AUTOMATED LINEAR FEATURE EXTRACTION IN SUPPORT OF RAPID DATABASE GENERATION

Richard Ley, Steve Wallace and Nick Davies
Space Department, Defence Evaluation and Research Agency,
Farnborough, UK.

ABSTRACT

Part of the UK Ministry of Defence element of the STOW programme investigated the time and cost drivers pertaining to the entire process of the rapid generation of Synthetic Natural Environments (SNE) databases. Data requirements, products, information and systems were analysed to identify bottlenecks and gaps. Traditionally, construction of SNE databases is a time consuming and very labour intensive exercise. It involves a very high degree of effort to generate the required source terrain and feature data, and significant further effort to convert source data into a compiled SNE database.

Standard military datasets are typically used to provide the bulk of the data for a SNE database (e.g. DTED and DFAD). However, such datasets may not be available for the specific area of interest, they may be at an inappropriate scale, they require augmentation and they are likely to be based on out-of-date mapping sources. An alternative worldwide and up-to-date source is required. The new series of Earth Observing satellites are creating a large archive of up-to-date geospatial data. The major blockage has moved down the value-added chain and it is the conversion of data into information that has become the major time and cost driver.

An approach to automated feature extraction from EO imagery is presented which uses an object-orientated geodata model as the framework to store contextual knowledge and to use this in the control of feature extraction routines. The problem of geographic extraction has proved complex and ideally requires the incorporation of contextual clues similar to those used by human interpreters of imagery. Often the feature recognition algorithms work at local levels and in a bottom-up fashion and lack the higher level control that would allow a more global understanding of parts of the image. The paper proposes a control strategy that incorporates both the global and local views.

The geodata model comprises a class hierarchy representing the features under study and their likely relationships. Each class of object within this model contains criteria that need to be satisfied in order to strengthen the belief that an instance of that object type has been recognised. The criteria cannot be rigid and the system must be able to control partial recognition of objects and identify conflicts. The system described will apply these ideas to the problem of geographic object recognition, focusing on the specific requirements of linear feature extraction.

Authors' Biographies:

Richard Ley is a GI/GIS and remote sensing expert who specialises in the generation and manipulation of geospatial information. His research interest focuses upon techniques that support the rapid generation of geospatial databases. Richard's current role is head of the Geospatial Information & Services Group and he is the technical leader of the ALFIE project.

Steve Wallace is a GIS and remote sensing specialist with experience gained in the capture, manipulation, and integration of all forms of spatial data. Particular expertise has been developed in the derivation of both feature and terrain data from spaceborne and aerial imagery, using semi-automatic and automatic approaches. Steve's current role as a Senior Scientist, involves the technical management of research projects.

Nick Davies trained in geospatial sciences, gaining an honours degree in Computer Science and GIS with particular experience gained in cartographic data design and production, and database programming. Nick is involved with all aspects of GIS system integration and development, including requirements analysis and system design.

AUTOMATED LINEAR FEATURE EXTRACTION IN SUPPORT OF RAPID DATABASE GENERATION

Richard Ley, Steve Wallace and Nick Davies
Space Department, Defence Evaluation and Research Agency,
Farnborough, UK.

BACKGROUND

Part of the UK Ministry of Defence element of the STOW programme investigated the time and cost drivers pertaining to the entire process of the rapid generation of Synthetic Natural Environments (SNE) databases. Data requirements, products, information and systems were analysed to identify bottlenecks and gaps (Wallace *et al*, 2000, Davies *et al*, 2000). Traditionally, construction of SNE databases is a time consuming and very labour intensive exercise. It involves a very high degree of effort to generate the source terrain and feature data required, and significant further effort to convert source data into a compiled SNE database. Standard military datasets are typically used to provide the bulk of the data for a SNE database (e.g. DTED and DFAD). However, such datasets may not be available for the specific area of interest, they may be at an inappropriate scale, they typically require thinning, generalising, and augmenting with additional data, and they are likely to be based on out-of-date mapping sources. An alternative world-wide and up-to-date source is required.

This lack of source materials for SNE generation has often been cited as one of the major stumbling blocks. However, the new series of Earth Observing satellites are creating a large archive of up-to-date geospatial data. The major blockage has moved down the value-added chain and it is the conversion of data into information that has become the major time and cost driver. The abundance and detailed content of this imagery will remain inaccessible unless the information content can be readily extracted through automation. Thus, a 3-year research project, entitled Automatic Extraction of Information from Geospatial Data (alternatively titled ‘Automated Linear Feature Identification and Extraction’ (ALFIE)) was initiated by the UK MoD in 1999 to investigate ways of addressing this blockage. The focus of the research is the automatic extraction and attribution of roads, railways and rivers.

This is an area that has attracted considerable funding in the past with the emphasis being on the development of extraction algorithms based upon identifying features through their radiometric properties in the imagery. The view taken here is that the use of radiometric properties alone will not provide sufficient information to extract and attribute linear features. The key is to utilise as much information pertaining to that feature as possible. For example, in the manual extraction of a linear feature, the operator does not simply look at the spectral response to determine that a feature is a road. The human brain assimilates the context in which that feature lies. In an urban area, a road is likely to have buildings along one or either side. The style of junctions, line curvature and width can also all give clues as to the nature of the feature. It is this information which ALFIE will utilise. It will incorporate rules based upon contextual information, implementing them through a combination of spatial analysis and intelligent agents within an object-oriented environment. It should also be stated that this project does not intend to generate any new linear feature extraction algorithms. Many of these algorithms have been developed over the last 20 years. Each algorithm has benefits and limitations. In general, each algorithm works well within the domain for which it was devised and tested. Thus, some algorithms work well on very high resolution imagery, but not on imagery with coarser resolutions, while others work well on SAR images but not on multi-spectral images. Therefore, the ALFIE project will bring together a suite of state-of-the-art algorithms to maximise the use of a variety of image types. Any resulting time and cost reductions in the extraction of linear features will be a significant benefit not just to synthetic environments but to all areas requiring access to geospatial data. The ALFIE research team includes staff at the UK’s Defence Evaluation and Research Agency (DERA), Laser-Scan, a GIS company and the Department of Geography, Nottingham University, UK. This paper investigates the methodology employed and describes the current status of the process flow model.

METHODOLOGY

The methodology followed must support the prime requirement of rapid database generation for any region in the world; timeliness and universality are fundamental considerations. There are four major elements to the approach being followed, namely the datasets employed, the exploitation of contextual information, the use of an object-oriented database, packaged and manipulated through a control strategy.

Datasets

Datasets	Imagery			Cartographic	
	Multispectral	*Panchromatic*	*Radar*	*Features*	*Elevation*
Minimal	7 bands 30m	10m		Vmap Level 1	DTED Level 1
	3 bands 20m	5m	20m		
Optimal	3 bands 4m	1m	8m	Vmap Level 2	DTED Level 2

Table 1. Datasets employed

A review of initiatives in automated feature extraction show that significant progress has been made, particularly in respect to roads and to a lesser extent 3D compact structures (Baumgartner *et al.*, 1997; Heller, 1998). Generally, the procedures adopted focus upon exploiting very high resolution imagery, either airborne or spaceborne. However, the timeframes required for collecting, handling and processing very high resolution imagery usually exceed those available for rapid database generation. Therefore, a reliance on very high resolution imagery is contrary to our aims.

Instead, developed techniques will be utilised that exploit those datasets that are universally available; these will tend to be medium resolution imagery. If timelines are longer, additional high resolution datasets will be exploited and techniques will be devised to accommodate the geo-processing of these (see Table 1). These datasets will probably need to be collected specifically rather than taken from archive. They will incur greater processing times but will provide far more detailed and dynamic information.

In addition, the approach includes the potential to exploit existing cartographic data. If available, feature and attribute data extracted from standard military products will be included. Again, depending upon timeframes, different levels of products will be utilised. It is assumed that there will be world-wide coverage of DTED and VMap at level 1 (approximately equivalent to 1:250 000 scale mapping) and that level 2 products (approximately equivalent to 1:50 000 scale mapping) may become available at a later date. The cartographic data will provide context and may be used to seed or train algorithms that can provide more complete datasets than are available within the products themselves. No over-reliance will be made of the cartographic sources and the system will be designed to work on the imagery alone. However, experience teaches that all available collateral information should be exploited to improve the probability of a successful extraction and attribution (Tonjes & Growe, 1998). The extraction process must be able to work at minimal and optimal levels of dataset availability.

The minimal dataset comprises Landsat multispectral and SPOT panchromatic imagery plus two level 1 standard military datasets whilst the optimal dataset is based upon the new generation of high resolution spaceborne imagery. Currently, this is limited to Ikonos imagery (both multispectral and panchromatic) but additional sources are expected in the near-future. Radarsat provides the high resolution Synthetic Aperture Radar (SAR) imagery. Level 2 cartographic products are included in the optimal dataset. An intermediate grouping comprises the Indian IRS multispectral and panchromatic plus the European Space Agency ERS SAR datasets.

Context

Automatic object recognition techniques are used very successfully in the fields of engineering drawing recognition and industrial object inspection (Priestnall *et al.*, 1996). Here the number of different object types under study is limited and their size, shape and dimensions are predictable, or the drawings conform to some standard accepted code of practice. Geographical objects do not conform to such standards and the same assumptions regarding size and shape cannot be made especially when different types of imagery at varying spatial resolutions are the sources. However, some ideas and techniques are transferable to the geospatial domain. One crucial concept that is transportable is the use of contextual clues such as the relationship between one object and its neighbours, or the containment of an object within a region.

The importance of context in an extraction regime can be demonstrated by an image sequence (see Figure 1.) At the pixel level in the first image, it is difficult to identify the object, but as the view is widened then more 'supporting evidence' from the immediate context of the object is incorporated and so the object can be identified more confidently.

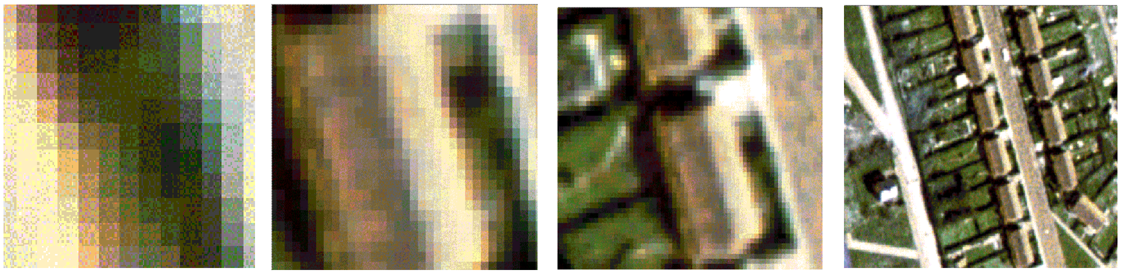

Figure 1: The use of context (© NRSC, 1996)

Spatial context can be considered at several levels (Baumgartner *et al.*, 1997). At the regional level, derivation of contextual clues of a spatial nature (such as containment) can be achieved using medium resolution imagery (20-30m pixels). Thus, context regions such as urban, rural and forest could be defined. This is important since the behaviour and relationship of objects may vary depending on the context region. On the more local level, the interrelationships between features can be defined in closer terms. These can be termed “sketches”. For example, one sketch may be “occlusion shadow” that might consist of two road segments broken by the shadow of some high object. Context at this local level is only really applicable when high resolution imagery is used. Within a linear feature extraction strategy, context regions can be used to limit the choice of acceptable feature recognition solutions or to refine the parameters used. Attempts at feature recognition often focus on the geometric character of one type of feature in isolation in order to manage the complexity of the problem as a whole (Wang & Howarth, 1991). However, there is much to be gained from attempting to incorporate more contextual information from the image in terms of interrelationships between different feature types. Junction features, for example, hold vital network topological details and can form good starting points for searches to complete linear segments and begin growing the network. In addition, features represented in high resolution imagery often appear as complex composite features where no single feature extraction algorithm will suffice. A framework will be developed in order to represent these complex geographic features and relationships in a hierarchical fashion so knowledge of a feature’s locus can be used to suggest likely feature recognition procedures.

Object-oriented database technology

The object-oriented database will contain a geo-model composed of a class hierarchy representing the features under study and their likely relationships (see Table 2 for an example of this model related to railways). Each class of object within this model will contain criteria (in the form of class members) which need to be satisfied in order to strengthen the belief that an instance of that object type has been recognised. The criteria cannot be rigid and the system must be able to control the recognition of features and identify conflicts. The ability to allow several partially recognised features to mutually confirm each other is a vital step towards the introduction of geospatial context. Another benefit of the object database will be the ability of object classes to contain knowledge of the types of feature recognition that are appropriate for a particular object type at a particular scale. This will allow a partially recognised feature to effectively attempt to classify itself.

The modelling of ‘real world’ objects encapsulates not only the geometry and attribution but also the structure and behaviour of these objects. Structure and context have been difficult to incorporate in classic ‘image processing’ systems but become a natural component of an object-oriented database. It does this by the use of intelligent agents. These agents are pre-defined to determine the most appropriate class (e.g. road, railway etc.) dependent on the description, behaviour and association of each line. The implementation of these agents is carried out transparently from the user’s perspective. In considering linear features, active agents will automatically determine, for example, the width and gradient of them and compare these with the parameters for each type of feature. Table 3 gives an example of the ranges of widths and gradients expected for roads and railways in a developed, non-mountainous country. Furthermore, the ability to rapidly experiment and prototype through use of classification methods attached to the object classes offers up a novel research environment.

The hierarchical data model and the intelligent agents are being implemented within an object-oriented database that utilises a Geographic Information System (GIS) as a front-end. Laser-Scan’s Integrated GIS (IGIS) and its GOTHIC spatial database were selected for the prototype system.

Class	Sub-class	Class members
Known	Railway	Maximum gradient permissible
		Maximum curvature permissible
		Width
		Junction type
		Connectivity
		Associated features

Table 2: Subset of the geodata model

Intelligent agent	**Contextual rules**							
	Railway		**Motorway**		**Main road**		**Minor road**	
	Max	*Min*	*Max*	*Min*	*Max*	*Min*	*Max*	*Min*
Gradient	1 in 50	None	1 in 10	None	1 in 7	None	1 in 4	None
Width	20 m	3 m	30 m	20 m	15 m	6 m	6 m	3 m

Table 3: Illustrative intelligent agents and their associated contextual rules

Control Strategy

The final element to the research is the definition and implementation of a control strategy. This exploits the use of *a priori* knowledge (e.g. class, context, containment, connectivity, duration) to undertake per object classification techniques. It then builds upon this by employing a variety of image recognition algorithms. There are a number of these algorithms already developed but their effective exploitation is hindered by the lack of a framework that automatically assesses which set of algorithms should be used in which set of circumstances. These include the type and resolution of imagery, the type of feature, the type of contextual region and the timeframe. The control needs to be self-assessing so that levels of confidence can be assigned to any outputs. Thus, a means of defining assessment metrics and confidence levels are being developed. The strategy allows output at various stages, ranging from a quick snapshot of readily classifiable features progressively through time to a more detailed classification of all the major linear man-made features.

SELECTED ALGORITHMS

A review of Defence and public domain algorithms was undertaken and a down selection was made (Ducksbury (1999). Five packages of algorithms will be incorporated in the future work. Two are in the public domain (G-Snakes and Multi-resolution edge linker), one is proprietary (Laser-Scan's V-trak) and the final two are internal to DERA (SUSAN and Linefinder). The algorithms have been selected to ensure that extractions can be made from a variety of image types and resolutions. For example, SUSAN is an effective edge detector and hence works well on medium resolution imagery where linear features are typically only two or three pixels wide. Linefinder, on the other hand, uses the Marr-Hildreth filter that detects centrelines and thus is only effective on high resolution imagery where both sides of a linear feature are identifiable. Linefinder also incorporates some clever processing to reduce the amount of clutter generated by the centreline filter. To ensure maximum benefit is gained from these algorithms, a flexible approach has been taken in porting them to the Commercial Off The Shelf (COTS) GIS. Hence, the use of an edge detector can be used instead of the centreline filter within the Linefinder algorithm if the circumstances demand it. This means that existing algorithms can be tailored to specific image types, thereby extending their use outside of the domain for which they were originally designed.

PROCESS FLOW

The process flow may comprise two complete passes; the first is at the GLOBAL level during which the major features will be extracted. Medium resolution datasets will be exploited. The results of this pass may be fed into the second pass at the LOCAL level. This only occurs if there is both more detailed source material available (i.e. high resolution imagery) and time to process it (see Figure 2).

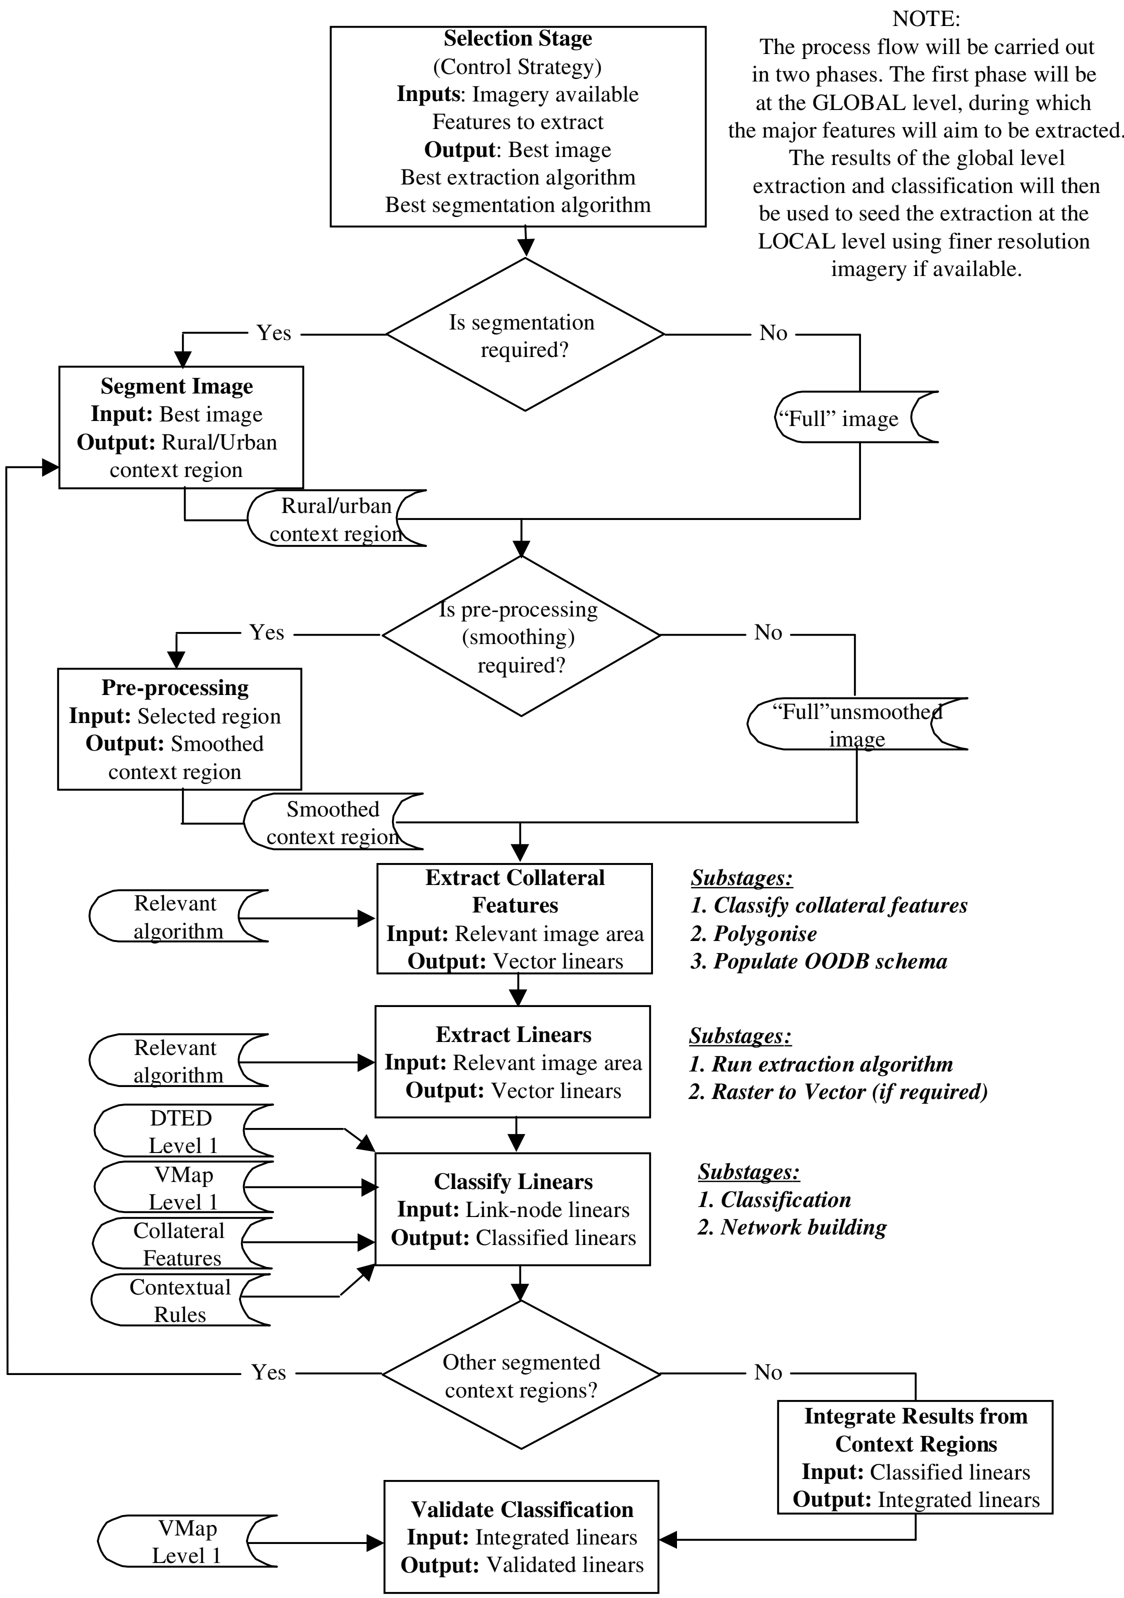

Figure 2. ALFIE Process Flow

Control Stage

The aim of the control strategy is to provide a framework from which the most appropriate algorithms (including the actual feature extraction algorithms as well as image segmentation and pre-processing algorithms) are selected automatically, given the type of imagery available and the features to be extracted. The algorithms will also be parameterised automatically. The control stage is sub-divided between segmentation, domain of interest and selection sub-stages.

Segmentation Sub-Stage. Contextual regions are defined within the imagery. The actual segmentation process to be employed is determined by the datasets available and the types and scale of features to be extracted. These are mainly based upon spatial differences although there may be some aspatial variation involving the quality of the imagery. Adjacent scenes from the same sensor often exhibit different tonal balances and so may require different pre-processing to optimise contrast etc. Rarely can this occur within a single scene. The spatial differences are dependent upon variations in land surface characteristics. These include variations in major land cover (i.e. water/land; urban/rural; open/woodland), terrain (gentle/steep sloping) and *a priori* knowledge of major communication corridors (e.g. the presence/orientation but not the exact location of roads, railways etc). Each of these may well have implications on not only the fine-tuning of the algorithms but also on the algorithm to be used. Contextual regions are over-lapping rather than mutually exclusive. They are segmented using the medium resolution datasets using different algorithms depending upon the type of region and the dataset used. During the LOCAL pass, contextual sketches rather than regions are extracted.

Domain Of Interest Sub-Stage. Transportation networks vary considerably throughout the world. For example, civil engineering constraints of maximum gradients and curvature for railway construction may be relaxed in areas of rugged terrain. The maturity of road networks will vary considerably depending upon population density, national wealth and alternative means of transport. The selection of the domain of interest is generally based upon *a priori* knowledge.

Selection Sub-Stage. Once the datasets have been divided into domains and homogeneous contextual regions, the selection process will determine both the optimal datasets to be exploited and the algorithms to be employed through a series of pre-determined look-up tables. Furthermore, the individual parameters of the algorithms will be specified in a similar manner. This may involve a multi-resolution approach given the selected combination of image type and linear feature width (and number of different widths). Finally, the control strategy will determine the required level of pre-processing of the imagery. This is very much linked to the set of algorithms that are to be employed.

Pre-processing Stage

The primary aim of this stage is to prepare the input imagery to aid the extraction algorithms. For example, an efficient smoothing algorithm that retains edges could be run over the input image to enhance the effect of an edge-extracting algorithm. In the same way, a Normalised Difference Vegetation Index (NDVI) image could be derived prior to the extraction of woodland and trees (which may be required for collateral information). The pre-processing stage is sub-divided between image processing and training sub-stages.

Image Processing Sub-Stage. The type of processing is determined by the type of datasets exploited, the features to be extracted, and the algorithms selected. The Control strategy pre-determines the level of pre-processing. This sub-stage may include improving the overall characteristics of the imagery (e.g. contrast or brightness) and/or focussing on improving specific features through edge enhancement or smoothing filters before running the selected algorithm. Furthermore, standard image processing techniques such as band ratioing may be used, first to enhance a particular feature of interest, and secondly it enables three or more multispectral bands to be integrated into a single band that emphasises the information contained within each of the input bands.

Interactive Training Sub-Stage. This is kept to a minimum but could provide illustrative geometry and attributes of the linear features present from the cartographic datasets. These, when compared to the imagery, could provide training information concerning the width, spectral consistency and contrast of the linear objects. However, the positional inaccuracies of the cartographic Level 1 products may reduce this sub-stage to one of identifying communication corridors only.

Extract Collateral Features Stage

The use of context within the extraction methodology requires some collateral information to be extracted. For example, trees often line suburban roads, while hedgelines often bound minor rural roads. Before this information can be utilised within the classification stage, it has to be extracted from the imagery. This then is the aim of the collateral features extraction stage. Rivers and canals will also be extracted at this stage. This is because water is relatively easy to extract from multispectral imagery (in comparison to roads, railways etc.) due to the near unique total absorption of near infrared light in water. This stage will be sub-divided between the extraction of the collateral features and a raster to vector conversion.

Extraction Sub-Stage. Features such as waterways and hedgelines are extracted. These will be used to 'mask' the lines representing these features from the set of lines provided by the linear feature extraction algorithm.

Polygonisation Sub-Stage. The extracted collateral features are vectorised to enable overlay on the extracted linear features.

Extract Linears Stage

It is at this stage that the actual linear feature extraction algorithms are run. The output is a group of discrete 'unknown' linears. Since a number of the algorithms output raster formatted data, this is sub-divided between the specific linear feature extraction followed by a vectorising sub-stage.

Extraction Sub-Stage. The selected datasets of an individual context region are passed through the chosen algorithms in order to extract the linears. The results of this stage are a number of disconnected raster polygons including both signal and noise. If either multi-resolution or multi-spectral imagery is employed, a series of overlapping polygons is created.

Raster-To-Vector Sub-Stage. The raster polygons are converted to vectors with some tagging of geometric (e.g. length, width, consistency of width and pattern of feature) and radiometric intelligence (e.g. homogeneity, specific spectral response).

Classify Linears Stage

From a situation where the system contains 'unknown' linears, the aim is to categorise and classify these unknown lines into specific classes (road, railways, rivers etc.). This stage is sub-divided between classification and network building sub-stages although the process is iterative rather than sequential.

Classification Sub-Stage. The vectors are analysed against the contextual rules that are held within the object-oriented database as intelligent agents. If the vector is long enough, it is given a radius of curvature and gradient to accompany any other attribution already assigned. Initially, the vector is given the object class 'unknown'. It is then tested against each of the rules pertaining to rivers, roads and railways and a running score or confidence level is recorded. Once this probability reaches a prescribed level in one of the three categories, its status changes to line 'known' (e.g. river, road or railway). Sub-categories exist within roads (e.g. motorways, number of lanes, hard or soft surfaced etc).

Network building sub-stage. As a result of the initial classification, confidence levels prioritise a cycle of further labelling and network building. Starting with high confidence features such as strong T-junctions, linear segments are followed, initiating spatial searches to gain further local contextual clues and to bridge gaps in the vectors. Extended but status 'unknown' networks are passed back to the classification sub-stage for further attribution and confidence level building. This loop is repeated either until all features are classified or until no further enhancement of confidence levels is created.

If there are contextual regions still to be analysed the process returns to the pre-processing stage. If all contextual regions at the current scale are processed, the validation stage is initiated.

Validation Stage

Details of this stage are currently being compiled. It is planned to incorporate statistical reporting of the process (e.g. the average confidence levels for each feature, the proportion of status 'knowns' to 'unknowns' and the connectivity of the derived networks). It may be possible to compare these with statistical summaries of existing geospatial information in similar geographical areas (e.g. density of various road categories per square kilometre). This type of information is being recorded elsewhere in order to generate geo-typical relief (Chapman 1999).

The state of recognition of the linear features is displayed by highlighting the status 'unknown' features along with their current confidence levels. These unknowns are classified with the 'safest' option as vehicle navigation could be based upon the resulting geospatial information.

After the GLOBAL pass is completed, the major lines of transportation are both extracted and classified. Time and dataset availability permitting, the cycle is repeated at the LOCAL level. GLOBAL features become available as additional contextual information to support processes performed at the LOCAL level.

Future work

Within the ALFIE research project, the second year work is concentrating upon defining the details of the Control Strategy, building the object-oriented database including the 'intelligent' agents and porting the algorithms into the target GIS. A prototype system will be built during year three. Upon the successful completion of the prototype, it is planned to move the extraction of linear features into the production domain and extend the research to include other features, such as buildings and industrial land cover.

REFERENCES

Baumgartner, A., Eckstein, W., Meyer, H., Heipke, C. & Ebner, H. (1997). Context-supported road extraction. In, Gruen, A., Baltsavias, E.P. and Henricsson, O. (eds) Automatic Extraction of Man-Made Objects from Aerial and Space Images (II). Birkhauser Verlag, Basel, pp299-308.

Chapman, C & Drysdale, J (1999), Integrated Geo-Typical Terrain Generator. Working Paper 3: Human Culture and Networks, Settlement Positioning And Transportation. Version 1.0. 297-32.WKP, Unpublished Working Paper.

Davies, N. Kodz, D. Jarvis, K & Jones, S. (2000). Research Study into Data Requirements and Architectures for Rapid Terrain Database Generation. Volume 2: System Architecture. DERA/CIS/CIS2/CR39000039/1.0/Vol2 Unpublished DERA Report.

Ducksbury, P. G. (1999). Short Survey - Linear Feature Detection Techniques. DERA Malvern, Sensors and Processing Sector. Unpublished, internal working paper.

Heller, A. J., Fischler, M. A., Bolles, R. C., Connolly, C. I., Wilson, R. & Pearson, J. J. (1998). An integrated feasibility demonstration for automatic population of geospatial databases: Annual progress report. <http://www.ai.sri.com/~apgd/papers/apgd-iuw98-pi.pdf> (August, 1999).

Priestnall, G., Marsden, E. & Elliman, D.G. (1996). Arrowhead recognition during automated data capture. Pattern Recognition Letters.

Tonjes, R., & Growe, S. (1998). Knowledge-based road extraction from multisensor imagery. Commission III, working group 4.

Wallace, S. Wilkins, H. Jarvis, K & Jones, S. (2000). Research Study into Data Requirements and Architectures for Rapid Terrain Database Generation. Volume 1: Data Requirements. DERA/CIS/CIS2/CR39000039/1.0/Vol1 Unpublished DERA Report.

Wang, J. & Howarth, P. J. (1991). Structural measures for linear feature pattern recognition from satellite imagery. Canadian Journal of Remote Sensing, 17(4), 294-303.