution value. This is shown in Figure 2 for a representative target chip.


Figure 2: Initial Shadow Region Growth

The region growth is an iterative process which is tried again at all pixels reclassified from an initial state of "unknown" to "shadow", and is finished when no more pixels change classification. Such a method can easily handled punctured regions, although it must be simply connected. Second-pass clean-up consists of a reclassification of pixels within the outermost shadow boundary that were initially not identified as shadow, but are not part of a cluster meeting certain dimension requirements, and can also include smoothing the outer boundary.
Figures 3 through 7 demonstrate the automatic thresholding portion of the algorithm. Figure 3 shows a clean view of the demonstration target, and Figure 4 shows the detected edges of the shadow used in next section of the algorithm.


Figure 3: Shadow Demonstration Target
The region grow program uses thresholds for both the interior/exterior decision and also the edge/not_edge decision. If the edge/not_edge threshold is set too low, the inputs to the region grow portion will be full of false alarms. The edge pixels are checked to determine if they belong to a line of some minimum length and a cluster of some minimum size, rather than just resulting from isolated noise points After this clean-up, the edge threshold is set to the minimum necessary for predominant coverage around a simply closed region. This part is still done by a human operator, so it is somewhat subjective. Certain leaks of a few pixels wide must be tolerated in the real edge in order to
eliminate false edges on the interior of the region.

A leakage potential is evident at the top of Figure 4, where the edge region thins, and the single pixel connecting the wider curves happens to fall below the edge threshold necessary for the region to grow to its "natural size". The natural size of the shadow is reached when the shadow boundary falls generally near the center of the thicker regions of pixels identified as possible edges. The "natural edge" pixels are the darkest ones evident in Figure 4. The lighter colored pixels are false edges, so the edge threshold is automatically adjusted to prevent these from stopping the growth process.


Figure 4: Detected Edges of Demo Target


Figure 5: Shadow Threshold Set Too Low
Figure 5 shows the target with the shadow threshold below the finally selected value. The uniform dark gray region of the image identifies the shadow region, and its border is the black band surrounding the shadow.

The shadow region fills most of the space inside the detected edges, but not quite all. Increasing the shadow threshold by 2 dB results in the shadow region shown in Figure 6. It shows a relatively small change for a large change in the threshold, which gives the gentle slope of the metric in Figure 1, before the final threshold is chosen. However, increasing the shadow threshold further, results in "spill-over" of the shadow and a dramatic increase in the region size reported by the growth algorithm.


Figure 6: "Natural" Shadow Region


Figure 7: Shadow Region Spill-Over From Top
This spill-over effect is clearly demonstrated in Figure 7 and corresponds to the large jump in the metric plotted in Figure 1. The leak point
can easily be identified as the thin part of the detected edge previously identified at the top Figure 4.
The algorithm starts with thresholds known to be too low, and increases until at least one-third of the pixels in the image are identified as shadows. This is known to be too high, based on the nature of the target chips used. The pixel counts corresponding to each threshold tested are examined for the largest jump. When the most rapid change in the region size is identified, the algorithm backs the threshold down to the last value before the jump. This is then the automatically selected threshold used for the final reported region.

Given that $F\left(x, y, \tau_{i}\right)$ represents a binary image pixel in an $N$ by $M$ image of ones and zeros where ones are selected based on the threshold criteria $\tau_{i}$, the threshold value, $T$, can then be determined by:

$$
T=\arg _{i i}\left\{\max \left[S\left(\tau_{\min }\right), S\left(\tau_{\min +1}\right), \ldots, S\left(\tau_{\max }\right)\right]\right\}
$$

Where $S\left(\tau_{i}\right)$ is defined as

$$
\begin{aligned}
& S\left(\tau_{i}\right)=\sum_{x=1}^{N} \sum_{y=1}^{M}\left|F\left(x, y, \tau_{i}\right)-F\left(x, y, \tau_{i-1}\right)\right| \\
& \text { for } \tau_{\min +1} \leq \tau_{i} \leq \tau_{\text {max }} .
\end{aligned}
$$

Here, $\tau_{\text {min }}$ and $\tau_{\text {max }}$ represent the minimum and maximum amplitude threshold used to segment the general region of interest.

## 2 Segmentation Scoring

The evaluation of image segmentation is an ongoing area of research, with a variety of currently used metrics. For this effort, three were chosen: percent pixels same (PPS) [2], partial-directed hausdorf [3] for omission errors (O-pdh) and commission errors (C-pdh), and the complex inner product (CIP) [4]. For comparison purposes, all three metrics are normalized to a scale of zero to one with one being perfect segmentation. The PPS metric provides a spatial measure of the percentage of pixels that are common to two different segmentations, but does not provide a measure of the shape matching. The O-pdh and C-pdh
are combined as a single metric for shape matching; they measure the distances between the two edges. The CIP is a scale-independent measurement of shape matching. In general, shape is the most important factor for many ATR algorithms.


Figure 8: Manual and automatic segmentation of test shadow. Top two sections compare automatic with two different manual segmentations, bottom section compares the two manual segmentations.

Table 1: Quantitative metrics for evaluating various segmentation processes

|  | Hand Segmentation \# 1 |  |  |  | Hand Segmentation \#2 |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | CIP | O-pdh | C-pdh | PPS | CIP | O-pdh | C-pdh | PPS |
| Hand Segmentation \# 1 | 1.00 | 1.00 | 1.00 | 1.00 | 0.23 | 0.25 | 0.27 | 0.79 |
| Hand Segmentation \# 2 | 0.23 | 0.25 | 0.27 | 0.79 | 1.00 | 1.00 | 1.00 | 1.00 |
| Automatic Segmentation | 0.39 | 0.14 | 0.14 | 0.59 | 0.52 | 0.03 | 0.04 | 0.53 |

Segmentations of clutter-filled SAR images are generally not easy to evaluate and score, since not only is the actual edge location not known, it may actually be a region with finite width. Additionally, the noise level and grainy resolution can combine to make inconsistent even a careful manual process inconsistent, without knowledge of the target and a calculated mask for desired shadow shape. An example is shown in Figure 8, which was segmented twice, on separate occasions, by the same individual. The top section has the results of automatic segmentation drawn as the inner contour and one of the manual segmentations as the outer contour, while the middle contrasts the automatic method with the second manual result. At the bottom of the figure, the difficulty of the task is clearly demonstrated by the differences between the two manual segmentations.

## 3 Results

The scores for the two hand segmentations and the automatic algorithm are shown in Table 1. A perfect score of 1.0 is obtained only when a certain result is compared with itself. The automatic method does not produce as large of an area as either of the manual processes, so it scores lower on PPS, O-pdh, and C-pdh. This can be expected, since these values depend on scale and spatial mass. The automatic segmentation does do well with the CIP metric, which scores on shape only. This suggests that the automatic technique is obtaining the key shape information from the SAR shadow. The maximum detected edges, however, are still somewhat "inside" the shadow as identified by a human evaluator.

## 4 References

[1] Wilson, K S, Ryan, P A, "Phenomenology Metric Development for SAR Scene Modeling Tools", Algorithms for Synthetic Aperture Radar Imagery VI, Zelnio, SPIE Vol. 3271, pp. 582-588, International Society for Optical Engineering, Orlando, 1999
[2] Power, G J, Weisenseel, R A "ATR subsystem performance measures using manual segmentation of SAR target chips", Algorithms for Synthetic Aperture Radar Imagery VI, Zelnio, SPIE Vol. 3271, International Society for Optical Engineering, Orlando, 1999
[3] Beauchemin, M, Thompson, KPB, Edwards, G., "On The Hausdorff Distance Used For The Evaluation Of Segmentation Results", Canadian Journal of Remote Sensing, pp. 38, Vol. 24, No. 1, March, 1998
[4] Power, G J, Awwal, A.A.S., "Optoelectronic complex inner product for image segmentation performance evaluation," Photonics Devices and Algorithms for Computing II, SPIE Vol. 4114, San Diego, 2000

