

THE UNIVERSITY OF HULL  
CARTOGRAPHIC INFORMATION SYSTEMS  
RESEARCH GROUP C.I.S.R.G.

DEPARTMENT OF COMPUTER SCIENCE · HULL HU6 7RX · UNITED KINGDOM  
TELEPHONE 0482 465295/465951 · FACSIMILE 0482 466666 · TELEX 9312134728 HUG

DR M VISVALINGAM  
HONORARY CO-ORDINATOR

Discussion Paper Series Editors: Dr M Visvalingam  
Professor M E Turner



C.I.S.R.G. DISCUSSION PAPER 12

Global dimensions for the recognition of  
prototypical urban roads in large-scale  
vector topographic maps

by

M Visvalingam, D A Varley  
and C P Wright

## Discussion Papers

1. Wade, P., Visvalingam, M. and Kirby, G. H. (1986)  
**From Line Geometry to Area Topology**, 48 pp.
2. Visvalingam, M. (1987)  
**Problems in the Design and Implementation of a GKS-based User Interface for a Graphical Information System**, 18 pp.
3. Visvalingam, M. (1988)  
**Standardising Basic Spatial Units: Problems and Prospects**, 19 pp.
4. Whyatt, J. D. and Wade, P. R. (1988)  
**The Douglas-Peucker Algorithm for Line Simplification: an introduction with programs**, 17 pp.
5. Visvalingam, M. (1988)  
**Cartography, Geographical Information Systems and Maps in Perspective**, 16 pp.
6. Visvalingam, M. and Whyatt, J. D. (1989)  
**The Douglas-Peucker Algorithm for Line Simplification: re-evaluation through visualisation**, 35 pp.
7. Visvalingam, M. (1989)  
**Trends and Concerns in Digital Cartography**, 37 pp.
8. Visvalingam, M. and Whyatt, J. D. (1990)  
**Cartographic Algorithms : problems of implementation and evaluation and the impact of digitising errors**, 24 pp.
9. Visvalingam, M. (1991)  
**Visualisation, VISC and Scientific Insight**, 21 pp.
10. Visvalingam, M. and Whyatt, J. D. (1992)  
**Line Generalisation by Repeated Elimination of the Smallest Area**, 16 pp.
11. Varley, D. A. and Visvalingam, M. (1993)  
**Area Topology for Road Extraction and Topographic Data Validation**, 34 pp.
12. Visvalingam, M., Varley, D. A., and Wright, C. P. (1993)  
**Global dimensions for the recognition of prototypical urban roads in large-scale vector topographic maps**, 20 pp.

## Special Issues

1. Visvalingam, M. (ed) (1986)  
Research based on Ordnance Survey Small-Scales Digital Data Proceedings of a meeting held as part of the 1986 Annual Conference of the Institute of British Geographers (Reading Jan 1986), **sponsored by the Ordnance Survey**, 79 pp.
2. Visvalingam, M. and Kirby, G. H. (1987)  
Director of Research and Development based on Ordnance Survey Small Scales Digital Data, **sponsored by the Ordnance Survey**, 38 pp.
3. Visvalingam, M. and Sekouris, N. M. (1989)  
Management of Digital Map Data Using a Relational Database Model, **sponsored by the Ordnance Survey**, 50 pp.

# Contents

Abstract	
1 Introduction	1
2 Background	2
2.1 Literature on Road Recognition .....	2
2.2 The Data .....	4
3 Research Strategy	5
3.1 Prototypical Roads .....	5
3.2 Global Dimensions .....	6
4 Observations	8
4.1 Key Dimensions .....	8
4.2 Selection of Cut-Off Values .....	10
4.3 Verification .....	12
5 Conclusion	17
Acknowledgements	18
References	19

## Abstract

This paper outlines a simple strategy for recognising prototypical cases of urban roads given only their geometric forms. Two new indicators, namely *exits* and *occupancy*, are used together with *average width* to label regions as roads. Even in its present simplistic form, the methodology proposed here proved to be useful for identifying prototypical cases. It is equally useful for flagging some roads as untypical. These were labelled as roads by a previous study on road extraction which exploited available semantic information. However, owing to the absence of logical links, some roads became combined with neighbouring regions making them untypical. The recognition strategy therefore provides an additional means for validating topographic data.

## 1 Introduction

This paper arose out of a SERC CASE project (Jan 1990–Dec 1992) undertaken in collaboration with the Ordnance Survey of Great Britain. It is concerned with the identification and labelling of areal objects given feature-coded vector topographic data. Geographic objects may be identified in a variety of ways. Varley and Visvalingam [1] used the method of extraction, which uses available semantic data. They distinguished extraction from recognition; the latter is based only on the outline forms and juxtapositions of regions of uncut space. Their study on road extraction used semantic codes associated with linear features, such as road metalling links, and road centre lines to extract roads and validate the input data. The extracted roads were then used to detect inconsistencies in the data and violations of the data specification. However, the process of extraction is affected by deficiencies and errors in the data. The aim of this study was to establish whether prototypical urban roads had distinctive geometric properties which could be used to provide an independent means for checking the results of extraction.

If we were to plot only the lines on 1:1250 Ordnance Survey maps in a single line-style in black-and-white and omit all the area symbolism, such as for roofed areas, map users would still be able to recognise and distinguish various classes of objects from their shapes and context alone. Although no road on an urban map is exactly the same as another, we would be able to correctly identify most of the roads. Yet, the automatic recognition of urban roads on large-scale maps has been a continuing topic of research. The identification and appropriate definition of roads is important to many urban Geographical Information Systems.

In the next, background, section we briefly review relevant research in this area. We then describe the ideas underpinning our recognition strategy. Previous studies have attempted to abstract a single set of rules for recognising all roads. Our experience with road extraction [1], indicated that even extraction requires quite different strategies for different categories of roads. Rosch's work [2] on categorisation and cognition encouraged us to concentrate initially on the prototypical cases. In edge-matched databases, these prototypical cases may be used to identify further roads by extension of these roads across map sheets, without reference to semantic information. Further study of the prototypical cases may in turn lead to the abstraction of patterns which identify less distinctive cases. We were interested in establishing whether prototypical cases could be identified directly using easily computed metrics.

This paper demonstrates that prototypical roads can be immediately and directly recognised using three easily calculated metrics, two of which are new to the literature. The ideas were tested using large-scale data for three urban environments but it needs further testing. The results suggest that a single pseudo-dimension may be sufficient for recognising prototypical cases. The other two indicators are useful for detecting those cases of extracted roads which deviate from the mental prototype in some aspects. It also

suggests that these metrics may be used within more complex rules for recognising less typical cases of roads.

## 2 Background

### 2.1 Literature on Road Recognition

The term, recognition, has been used quite loosely to describe a variety of quite dissimilar approaches to object identification. This is inevitable given that some studies, such as that by Marr [3] and others in Humphreys [4] are more concerned with understanding human vision while other studies have been oriented instead towards solving practical problems. There are a number of factors which have led to a variety of approaches. These include:

- **The nature of the phenomena itself.** As discussed later, even roads vary in character and require different strategies for recognition.
- **The nature of the input data.** The majority of studies focus on the recognition of roads on images [5]. McKeown et al [6] and Quam [7] have reviewed and contributed algorithms for tracking inter-urban highways on medium-scale images of scenes. Others have been more concerned with identifying symbolised roads on images of small and medium-scale maps.

The present study is based on vector data. Vector databases also vary in specification. It is not possible therefore to use the methods proposed in this paper with unstructured data, in which line intersections are graphically implied rather than mathematically defined.

- **Scale of mapping.** Small-scale maps only show a schematic network of road centre lines. Zhu and Kim [8], for example, were able to identify these using a variety of techniques including connected component labelling. Medium-scale maps, such as those used by Nagao et al [9] and Suzuki et al [10], show roads as having roughly parallel sides but not necessarily as closed regions. In large-scale maps such as those used in this study, roads are defined by their boundaries. The boundaries of roads describe a variety of sub-features, such as lay-bys, roundabouts, entrances to drives and bus stops. These maps also tend to be dense and detailed.

Thus techniques for road recognition are not universally applicable. In this study we have attempted to formulate a strategy which is more dependent on the general nature, rather than precise form, of urban roads.

The term, map scale, is used here in a relative sense since national mapping is undertaken at a variety of scales. The basic scale of urban mapping undertaken by the Ordnance Survey of Great Britain is 1:1250. These 1:1250 large-scale urban maps are black and white and use only two types of lines, solid and pecked. Since these large-scale maps portray the outlines of real world entities, map readers are able to recognise many objects from their

forms and context alone. This semantic information is recorded by feature codes associated with vector-digitised points and lines. The process of road extraction, described by Varley and Visvalingam, used fragmentary feature codes to identify roads. The present study on road recognition attempts to perceive roads in the way a map reader would given Figure 1.

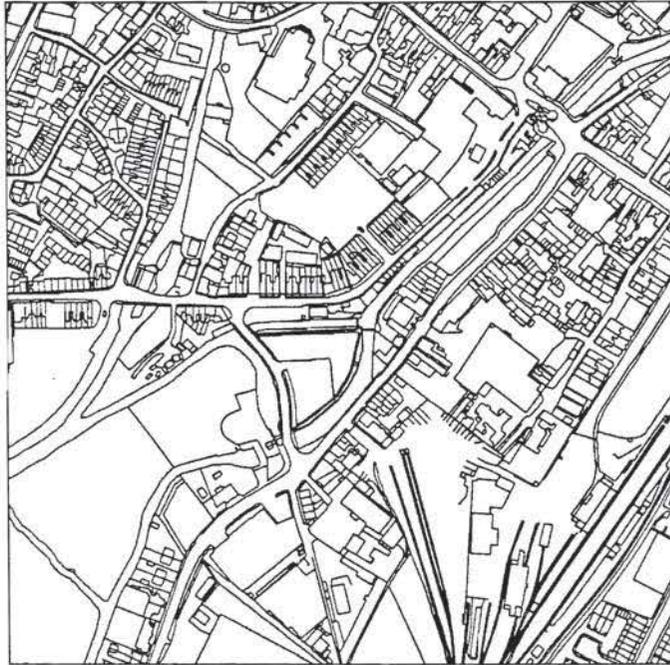


Figure 1: An example of a 1:1250 topographic map produced by the Ordnance Survey of Great Britain (Data: Crown Copyright).

Only two other British studies have been based on data digitised from the OS 1:1250 maps although their digital data conformed to different specifications. Of these, only the study by De Simone [11] is noteworthy; it provided the inspiration for this work. De Simone devised an elaborate staged strategy for recognising objects in three groups of topographic entities, namely railways, roads and land parcels. Each group was characterised by a complex of objects which he called superstructures. Since elements of railway superstructures have the same patterns as those in road complexes, his strategy required the recognition and elimination of railway superstructures prior to road recognition; rails are easily distinguished by their parallel configurations and narrow gauges.

The full description of De Simone's rather complicated and sometimes ad-hoc recognition strategy is outside the scope of this paper. Having eliminated the railways, the regions which were "long, with fairly uniform width, parallel edges and a high degree of straightness" were selected as potential candidates (pp. 141–142). The selected regions may contain objects other than roads. So the context of the road candidate was taken into account. The rules for this were not clearly explained. Our interpretation of them

is as follows. Adjacent shapes, which looked like pavements (narrow, long and thin with parallel sides) and which had 40% or more adjacency with the road were examined. If the sum total of adjacent 'possible pavements' exceeded 30% of the road length, then the shape was labelled as road and the adjacent shapes which lent context were labelled as pavements/verges. The entire process involved a great deal of computation, numerous parameters and cut-off values. Roads without the anticipated context were not identified. A couple of others without pavements (p. 163) were adjoined by narrow front gardens, without dividing fences. These were reckoned as pavements even though they do not look like them. Consequently, the roads were recognised but part of the adjoining land parcels were erroneously classified as pavement; the adjoining houses and their land parcels were not recognised as a result.

The thesis does not refer to these problem cases nor does it indicate how roads without the anticipated context would be recognised. We rejected De Simone's approach for three reasons. Firstly, De Simone edited his maps to ensure that the data conformed to expectations. Given the enormous size of national databases, it is unrealistic to assume perfect data. Our aim was to develop methodology for validating data. Secondly context, as used by De Simone, is not reliable especially with imperfect data. Finally, De Simone's approach would have been difficult to program and verify and appears to be unnecessarily complicated. Map users do not rely on the prior recognition of railway objects when recognising roads.

## 2.2 The Data

Our initial study was based on the same OSBASE data used by Varley and Visvalingam [1], who describe the dataset in some detail. The data were for two contrasting urban areas, namely Birmingham and Canterbury. The data for Canterbury was a more recent product. OSBASE was a prototype database which was created by the Ordnance Survey (OS) of Great Britain for its own internal experimentation and for market research. This database is now obsolete and has been superseded in turn by other experimental products. The recognition strategy was also tested on a recent experimental database for Ashford (Kent), which conforms to a more recent object-based prototype called Project 93.

OSBASE is a vector database consisting of point, line and text features. Varley and Visvalingam explained how the lines were re-organised into the data structures which articulate the Disassociative Area Model (DAM) [12]. DAM is a topological model which makes explicit the geometric topology of lines and areas. Varley and Visvalingam used topological concepts to formulate reliable and efficient algorithms for extracting roads using available semantic clues, and then for validating and extending the semantic coding. This study is based on their classification of primitive regions into roads

and other types of objects. This catalogue served two functions. It enabled us to recognise patterns in the data and formulate recognition strategies and rules. It also allowed us to locate errors of commission and omission for further study.

### 3 Research Strategy

Figure 1 shows a sample map. Although this map is devoid of text and symbolism, map users can still pick out the major roads on the map from their forms and juxtaposition vis-a-vis other regions. In this section, we consider how these prototypical urban roads may be described in qualitative and quantitative terms.

#### 3.1 Prototypical Roads

The principles underpinning categorisation and classification have been explored by psychologists, linguists and anthropologists in Rosch and Lloyd [13]; some of these ideas continue to influence the literature on Visual Cognition [14]. A category is a set of objects which are equivalent in some respects. A taxonomy is a system by which categories are related to one another by class inclusion. Categories and their members are distinguished from other categories by a set of defining characteristics.

Rosch [2] suggested that basic categories are the categories that best mirror the correlational structure of the environment and that objects are first seen and recognised as members of their basic category and that only with the aid of additional processing can they be identified as members of a super- or sub-class. The computational models for road recognition normally place urban roads at a basic level of the taxonomy as distinct from others, such as buildings. We differentiate basic objects by their perceptual and functional attributes. At a higher level, roads may belong to a superclass of linear networks which would include other basic objects, such as rivers. Superclasses generally share fewer attributes. Thus, De Simone was able to differentiate between railways and roads according to their width. At a lower level, we can sub-categorise roads on several criteria. Motorways, freeways and rural roads differ from urban roads. Equally, roads in newer planned settlements and suburbs have a different profile compared with those in the core of historic towns, such as Canterbury.

Although objects are often seen as distinctly different in their geographical context, they are not necessarily discontinuous or clustered in statistical space. Since the categories in the taxonomy seldom have clear-cut boundaries in this property space, it is difficult to recognise them by a formal set of rules. Rosch [2] noted that we appear to circumvent this problem by thinking of each category in terms of its clear, or prototypical, cases rather than in terms of its statistical boundaries. People tend to easily agree on whether a case belongs to a category even if there is some disagreement over the precise location of these boundaries.

Rosch's experimental studies also indicated that learning, recognition, recall and classification were more easily accomplished through the use of prototypical instances of classes. Past research on road recognition has not sought to distinguish between prototypical and less typical cases but have attempted to formulate a single set of rules to identify all cases of urban roads. Varley and Visvalingam found it necessary to distinguish between trivial and more tricky cases when formulating algorithms for road extraction. Even map users can recognise only some roads instantly; a greater degree of cognitive effort is needed to resolve others. It was quite clear that some of the problems facing extraction would also affect recognition. We therefore decided to focus initially on prototypical cases of roads.

Before we can recognise prototypical cases, we need to specify what we mean by them. When we think of urban roads we think of arterial branched networks. No doubt, there are many other networked objects since branching is a quality of the superclass of linear networked objects. However, roads are recent superimpositions on the urban scene and they often cross over waterways and railways. Roads therefore tend to segment other objects and form the most extensive networks connecting intra- and inter-urban spaces. Roads also penetrate urban space and extend beyond the map sheet.

The mental models we use for holistic recognition of urban roads would be different to those we would adopt for recognising rural roads or for incremental tracking of motorways. This suggests that urban roads form a separate basic category. The Department of Transport classification of roads usually applies to component parts of a connected network and additional processing is needed to identify these sub-categories. Within the basic category, the difficult cases would tend to be those which do not match our mental model of urban roads, as seen later.

### 3.2 Global Dimensions

The next problem was to determine the characteristic properties of prototypical cases. We pursued some context-based analysis initially but found it difficult to recognise even fairly prominent roads. This was partly because the data specification does not require that all objects are defined by closed polygons. Consequently, several objects—including roads, pavements, car parks and fields—can become amalgamated with some of their neighbours. De Simone manually edited his maps to separate these features. It must be accepted that mass digitised data is likely to exhibit this problem since the boundaries between these objects are often conceptual rather than physical and may not be indicated on the visual map. Manual digitisation of these logical links is unreliable. Moreover, if the input to the recognition process consists of auto-vectorised scanned images, they are likely to be missing. We were keen to establish whether the main arterial roads could be recognised despite the absence of these logical links. The long-term intention was to devise procedures for automatically separating these objects and truncating minor features on roads as part of the research on map generalisation and scale-free mapping [15].

Since contextual information could not be determined reliably at this stage of the analysis, we decided to explore the attributes of objects instead. Garner [16] made a distinction between features and dimensions. He used the term, feature, to refer to some distinguishing component of an object which may be either present or absent. For example, the presence or absence of a single line can serve to distinguish an O from a Q. Garner proposed that the visual system prefers feature descriptions in many information processing tasks. Suzuki et al [10] and De Simone [11] incorporated features, such as right-angled corners and parallel sides, in their recognition systems.

However, the identification of in-line features is computationally demanding. In his formal definition of terms, Garner differentiated between these optional features and dimensions. He defined dimensions as potentially variable properties of components which are always present. Objects also have global dimensions. Garner was aware that aspects of stimuli could be studied using either dimensions or features. For example, the letters E and F are distinguished both by the dimension, number of horizontal lines, and by the presence or absence of the lowest horizontal line. Some features can be tracked more easily through such pseudo-dimensions.

The choice of the term, dimensions, to denote these intrinsic aspects of an object is somewhat confusing since the word dimension has specific meanings which are different in physics and mathematics. We used the word to imply the phrase, *global dimension*, which refers to the intrinsic properties of the object. The study reported in this paper was concerned with assessing whether easily computable global dimensions enable direct recognition of prototypical urban roads and reveal the defining aspects or essence of these roads.

Since large-scale maps depict roads by their boundaries rather than by stylised conventions, a single dimension, such as width at the map edge, is inherently unreliable. In his textual analysis, De Simone noted that roads were extensive but his recognition strategy did not quantify this adequately. He therefore had to use contextual information, such as adjacency, full and partial containment, line continuation, alignment of shapes and shape combination.

The problem in recognition is one of identifying and quantifying the definitive aspects or the essence of the class. We investigated the following dimensions.

- *area*, of the region
- *perimeter*, of the region
- *shape* -  $\frac{\text{perimeter}^2}{\text{area}}$  - a measure of elongation.
- *average width* -  $\frac{2 \times \text{area}}{\text{perimeter}}$  - provides a crude indicator
- *exits*: the number of exits may be regarded as a pseudo-dimension since it indicates the presence of branches. Thus a high number of exits is indicative of branching.

- *envelope*; as a crude indicator of extent
- *occupancy* -  $\frac{\text{area}}{\text{envelope}}$  - again this is only a crude indicator of branching. It also captures the fact that roads feed but do not flood urban spaces.

The first four dimensions were also used by De Simone; whereas De Simone attempted to measure these dimensions as accurately as possible, we were only concerned with quantifying them as crude indicators. We introduced the last three indicators; *exits* and *occupancy*, in particular, were designed to capture the intrinsic nature, rather than precise form, of roads. Except where a section of a road becomes truncated by overhead features, such as bridges, roads by their very nature will eventually connect with those on adjacent sheets at the map edge. The feasibility study on road extraction initially selected only those regions with at least one map edge link as candidate regions since urban roads tend to extend beyond a single map sheet. Indeed, in the 35 sheets we have studied to date, there were only 2 detached sections of road. This led to the insight that arterial roads have a number of *exits* (or intersections with an imposed window) and that this could be used as one indicator. As a corollary, an arterial urban road would tend to extend over a large part of the map. To test this we included the *envelope* or bounding box of the region as another indicator. This is not entirely reliable. Since urban roads will branch out in several directions and will segment other objects, it was worth considering. Although it proved to be unreliable on its own, it led to the inclusion of a third dimension, *occupancy*. With an extensive and branched object, as gauged by *exits*, *occupancy* is indicative of the netted character of roads.

Note that all of these dimensions are easily gathered. Area, perimeter and envelope feature in the data structures of many GIS since they are used by many applications. The number of exits is also deduced by reference to the outermost enclosing hole in the left/right boundary fields in the link records [12]. So the co-ordinate files need not be accessed from disc to compute these dimensions.

## 4 Observations

In this section we describe the strategy for visual analysis of statistics which was used to identify the key dimensions of prototypical roads and their cut-off values.

### 4.1 Key Dimensions

There are 5277 regions with at least one *exit*. Of these, 100 are roads (see Table 1). Figure 2 only shows the 183 regions (3% of all regions) with three or more *exits*; these include 44 roads. This selection was justified on the grounds that arterial roads are likely to have more than 2 *exits* in a 500m x 500m area. Three *exits* is also a definitive proof of the existence of at least one branch. Furthermore, these 44 road regions account for over 93% of the

Number of Exits	Roads	Neighbours	Others	All
More than 9	20	1	0	21
3-9	24	83	55	162
Less than 3	56	833	4200	5059
Total	100	917	4255	5272

Table 1: Counts of regions with at least one exit.

Roads (44 Regions)				
Dimension	Mean	Std. Dev.	Min	Max
Area (m <sup>2</sup> )	16229	12962	397	49678
Perimeter (m)	4806	3568	384	12063
Shape	1485	1086	139	3773
Width (m)	6.45	1.50	2.08	8.84
Exits	9.14	5.54	3	23
Envelope (m <sup>2</sup> )	147075	98851	3765	250000
Occupancy (%)	11.9	5.2	5.3	26.6

Neighbours (84 Regions)				
Dimension	Mean	Std. Dev.	Min	Max
Area (m <sup>2</sup> )	2109	4100	25	25989
Perimeter (m)	1066	1290	45	7063
Shape	712	751	19	3692
Width (m)	3.41	3.25	1.08	29.98
Exits	3.98	2.33	3	23
Envelope (m <sup>2</sup> )	21314	31835	49	146483
Occupancy (%)	15.9	13.9	2.7	62.9

Others (55 Regions)				
Dimension	Mean	Std. Dev.	Min	Max
Area (m <sup>2</sup> )	3514	8712	3	41245
Perimeter (m)	685	842	14	4657
Shape	420	703	24	3465
Width (m)	7.31	11.62	0.48	70.74
Exits	3.85	1.37	3	8
Envelope (m <sup>2</sup> )	17769	32618	7	134269
Occupancy (%)	31.0	19.8	1.7	80.8

Table 2: Summary statistics for regions with more than 2 exits.

area of all roads. The bulk of the remaining cases with one or two exits do not exemplify a mental model of urban roads. They are segments resulting from the superimposition of a viewing window on the urban scene, in this case the map sheet boundary; most of these are very small segments which just intrude into the map. The road regions with zero exits are those which have become detached by overhead features.

Figure 2 shows the univariate distributions of 6 of the global dimensions for roads, their neighbours and other objects as labelled by Varley and Visvalingam [1]. In each diagram, the regions in that class are sorted on the depicted dimension. The point symbols record the values for the regions arranged in ascending numeric order along the x-axis. The same y-scale is used to facilitate comparisons of values across the graphs for all categories. These plots show that *area* and *shape* (Figures 2a and 2b) are the least useful for constraining the search. Figures 2c and 2d indicate that *average width* and *occupancy* are useful for eliminating unlikely candidates. Although both *average width* and *shape* are based on area and perimeter, *average width* is a better measure which also has the advantage of having a physical meaning relating to common knowledge. *Perimeter*, *envelope* and *exits* have similar univariate distributions and are all useful for distinguishing the larger roads. *Perimeter* is less helpful since it is difficult to select statistical boundaries on a priori grounds. An *envelope* of 250000 metres sq, covering the map sheet, is particularly distinctive since roads are likely to segment other objects (Figure 2e). However, with one notable error of commission, *exits* is equally discriminating (Figure 2f). Since *envelope* is only a crude indicator, *exits* was favoured as a more meaningful and reliable dimension for focusing on the most distinctive cases.

In summary, *exits* appears to be the best indicator of networked, arterial roads. *Average width* and *occupancy* emerge as useful dimensions for constraining the search for other prototypical roads.

## 4.2 Selection of Cut-off Values

We then considered how the selection of cut-off values could become independent of human judgement. Both visual and statistical summaries (see Table 2) suggest that some roads are very distinctive and that they may be recognised using the mean value of 9 *exits* as the cut-off value. Since the statistical distribution is skewed, this is likely to omit other prototypical cases but this does not matter at this stage. It is also likely to incur errors of commission although there was only one rogue case in our sheets, namely a railway shunting yard. The reason why this railway has so many *exits* is because railway detail is not link-and-node structured except at the map edge. Consequently a number of adjacent regions have become amalgamated. Since roads segment other large objects, the latter must be limited in their extent resulting in higher *occupancy*. Roads with more than 2 exits only occupy between 5 and 27% of their *envelopes*; the shunting yard has a value in excess of 40%.

Although *average widths* of roads are constrained by function and land

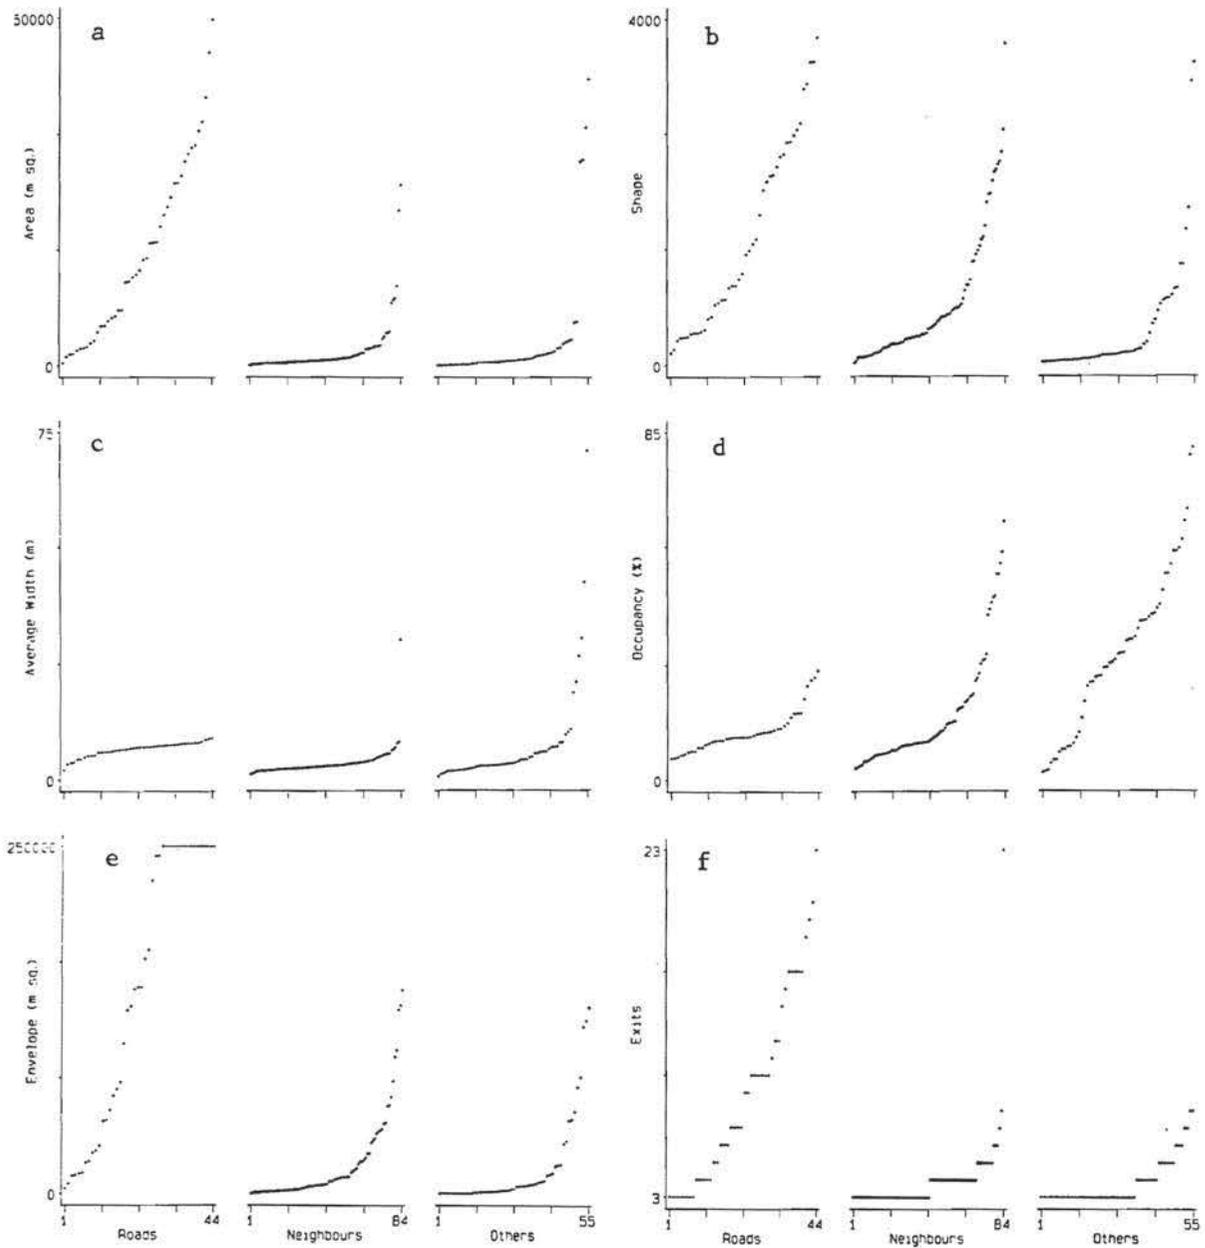


Figure 2: Univariate plots of a) area b) shape c) average width d) occupancy e) envelope f) exits.

value to lie within a restricted range, specific values can differ from one urban environment to another. For example, the minimum *average width* for a road with more than 2 *exits* is only 2 metres. This is not surprising since Canterbury is a historic town. The formula also under-rates the *average width*, particularly for more compact regions. Also, *average width* is sensitive to leaks. Thus, both *average width* and *occupancy* are environment- and data-dependent. It is therefore difficult to guess cut-off values but guesstimates may be available from transport authorities. However, the fact that these indicators are sensitive to data configurations, makes them suitable for picking out outliers.

Figure 3a shows all regions with more than nine *exits* except for the shunting yard (which falls outside the range of values shown). All 20 regions are roads by virtue of the fact that they include road metalling links. The roads which fall in the top left quadrant of Figure 3a conflate with our expectations of roads. All cases which fall outside this quadrant do not do so. If we assumed that prototypical arterial roads are unlikely to have occupancy of more than mean + 3(std dev), then objects with values in excess of 27% (well outside the range for roads) are unlikely to be typical. This is sufficient for eliminating the shunting yard. Figure 3a shows that all except 3 roads occupy much less than 20% of their *envelope*. A plot of these exceptional roads showed that they had merged with other entities (Figures 4a and 4b).

Figure 3a also shows that there are some roads with very narrow *average widths* for extensive roads. The cut-off of 5.6 metres is arbitrary but it illustrates that the *average width* dimension may also be used to locate other types of composite objects. For example, the light grey region in Figure 5a is a road with an *average width* of only 4.3m and it is quite clear that it has merged with pavements and paths; the same is true of the road in Figure 5b. Unusually high values for *occupancy* appear to be due either to structuring errors, giving rise to convoluted regions, or to roads which have merged with fields and car parks. Low *average widths* indicate merging with adjoining pavements, back lanes and paths in the urban-rural fringe. Some of these merged regions can be automatically detected using simple topological rules. Given the data structures in DAM [12], a link which has the same boundary on either side, flags an error [1]. However, this will only identify some inconsistencies in the data. The above checks on prototypical attributes help to locate other cases.

### 4.3 Verification

Although successful, roads with more than 9 exits account for less than 50% of roads with more than 2 *exits* (20% of all roads). We therefore investigated the utility of these three dimensions for classifying regions with 3 to 9 *exits*. As the number of *exits* decreases, it ceases to be a defining variable and becomes a mere filter. All the roads recognised as roads (top left quadrant in Figure 3b) appear to be prototypical cases. Only one other networked object,

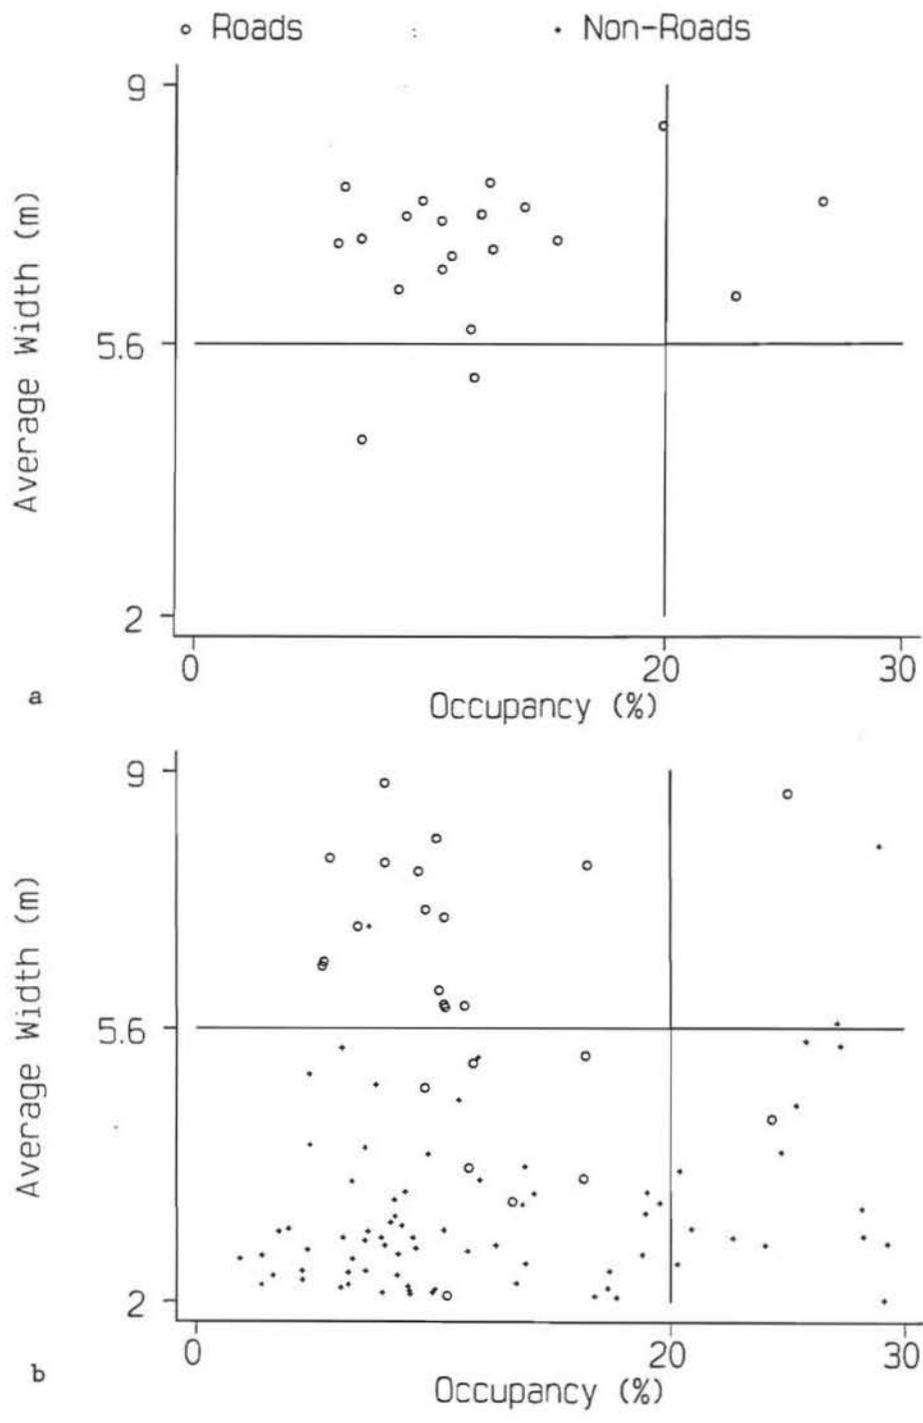


Figure 3: Bivariate plots of a) more than mean number of *exits*  
 b) between three and mean number of *exits*.

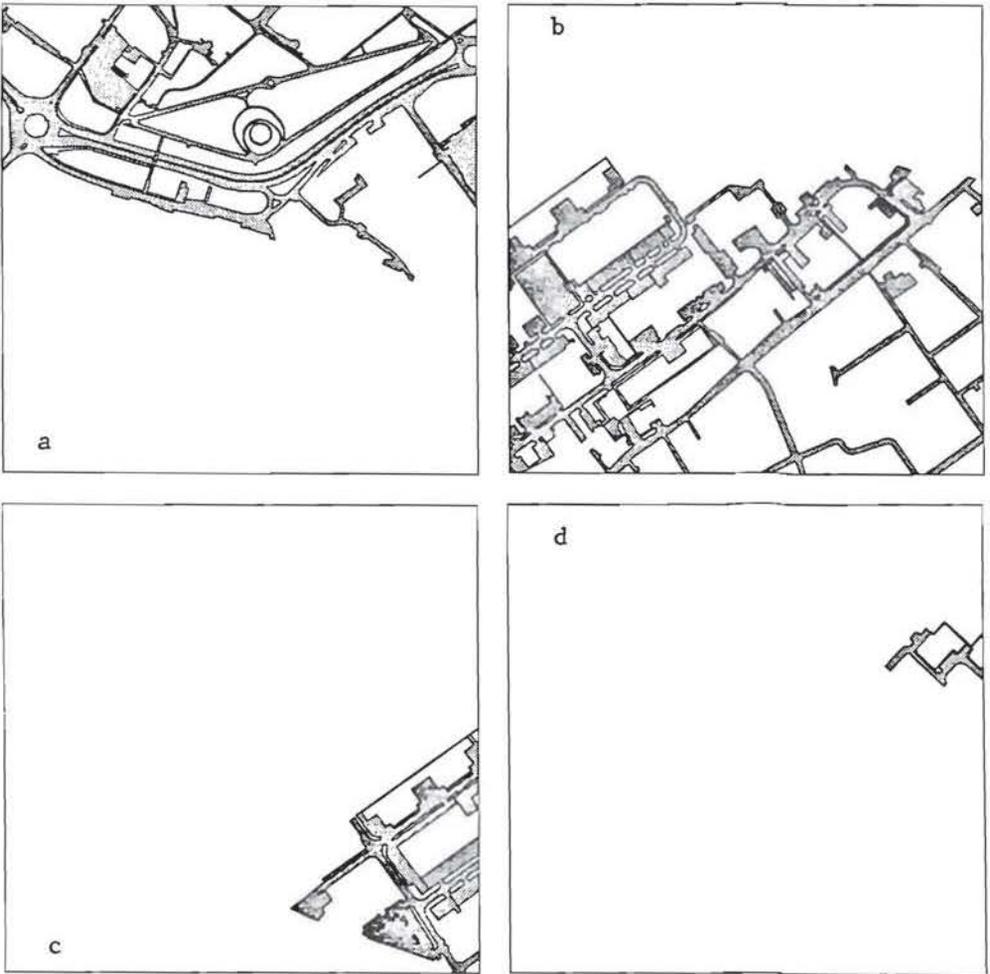


Figure 4: Roads with high *occupancy*.

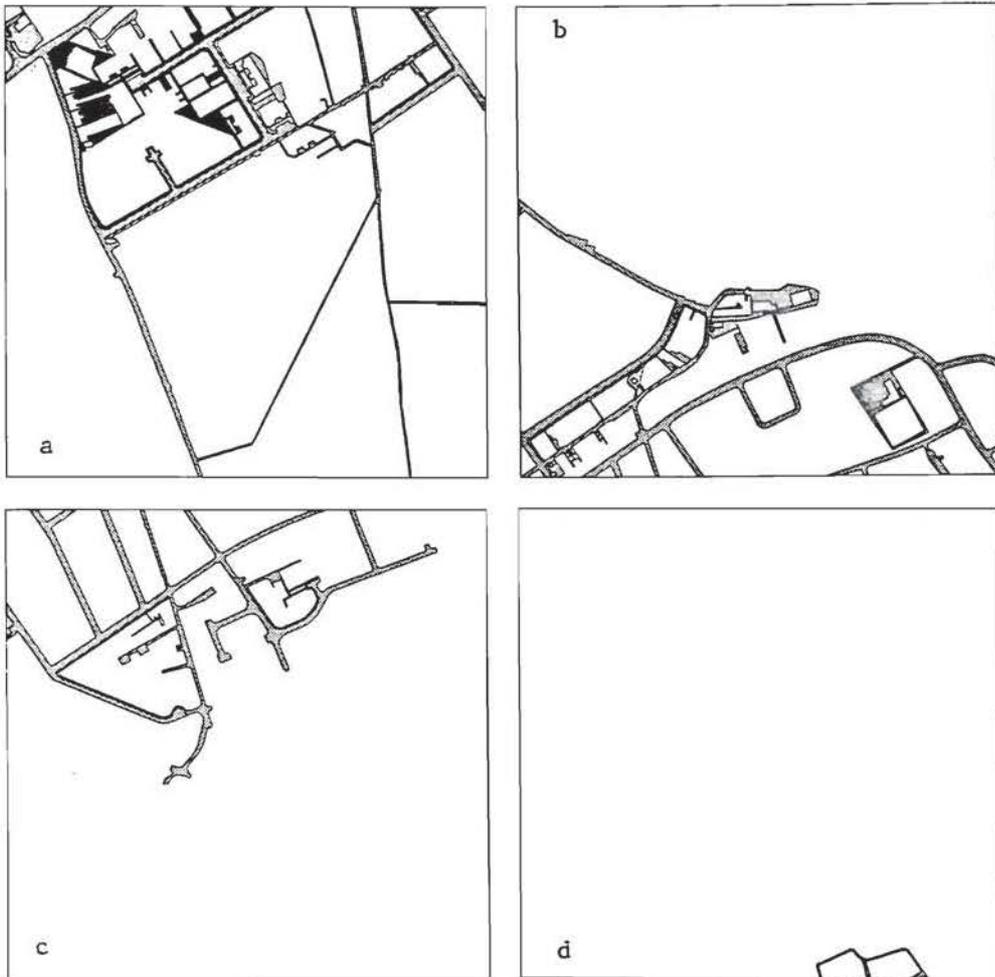


Figure 5: Examples of roads with low *average width*.

a river with 3 *exits*, is mislabelled as a road. We can tell from Figure 6 that it is not a road because it has sections which are narrower, and banks which are not as smooth as roads. Both characteristics can be discerned but only with more processing and this is outside the scope of this study. However, this case indicates that a small number of exits may pick the superclass of networked objects rather than the arterial urban roads alone.

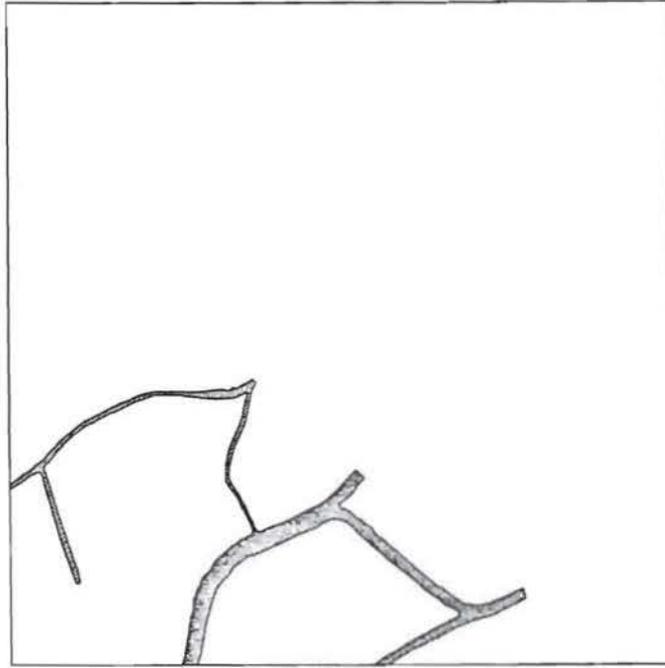


Figure 6: The only error of commission—a river.

The roads falling outside the top left quadrant in Figure 3b may be classified as follows:

- Roads which have combined with broad, spacious regions such as car parks or fields (Figures 4c and 4d); These appear in the quadrants on the right hand side. The case shown in Figure 4d with high *occupancy* and low *average width* is particularly interesting since it is a road which has three *exits*, only because it has leaked.
- Roads which have become merged with narrow regions such as pavements or back alleys (these regions appear in the top half of the lower left quadrant of Figure 3b). Unlike the regions in the next category, these are mainly roads whose global dimensions have become distorted by inclusion of other objects (for example, see Figure 5c).
- Pavements, paths and back alleys which have been labelled as roads by the extraction process, only because they have included small segments of road that just intrude into the map-sheet. These regions tend to have very low *average width*, and appear in the bottom half of the

lower-left quadrant (for example, see Figure 5d and the dark grey region in Figure 5a). The consistency checking undertaken by Varley and Visvalingam would have noted some inconsistencies in the data but the recognition process is able to infer that these polygons look more like pavements and paths than roads.

Again there seems to be a pattern: inclusion of objects such as car parks gives rise to high *occupancy*, while the inclusion of narrow regions results in low *average width*. Future work can attempt to establish whether the patterns for untypical roads, including those with less than 3 *exits*, can be resolved automatically and efficiently but in a simple and meaningful manner.

These three key dimensions were also applied to a recent experimental database, consisting of 7 map sheets in the Ashford area. The data specification for this prototype required the inclusion of logical boundaries. There were 13 roads in all; only 6 of these had 3 or more *exits*. The 3 dimensions were sufficient to discriminate unambiguously between these 6 roads and other objects in this subset.

## 5 Conclusion

This paper has made a number of contributions. First, it has conceptualised the form of urban roads in terms of their context and functions. Their transport and access functions specify an 'arterial' networked form. In the current state of transport development, roads appear to be imposed on other urban objects to link and connect urban spaces in an extensive way. Roads serve urban spaces without swamping them.

Second, this paper has shown that given topologically structured data, it is possible to recognise prototypical urban roads simply and efficiently through use of novel global dimensions. The conception of dimensions involved a great deal of mental visualisation aided by data visualization with the latter serving the former. The key characteristics which seem to typify urban roads are extent, netness, low occupancy and width of roads.

Third, the paper has shown that these characteristics may be quantified as three global dimensions of regions, namely *average width*, *exits* and *occupancy*. Two of these indicators are new to the literature on road recognition. They are *exits* and *occupancy*. *Exits* is a pseudo-dimension since its main function is to detect the presence or absence of features, such as branches and connections with external spaces. They are both dimensionless variables. Unlike *average width*, they should be applicable in a wide range of urban environments since they are indicative of extensive networks, rather than measures of the precise size and form of roads. These dimensions are easily calculated using data normally held in attribute tables by GIS without accessing co-ordinates on disc.

Fourth, the paper has suggested that *exits* alone is sufficient for identifying extensive arterial roads with more than 9 *exits*. The other dimensions

served to pick out deficiencies in data or roads which had become combined with neighbouring regions. All three dimensions have to be used to distinguish between roads and other objects with three to nine *exits*.

Fifth, it has shown that the method is sensitive to data conditions and that it is capable of picking out untypical cases for investigation. With one exception, the untypical cases of extracted roads are either not roads or are amalgams of objects. The method is thus useful both for flagging the absence of logical links in road boundaries and for locating some residual structural and semantic errors. This method of recognition is more reliable than extraction. It correctly rejected objects, such as pavements, which had been extracted as roads since they contained erroneous road metalling links. This suggests that the process of road extraction must be even more intelligent than that described by Varley and Visvalingam [1].

Finally, it has demonstrated the value of visualization showing how visual and statistical summaries may be used in conjunction with maps to assist in the processes of ideation and verification as suggested by Muehrcke [17].

In conclusion, we would like to stress the following. The recognition of prototypical roads forms only part of a programme of study on spatial data models and algorithms for topographic data and just one of the methods used for identifying roads and validating data. The ideas presented have only been tested with link-and-node structured data for British urban areas as modelled on Ordnance Survey 1:1250 maps. It must also be remembered that two of the three dimensions, namely *exits and occupancy*, are both dependent on the dimensions of the viewing window. The map sheet boundary provides a convenient window. It can be varied to enhance recognition but it should not be so small that it no longer measures the extensive, arterial nature of urban road nets. Neither should it be so large that it misses the urban environment and intersects inter-urban highways instead. Equally, rectilinear windows may not be optimal for detecting grid patterned roads.

## 6 Acknowledgements

We would like to thank the UK Science and Engineering Council for award of a CASE studentship to Dominic Varley and a quota award to Chris Wright. We are also grateful to the Ordnance Survey of Great Britain (OS), the collaborating body on the SERC CASE project, for providing access to their digital topographic data. We are particularly grateful to John Farrow, formerly of the OS for his support and encouragement.

## References

- [1] D.A.Varley and M.Visvalingam, Road Extraction and Topographic Data Validation Using Area Topology, *Computer J.*, **37(1)**, 1994, pp. 3-15.
- [2] E.Rosch, Principles of Categorization, in *Cognition and Categorization*, (E.Rosch and B.B.Lloyd Eds.), pp. 27-48, Lawrence Erlbaum Associates, 1978.
- [3] D.Marr, *Vision*, W.H.Freeman and Company, 1982.
- [4] G.W.Humphreys (Ed.), *Understanding Vision: An Interdisciplinary Perspective*, Blackwell, 1992.
- [5] P.Suetens, P.Fua and A.J.Hanson, Computational Strategies for Object Recognition, *ACM Computing Surveys*, **24(1)**, 1992, pp. 5-61.
- [6] D.H.McKeown and J.L.Denlinger, Cooperative Methods for Road Tracking in Aerial Imagery, in *Proceedings, IEEE Conf. CVPR '88*, 1988, pp. 662-672.
- [7] L.H.Quam, Road Tracking and Anomaly Detection in Aerial Imagery (SRI International AI Center Tech. Note 158), in *Proceedings, ARPA Image Understanding Workshop*, 1978, pp. 51-55.
- [8] Z.Zhu and Y.Kim, Algorithm for Automatic Road Recognition on Digitized Map Images, *Optical Engineering*, **28(9)**, 1989, pp. 949-954.
- [9] T.Nagao, T.Agui and N.Masayuki, An Automatic Road Vector Extraction Method From Maps, in *Proceedings, 9th Int. Conf. Pattern Recognition*, **1**, 1988, pp. 585-587.
- [10] S.Suzuki and T.Yamada, MARIS: Map Recognition Input System, *Pattern Recognition*, **23(8)**, 1990, pp. 919-933.
- [11] M.D.De Simone, *Data Structures and Feature Recognition: From the Graphic Map to a Digital Data Base (Unpublished thesis)*, N.E. London Polytechnic, 1985.
- [12] G.H.Kirby, M.Visvalingam and P.Wade, Recognition and Representation of a Hierarchy of Polygons with Holes, *Computer J.*, **32(6)**, 1989, pp. 554-562.
- [13] E.Rosch and B.B.Lloyd (Eds), *Cognition and Categorization*, Lawrence Erlbaum Associates, 1978.
- [14] G.W.Humphreys and V.Bruce, *Visual Cognition: Computational, Experimental, and Neuropsychological Perspectives*, Lawrence Erlbaum Associates, 1989.

- [15] M.Visvalingam and P.Williamson, *Generalising Roads on Large-Scale Maps: A Comparison of Two Algorithms*, (CISRG Discussion Paper 13), University of Hull, 1994.
- [16] W.R.Garner, Aspects of Stimulus: Features, Dimensions and Configurations, in *Cognition and Categorization*, (E.Rosch and B.B.Lloyd Eds.), Lawrence Erlbaum Associates, 1978, pp. 99-133.
- [17] P.Muehrcke, Maps in Geography, *Cartographica*, **18(2)**, 1981, pp. 1-41.
- [18] D.A.Varley, *Road Extraction and Recognition for Validation of Large Scale Topographic Data (Unpublished thesis)*, University of Hull, 1994.

