

AUTOMATIC EXTRACTION OF ROAD NETWORKS IN URBAN AREAS FROM IKONOS IMAGERY BASED ON SPATIAL REASONING

J. Gao and L. Wu

School of Geography and Environmental Science, University of Auckland, Auckland, New Zealand
jg.gao@auckland.ac.nz

Commission III, Working Group III/4

KEY WORDS: Remote Sensing; Mapping; Extraction; IKONOS; Automation; Urban

ABSTRACT:

In this study we developed a spatial reasoning-based method of automatically extracting roads in a densely populated suburb of Auckland, New Zealand from IKONOS data. First, all of the four multispectral bands were grouped into 20 clusters in an unsupervised classification, two of which corresponded to road networks. This intermediate result was then converted into a binary image of road and non-road pixels. This binary image was then further processed with spatial reasoning in two ways. First, all isolated or small clusters of pixels were examined spatially to determine if there were other isolated pixels in their immediate vicinity. If no neighbouring pixels were found, they were considered as noise and removed from the image. If neighbouring pixels were found, their position in relation to the pixel under consideration was further analyzed. If they were aligned with existing pixels along a certain orientation, then they were regarded as a portion of a disjointed road and retained in the output image. Second, these disjointed road segments were later joined together to form a road network. The extracted road network was unified to a constant width because trees planted along both sides of a road caused its width to vary in different sections. The detected results using a threshold of six pixels show that most roads can be extracted at a reasonable accuracy level.

1. INTRODUCTION

Roads in dynamic cities tend to change very frequently even within a short period of time. Road maps of these areas have to be updated periodically, preferably from current satellite images to meet the urgent need of urban planners. With the advances in remote sensing, more and more high quality and fine spatial resolution satellite images have become available from different platforms. For instance, the recently emerged IKONOS satellite imagery has a spatial resolution of 4 m in the multispectral mode and of 1 m in the panchromatic mode. These images enable the extraction of even minor streets in urban areas. They have raised a renewed possibility of timely and efficiently updating changed road networks in urban areas.

Extraction of road networks from remote sensing images can be accomplished either manually or automatically. Manual extraction is subject to the analyst's experience and skills. Roads can be recognized reasonably well even from noisy images that contain incomplete information about roads if s/he is familiar with the study area. However, this manual method is expensive and time-consuming. By comparison, automatic extraction of road network information involves significantly less time and expense, even though it is more complex methodologically.

Automatic extraction of roads from satellite images faces several challenges because the image appearance of roads depends upon the spatial resolution of the satellite images. In addition, the extraction is hampered by noise on satellite images. Ground objects such as trees along a street can obstruct the image of roads. Vehicles on the road may cover certain parts of a road and make it difficult to detect on the image.

So far various automatic methods have been developed to extract roads from satellite images. These methods fall into five broad categories: ridge finding, heuristic reasoning, dynamic programming (DP), statistical tracking, and map matching (Xiong,

2001). Ridge finding is a classic method in which an input image is edge-filtered to obtain the magnitude and direction of linear features, including roads (Nevatia and Babu, 1980). Wang et al. (1992) developed a way of detecting ridges from SPOT data. In this gradient direction profile analysis method four predefined directions for each pixel are calculated first and the gradient direction for a pixel is the direction of the maximum slope among the four defined directions around the pixel. The road segments have the same ridge direction and it is perpendicular to the gradient directions of the pixels with the bridge. Analysis of the gradient profile will generate the ridge pixels. The road network or segment can be obtained by linking all the ridge points. Steger (1996) introduced differential geometry to ridge finding. This method uses curve or surface fitting techniques to locate ridges on remote sensing imagery. If the image intensity surface is represented by a mathematical equation, the first and second derivatives of the equation can be analyzed to locate edges.

In DP, roads are modelled as a set of mathematical equations. The derivatives of the grey values perpendicular to the direction normal to the road tend to be maximized, while derivatives along the road direction are minimized. Roads appear to be straight lines or smooth curves. Their local curvature has an upper bound. DP is advantageous in finding curves in noisy pictures, for it can bridge weakly connected feature elements automatically while the program searches for optimal solutions (Gruen and Li, 1995).

Statistical inference models are particularly suitable for detecting roads with complexity and uncertainty (e.g. bridges, road width variation, vehicles and shadows on the roads and image noises, etc). Barzohar and Cooper (1996) explored the method further and developed a stochastic approach that can be applied to automatic extraction of highly sophisticated roads. A geometric-stochastic model formulates road width, direction, grey level intensity and background intensity as a

stochastic process using Gibbs distributions. Roads are found by maximum a posteriori probability estimation.

Geman and Jedynak (1996) developed another statistical model to track roads through hypothesis testing. This approach uses the testing rule that is computed from the empirical joint distributions of tests (matched filters for short road segments) to determine whether the hypothesis (road position) is true or not. The tests are performed sequentially and an uncertainty or entropy minimization procedure is devised to facilitate testing decisions so that new tests can be analytically identified. Although this method works best for coarse resolution images, it is also adaptable to fine-resolution images (Xiong, 2001). Tupin *et al.* (1998) proposed a two-step algorithm for almost unsupervised detection of linear structures which are treated as road segment candidates. During the first step, linear features which are treated as road-segment candidates are extracted. In the second step, genuine roads among the segment candidates are identified by defining a Markov random field (MRF) on a set of segments, which introduces contextual knowledge about the shape of road objects.

The map matching method is adaptable for road network extraction in case of a large amount of data on road systems in many parts of the world (Maillard and Cavayas, 1989). It consisted of two major algorithms. The first algorithm focuses on image-map matching to identify roads that can be found on the map and the image. The second algorithm searches new roads based on the assumption that these new roads are connected to the old ones. This automatic approach is suitable for revising 1:50 000 scale planimetric data using panchromatic SPOT imagery. Stilla (1995) developed a syntax-oriented method that uses map knowledge as a supportive aid for image interpretation. Road network structures are obtained first through map analysis. Then image object models are defined and utilized to search for objects that fulfil model expectations with a given tolerance. Assessment on image objects with respect to its correspondence to the map representations results in road object identification for a given image scene.

Without pre-defined parameters or setting any threshold or to describe statistically the classes to be extracted in pattern recognition, a neural network is an alternative way in road network extraction. The commonly used is back-propagation network. Fiset and Cavayas (1997) described a map guided procedure to automatically extract road network using this network. Bhattacharya and Parui (1997) improved back-propagation neural network for detection of linear feature from IRS and SPOT images with a satisfactory result.

Differential shapes, also known as deformable contour models, is a novel approach for the integration of object extraction and image-based geospatial change detection (Agourls, 2001). In this method, the extraction is carried out at three levels. At the low level of analysis images are characterized and information extracted on a pixel by pixel basis using only reflectance/emission measurements (e.g., filtering, edge detection, segmentation, etc.). At the intermediate level of analysis low level results are symbolized and fused to form data structures for the high level analysis. The high level analysis involves not only imagery but also domain specific knowledge, ancillary sources, and symbolized data from the intermediate level results to establish reliable interpretation. The program first runs a template-matching procedure that localizes potential road pixels, followed by the optimization to identify potential road segments. To allow a more inclusive search, a relaxed road model is utilized during this search. After this search processing, all segments found will be considered as a road candidate. Then the supervised ISODATA

classification procedure is applied to identify whether or not a candidate is indeed a road.

Also known as rule- or knowledge-based method, the heuristic method makes use of the human vision system. Meisels and Mintz (1990) developed a three-stage reasoning method for the extraction of simple man-made objects from aerial photography. In the low level, image primitives are considered as the building blocks of road and identified the value by checking the neighbors' value. During the intermediate level analysis image primitives are combined with line segments by using the reasoning mechanism. At the high level of processing gaps are filled and segments grouped by taking into consideration of distance, brightness and uniformity among them. This method is flexible when the problems concern linear feature alignment and fragmentation.

The above review indicates that a large number of studies have been carried out to detect roads from remote sensing images. Some of them have produced satisfactory results. However, they are limited in that they are designed for the extraction of roads from a particular type of imagery. Whenever a different kind of data or geographic region is involved, these methods are utterly incapable of the extraction.

In this study we propose a new spatial reasoning-based approach to extraction of road networks in urban areas from 4 m resolution IKONOS imagery. In this environment other urban built-up areas can be easily mixed with roads. Besides, urban vegetation can also obscure image properties of roads. At such a fine resolution level, roads of various widths and conditions have their unique spectral properties on the imagery. All of these factors make their extraction challenging. The challenge is overcome with spatial reasoning which takes the spatial arrangements and spatial relationship among road pixels into decision making. This extraction has been implemented in an environment which is not subject to the spatial resolution of the input imagery.

2. METHODOLOGY

2.1 Image properties of roads

As a kind of artificial entities, roads differ from other human-made and natural objects in their surface material and geometric configuration. Namely, they are usually paved with asphalt or concrete, causing them to reflect incident solar radiation strongly. A road is linear in shape and has a uniform width which varies with its significance. Highways and artillery roads are wider than inner city streets.

On remote sensing imagery roads are recognizable from their distinctive tone, texture and shape. Strong reflection of the incident radiation causes them to have a uniform light tone on remote sensing imagery. Consisting of the same paving material, roads have a smooth and identical texture. There is a drastic change in tone and texture across road boundaries. Geometrically, roads are elongated and spatially continuous. They also have a uniform width with small curvatures. Their physical appearance on an image, however, is affected by image spatial resolution. On a coarse resolution image, a road may not even register or appear to be a narrow line represented by a string of pixels. On a fine resolution image, however, the same road may appear as an array of cells with two parallel boundaries. The spatial continuity of a road

network may be disrupted by the presence of trees and vehicles on the ground, causing it to be intermittent. Any reliable extraction of road networks must take full advantage of these characteristics.

2.2 Data preprocessing

The original image is a subscene multispectral IKONOS image. It covers a densely populated suburb of Herne Bay in Auckland, New Zealand. The raw image was processed using the unsupervised classification method during which all pixels were grouped into 20 clusters in ERDAS Imagine. During post classification it was found that two of these clusters corresponded to roads. The classified image was then converted into a binary image of road and non-road pixels, and saved in the bitmap format, the commonly used and recognized image format by most image processing systems (Figure 1).




Figure 1. A binary image obtained from unsupervised classification of the raw multispectral IKONOS image into 20 classes in ERDAS Imagine.

This binary image shows clearly the outline of major streets. Due to the complexity of the scene (many buildings and cars parked along streets) and the limitations of the per-pixel based unsupervised classification method, noises of non-road pixels are quite common and widespread in this binary image. Appearing as isolated pixels or pixels in small clusters, these noises can be eliminated through spatial filtering. On the other hand, many road segments, which are spatially contiguous in reality, appear to be disjointed with a varying width in the image. These imperfections will be improved through subsequent spatial reasoning to make the extracted road network more reliable.

3. IMPLEMENTATION OF SPATIAL REASONING

The proposed method of road network extraction based on spatial reasoning consists of three stages after the initial unsupervised classification. They are noise removal, road segment joining and thinning.

3.1 Noise removal

Roofs of buildings in urban areas have an image tone similar to that of roads. Consequently, they have also been grouped into road pixels during the unsupervised classification. A close inspection of the results (Figure 1) reveals that building pixels are spatially isolated, or are in small clusters. These clusters, however,

do not contain too many pixels. A search of nearby pixels shows that these non-road pixels are not aligned with road pixels in any direction. If cluster size is taken advantage of to distinguish road pixels from non-road pixels, the distinction relies on a specified threshold for the number of pixels contained in a cluster and their spatial arrangement with another cluster of pixels. Care must be exercised in determining an appropriate threshold before any further processing is undertaken. The larger the threshold, the greater the number of noisy pixels, and the more interpretable the image becomes because any clusters of spatially contiguous pixels with a membership below this threshold will be regarded as noise and removed. On the other hand, however, a larger threshold may lead to the loss of information in the output image as a broken road segment may be made up of a small number of pixels.

All the road segments having fewer than the specified threshold of pixels will be treated as noise and subsequently removed. During this process, all the isolated pixels or those in small clusters are removed. All those remaining clusters contain enough spatially contiguous pixels above the threshold. Their longest length is then calculated and compared with the threshold. If the length is shorter than the specified threshold, then all pixels in the cluster will be removed. Those remaining pixels are considered to represent true roads. This noise removal procedure consists of several steps. The first is to set up a threshold for the shortest length of a road segment in the image. Length is defined as the longest dimension of a cluster of spatially contiguous pixels. This user-defined threshold governs the cluster size of noise pixels. Assume the pixel under consideration is located at (i, j) in which i varies from 0 to image column minus 1, and j from 0 to image row minus 1. The input image is processed pixel by pixel iteratively. During each iteration, i is incremented by one until it reaches image column minus 1. Within each loop j is incremented by one until it reaches image row minus 1. Through these iterations all pixels will be processed.

There are two public arrays of variable in the noise removal procedure, one storing pixel value and the other storing a Boolean value that indicates whether it is a road pixel. During every iteration, each of the four neighbouring pixels of (i, j) , $(i-1, j-1)$, $(i-1, j)$, $(i-1, j+1)$ and $(i, j-1)$, is examined in turn to determine whether pixel (i, j) is a road pixel. If one of the four neighbouring pixels is a road pixel according to the properties described previously, then pixel (i, j) is also considered a road pixel. Otherwise, the search algorithm is activated to search the other four directions: $(i+1, j-1)$, $(i+1, j)$, $(i+1, j+1)$, $(i, j+1)$ to check whether pixel (i, j) is the start of a road segment.

Determination of the length of a cluster of spatially contiguous pixels is accomplished recursively. The structure of the recursive search is illustrated in Figure 2. Pixel $(i+1, j-1)$ is evaluated first. If it is a road pixel, then its coordinate is set to (i, j) and the threshold of road length is decreased by 1. Afterwards the program recalls itself. If pixel (i, j) is not a road pixel or lies outside the image bound, then the recursion returns a false value and starts the next direction. If the threshold is 0, then it returns a true value. Four directions of search are considered: $(i+1, j-1)$, $(i+1, j)$, $(i+1, j+1)$, $(i, j+1)$. After all pixels in the input image have been searched, all the contiguous pixels will be calculated. If their number falls below the specified threshold, then all of them will be regarded as noise and removed from the output image.

3.2 Joining road segments

There are many methods by which the retained road segments can be joined together, such as Hough transform (global), optimal search (local), direction filter (local), overlaying with GIS layers, and so on. Because of the characteristics of the input image, the directional cone search method was used in this study. The joining procedure consists of several steps. The first step is to identify the end of a road segment through iteration under the assumption that both the end pixels of a segment are the ends of a road. All pixels in the image are searched for to determine whether they are the ends of a road. If a pixel is a road pixel, then a search is carried out further to determine if there are pixels to its left and its right. If so, then this pixel is not the end of a road segment. If pixels are found to lie next to the pixel under consideration, either to its left or to its right, but not to both of them, then the pixel under consideration is confirmed as the end of a road. A search is then carried out to determine whether there is another road end in its vicinity. If so, and if both clusters are oriented in the same direction and the distance between them is smaller than the specified threshold, then they are regarded as different segments of the same road, and will be joined together to form a longer road.




Figure 2. Structure of the recursive search.

This operation to determine whether pixel (i, j) is the middle of a road or its end is executed like this. If pixel (i, j) is a road pixel, then at least one pixel around it must be a road pixel. The pixel above it $(i-1, j)$ will be evaluated first. If it is also a road pixel, then pixels $(i, j+1)$, $(i+1, j+1)$ and $(i+1, j)$ will be evaluated, as well. If one of them is a road pixel, then pixel (i, j) is considered as the middle of a road. No further processing is undertaken. The program moves on to the next pixel. If none of pixels $(i, j+1)$, $(i+1, j+1)$ and $(i+1, j)$ is a road pixel, then pixel (i, j) must be the end of a road. The road segment joining operation is performed on it. If pixel $(i-1, j)$ is not a road pixel, the program will move on to

pixel $(i-1, j+1)$. If it is a road pixel, pixels $(i+1, j)$, $(i+1, j-1)$ and $(j, j-1)$ will be evaluated to determine if any of them is a road pixel. If one of them is a road pixel, then pixel (i, j) forms the middle of a road and the iteration shifts to the next pixel. Otherwise, pixel $(i, j+1)$ will be processed in the same manner as before. This process is repeated until the last pixel inside the operating window. Because (i, j) is not an isolated pixel, one of its eight neighboring pixels must be a road pixel.




Figure 3. The downward road segment to be joined with the end of a road segment linking (i, j) .

If the pixel is found to be the end of a road, the road segment linking operation is performed on it. In this operation, the cone along the direction is first defined. If pixel (i, j) is a road pixel and pixel 1 is also a road pixel (Figure 3) and none of pixels A, B and C is a road pixel, then the search join is downwards. In Figure 3 the directional cone is defined as $D \rightarrow E, F, G \rightarrow H, I, J, K, L \rightarrow M, N, O, P, Q, R, S \rightarrow \dots$. The join search is terminated by either finding the road pixel or meeting the designated threshold. If a road pixel is found during the search process, then its coordinates will be returned as $((i, j)$ and the pixel found) for the subsequent connection. If no road pixel is found and the iteration has reached the threshold, then this pixel is the end of the road, and the join operation will move on to the next pixel. Similar road segments in the directions of north, east, west, southwest, northwest, northeast, and southeast were also taken into account in the join operation.

After linking, dangling road segments were joined with those in their immediate vicinity.

3.3 Thinning

Noises in remote sensing images can make the width of a road vary in different parts. They can also make roads thicker than the unit width in remote sensing images. Such variations can be detected by humans in the manual method of road extraction. In the automatic method of extraction, width of road networks can be unified through thinning. There are five requirements for road thinning (Choy, 1994). Connected image regions should be thinned to connected line structures (connectivity preservation); approximate end-line locations should be maintained (no excessive erosion); thinned results should be minimally 8-connected (unit – width skeleton); the thinned results should approximate the medial lines (medial line approximation); and extraneous spurs caused by thinning should be minimized (boundary noise immunity). There are

many different methods to implement the thinning algorithm. The most common one is to successively “peel” the outermost layers off the objects until they are one pixel wide on the connected network. Usually a 3×3 window is applied to the image to decide if the central pixel should be removed.

The thinning algorithm adopted in this study is a two-stage, single-pass parallel operation within a 3×3 moving window. The first stage involves removal of redundant pixels to generate a medial line according to the set of the templates. The templates selected to remove the central pixel P should have the following two properties: (a) removal of P should not disrupt the local connectivity of the 3 by 3 pattern; and (b) removal of P should not incorrectly deform the shape of the pattern. The purpose of this stage is to determine if pixel P should be removed or retained without destroying the local connectivity. If the removal of P destroys the global connectivity, then P should be restored. Therefore, the second stage is the set of templates to restore P to maintain global connectivity. In this set four additional pixels located at above, below, and on both sides of P are used to determine global connectivity. If the central pixel P has the coordinates (i, j), then the coordinates of the four pixels are (i+2, j), (i, j+2), (i-2, j); and (i, j-2). Since the iteration order is from the top left corner to the bottom right, the two additional pixels on the left of and above P do not need to be considered separately.

Thinning will lead the input road segments to a skeleton of only one pixel wide (Figure 4). After the thinning algorithm was applied to the image that had been processed for noise removal and road segment joining, it was saved as a file.




Figure 4. Appearance of three representative clusters of spatially contiguous pixels before and after thinning.

4. RESULTS

After thinning, all roads have a uniform width of one pixel. Compared with the original image, the thinned result (Figure 5) is more interpretable. It contains nearly all major streets. However, the output result suffers from three main problems. First, some road segments shorter than the specified threshold still remain in the result (Figure 6). They should have been removed as noises, but were not. This is because after the threshold of removal was set, the length of the longest connected pixels was calculated as the length of the road. After the thinning process, some pixels had been removed to make the road uniformly wide. This may have caused the length of the road segment to be shorter than the threshold set in noise removal. On the other hand, some road segments that form a longer road have not been joined together. This is caused by the gap of larger than 4 pixels between them.

Apparently, this allowable gap between two segments is too conservatively defined in light of the obtained results. This threshold needs to be relaxed to further improve the quality of the mapped road networks.



Figure 5. Results of detected roads. This output result has been thinned with a threshold of 5 pixels in noise removal and a threshold of 4 pixels in road segment joining.

Second, some extracted roads show up as two parallel lines or even as a web (Figure 6). A comparison with the original image reveals that these roads are wide and have multiple lanes (e.g., major artery roads or motorways). Their extraction is not satisfactorily because of the improper assumption that all roads are one pixel wide. Under this assumption an operation window of 3×3 was applied. Within such a window size only eight neighboring pixels can be taken into consideration, resulting in the double line and web problem.



Figure 6. An enlarged portion of the extracted road network.

Road segments shorter than the removal threshold after thinning have not been removed. Roads wider than one pixel in the input image have some branches on the edge and appear as a web inside.

This problem could have been avoided using other algorithms and a varying operating window size. During the extraction the width of a road is detected first. If its width is larger than one pixel, then the operation window will be set larger accordingly. For example, if the road has four lanes, the operation window is increased to 5 by 5 ($2*2+1$) pixels. If the road has six lanes, the operation window is increased further to 7 by 7 ($2*3+1$) pixels, and so on. In order to accommodate a varying window size, dynamic templates need to be

developed. This topic is very complex and beyond the scope of this paper.

Third, a closer inspection of the result (Figure 5) reveals that some extracted roads are not so smooth. Some of them even contain branches (Figure 7). This problem is caused by the thinning algorithm which can thin most parts of a road except corners where small branches are found. Figure 7 shows the input rectangle in the left and its corresponding output in the right. Ideally, the output road should be a rectangular outline of one pixel wide. However, it contains two tiny branches in two of the four corners. The reason for this limitation is the order of iteration during which the search is carried out in the order from the first column of the top left corner down to the first column in the bottom left, then the second column from the top left to the bottom left, and so on until the bottom right. If the iteration order is changed from right to left, then the output road will just be a reversal of that shown in Figure 7. This problem can be overcome through a further process of filtering to remove the branches attached to a road.



Figure 7. A rectangular road block and its thinned output. The output contains two tiny branches because of the sequence of iteration. They will shift to the left if the order of iteration is reversed.

5. SUMMARY AND CONCLUSIONS

In this study we developed an automatic method of extracting road network information from hyperspatial resolution multispectral IKONOS data in a densely populated urban area. The introduction of spatial reasoning into the extraction is able to overcome the problems commonly associated with existent methods of road extraction, over which this proposed method based on spatial reasoning has a few advantages. The first is that it is highly flexible. No limit is imposed on the size of the area under study. There is no assumption about the analyst's familiarity with the research area. Second, it is comprehensive and covers the entire process of road network extraction, from unsupervised classification to create a binary image, to noise removal, road segment joining and road thinning. This method has been implemented in a flexible environment in which the user can specify some parameters during the extraction, such as the threshold for removing noise and the threshold for joining road segments together. Finally, the method can be applied readily with a wide range of image formats. Anyone who has access to an unsupervised classification package can perform the extraction. During undertaking of spatial reasoning the default file format is the commonly used bitmap, which is widely available and acceptable by most software packages. Apart from the IKONOS images, this method also works with other types of satellite imagery. Understandably, the finer the spatial resolution of the image, the better roads show up on the image, and the more accurate their extraction is. It must be pointed out, nevertheless, that all extracted roads are restricted to the uniform width of only one pixel.

References

Agouris, P., Stefanidis, A. and Gyftakis, S. 2001. Differential snakes for change detection in road segments. *Photogrammetric Engineering and Remote Sensing*, 67(12), pp. 1391-1399.

Barzohar, M. and Cooper, D. B. 1996. Automatic finding of main roads in aerial images by using geometric-stochastic models and estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(7), pp. 707 – 721.

Bhattacharya, U. and Parui, S. K., 1991. An improved backpropagation neural network for detection of road-like features in satellite imagery. *International Journal of Remote Sensing*, 18(16), pp. 3379-3394.

Choy, S.S.O., Choy, C. S.-T. and Siu, W.-C., 1995. New single-pass algorithm for parallel thinning, *Computer Vision and Image Understanding*, 62(1), pp. 69–77.

Fiset, R. and Cavayas, F. 1997. Automatic comparison of a topographic map with remotely sensed images in a map updating perspective: the road network case. *International Journal of Remote Sensing*, 18(4), pp. 991-1006.

Geman, D. and Jedynak, B., 1996. An active testing model for tracking roads in satellite images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(1), pp. 1–14.

Gruen, A. and Li, H., 1995. Road extraction from aerial and satellite images by dynamic programming, *ISPRS Journal of Photogrammetry and Remote Sensing*, 50(4), pp. 11-20.

Maillard, P. and Cavayas, F. 1989, Automatic map-guided extraction of roads from SPOT imagery for cartographic database updating. *International Journal of Remote Sensing*, 10(11), pp. 1775-1787.

Meisels, A. and Mintz, D., 1990. Symbolic reasoning in object extraction. *Computer Vision, Graphics, and Image Processing*, 52(3), pp. 447-459.

Nevatia, R. and Babu, K. R., 1980. Linear feature extraction and description. *Computer Graphics and Image Processing*, 13(3), pp. 257-269.

Steger, C., 1996. Extraction of curved lines from images. *Proceedings of the 13th International Conference on Pattern Recognition*, 25-29 Aug. 1996, vol. 2, p. 251-255.

Stilla, U., 1995. Map-aided structural analysis of aerial images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 50(4), pp. 3-10.

Tupin, F., Maitre, H., Mangin, J. F., Nicolas, J. M. and Peckersky, E., 1998. Detection of linear features in SAR images: application to road network extraction. *IEEE Transactions on Geoscience and Remote Sensing*, 36(2), pp. 434 – 453.

Wang, J. Treitz, P. M., and Howarth, P. J., 1992. Road network detection from SPOT imagery for updating geographical information systems in the rural-urban fringe. *International Journal of Geographical Information Systems*, 6(2), pp. 141-157.

Xiong, D., 2001. Automated road network extraction from high resolution images, National Consortia on Remote Sensing in Transportation: Safety, Hazards, and Disaster Assessment, Albuquerque, New Mexico, May 2001, pp. 1-4.