Institute of Terrestrial Ecology

[redacted]

Centre for Ecology & Hydrology

Natural Environment Research Council

ITE has six Research Stations throughout Britain, which allows the efficient use of resources for regional studies and provides an understanding of local ecological and land use characteristics. The Institute's administrative headquarters is at Monks Wood.

This report is an official document prepared under contract between the customer and the Natural Environment Research Council. It should not be quoted without the permission of both the Institute of Terrestrial Ecology and the customer.

ITE sites

Monks Wood
(Admin HQ)
Abbots Ripton
HUNTINGDON PE17 2LS
Telephone 01487 773381-8
Fax 01487 773467
Email MONKSWOOD@ITE.AC.UK

Merlewood Research Station
GRANGE-OVER-SANDS
Cumbria LA11 6JU
Telephone 015395 32264
Fax 015395 34705
Email MERLEWOOD@ITE.AC.UK

Edinburgh Research Station
Bush Estate
PENICUIK
Midlothian EH26 0QB
Telephone 0131 445 4343
Fax 0131 445 3943
Email BUSH@ITE.AC.UK

Furzebrook Research Station
WAREHAM
Dorset BH20 5AS
Telephone 01929 551518-9, 551491
Fax 01929 551087
Email FURZEBROOK@ITE.AC.UK

Banchory Research Station
Hill of Brathens
Glassel, BANCHORY
Kincardineshire AB31 4BY
Telephone 01330 823434
Fax 01330 823303
Email BANCHORY@ITE.AC.UK

Bangor Research Unit
University of Wales, Bangor
Deiniol Road
BANGOR, Gwynedd LL57 2UP
Telephone 01248 370045
Fax 01248 355365
Email BANGOR@ITE.AC.UK

Details about the Institute are available on the Internet via the World Wide Web (http:/www.nmw.ac.uk/ite)

INSTITUTE OF TERRESTRIAL ECOLOGY
(NATURAL ENVIRONMENT RESEARCH COUNCIL)

Project. T02083j5

Countryside Survey 2000
Module 7

LAND COVER MAP 2000

SECOND INTERIM REPORT

incorporating the

Fourth Quarterly Progress Report

CSLCM/Int2/Prog4

R.M. Fuller, G.M. Smith, J.M. Sanderson, R.A. Hill and A.G. Thomson,
Institute of Terrestrial Ecology
Monks Wood
Abbots Ripton
Huntingdon
Cambs. PE17 2LS

Corresponding author:
R.M. Fuller
Section for Earth Observation
Institute of Terrestrial Ecology
Monks Wood
Abbots Ripton
Huntingdon
Cambs PE17 2LS

Telephone:	01487 773 381
Fax:	01487 773 467
Email:	R.Fuller@ITE.AC.UK

25 March 1999

MAIN REPORT

EXECUTIVE SUMMARY AND QUARTERLY PROGRESS REPORT

- This is the second Interim Report on Land Cover Map 2000 (LCM2000), a part of Countryside Survey 2000 with links to the Northern Ireland Countryside Survey. The Report covers work done in Britain, to 19 February 1999.
- LCM2000 aims to make a census survey of the land cover / widespread broad habitats of the United Kingdom using satellite imagery and automated image processing techniques to achieve a classification accuracy of 90% for target classes.
- ITE has purchased winter TM imagery and matching summer TM and IRS imagery of most of England and Wales. ITE have also purchased a digital terrain model (DTM) for improved analyses.
- Image searches have been completed for half the winter of 1998-9. The acquisition rate is poor reflecting the continuing wet and cloudy weather over Britain. Even if no further scenes are added, acquisitions in winter 1997-98 give near complete winter coverage.
- LCM2000 plans to map widespread examples of Broad Habitats, reported in detail in the First Interim Report.
- The early stages of LCM2000 have involved procedural developments. Some have been software refinements; many have required the development of macros to automate the processes of image analysis; others simply formalise and make objective the complete sequence of activities.
- Geo-registration of multi-sensor images and resampling algorithms for segmentation outputs have demanded investment of time in methodological developments. Other pre-processing facilities have included automated procedures to identify and mask out clouds and cloud-shadows (ITE macro). Otherwise, the Specification envisaged limited pre-processing. However, further improvements in the processing stream include atmospheric corrections (new ITE software) and terrain-induced illumination corrections (Cambridge University subcontract).
- The various pre-processing improvements have demanded a more complex approach to image pre-processing, delaying the start of production mapping; however, the improvements will offer consequent time savings at post-classification correction stages.
- Laser-Scan has implemented a fully operational, Unix version of the prototype CLEVER-Mapping segmentation software, designed by Cambridge University colleagues.
- Segmentation consists of two separate stages: edge-detection to identify boundary features; and region growing from seed points.
- Experiments have examined: band selection for edge-detection and segmentation; choice of edge-detector; threshold values to identify edges; thresholds to control the degree of segmentation; post-segmentation boundary rejection and generalisation.
- After segmentation, raster-to-vector conversion draws boundaries between segments.
- The field reconnaissance data were collected in 1998 to cover analyses of most stock-images. The data are used to train the classifier, with field-mapped land parcels being identified on segmented images as training or validation polygons.
- Training for per-parcel analyses is objectively based on the segments, identified individually by pointing to them, without the need to painstakingly draw the outline. The process of training is probably 20-50 times quicker than per-pixel training, allowing a much larger sample and the definition of spare 'check' parcels for use in validation.
- Another software refinement allows the review of training areas using display image 'chips' to show the quality of the remotely sensed data in each training area. The operator

can compare training areas, define spectral subclasses, reject odd examples and flag training or validation polygons. The review facility saves much time in classification.

- The object-oriented procedures of CLEVER-Mapping allow a classification of one scene to be 'rolled over' to an overlapping neighbour to extend labelling to polygons on the new image. The training data review procedure ensures that the training set matches standards which an operator would apply.
- Classification requires that the training polygons are used to interrogate the image to derive statistical measures for reflectances in each chosen band and for each spectral subclass. CLEVER-Mapping uses a shrinking procedure, when extracting raster data for polygons, to avoid edge pixels and ensure the use of 'pure' core pixels of a cover type. The shrinkage is a dynamic process minimising edge effects while ensuring an adequate sample in smaller polygons.
- The classification procedure applies a maximum likelihood algorithm to the polygon mean statistics to select the most likely class in statistical terms. CLEVER-Mapping in IGIS also records the probabilities for all other subclass options. Per-pixel classifications record the natural heterogeneity associated with polygons.
- Validation has been designed as a two-stage process: first, using the check-polygons to validate the results derived from training polygons; second, using field survey data to validate and calibrate results. The former will assess how far extrapolation attaches the correct labels to areas of known cover. The latter will offer a method for 'translation' between detailed field classes and the more generalised target classes of LCM2000.
- Production has involved processing 5 scenes in parallel from initial atmospheric corrections through geo-registration, illumination correction, co-registration of summer-winter data and segmentation stages. Preliminary classifications will be produced in March 1999.
- Initial classifications have shown that, in principle, the Broad Habitat classification can be achieved. Overall success will only become clearly apparent once finalised classifications result. Validation based on check polygons will provide an early and objective measure of success.
- The GANNT (Figure 1) records progress in scene-pair equivalents. The assessment of progress has taken account of the development work on pre-processing and its contribution to overall progress.
- Certain tasks are a little ahead of schedule: image searches are half-completed for winter 1998-99. Image purchases cover 50% of Britain, against an original expectation of 40% coverage.
- Methodological developments have been extended to refine methods, automate procedures, and reduce post-classification analyses.
- On balance, progress is perhaps a month behind schedule but increased throughput, resulting from refined pre-processing procedures, will quickly make good that delay.

Figure 1. GANNT - The timetable for Land Cover Map 2000 (UK) and estimated progress in production (to the end of February 1999)

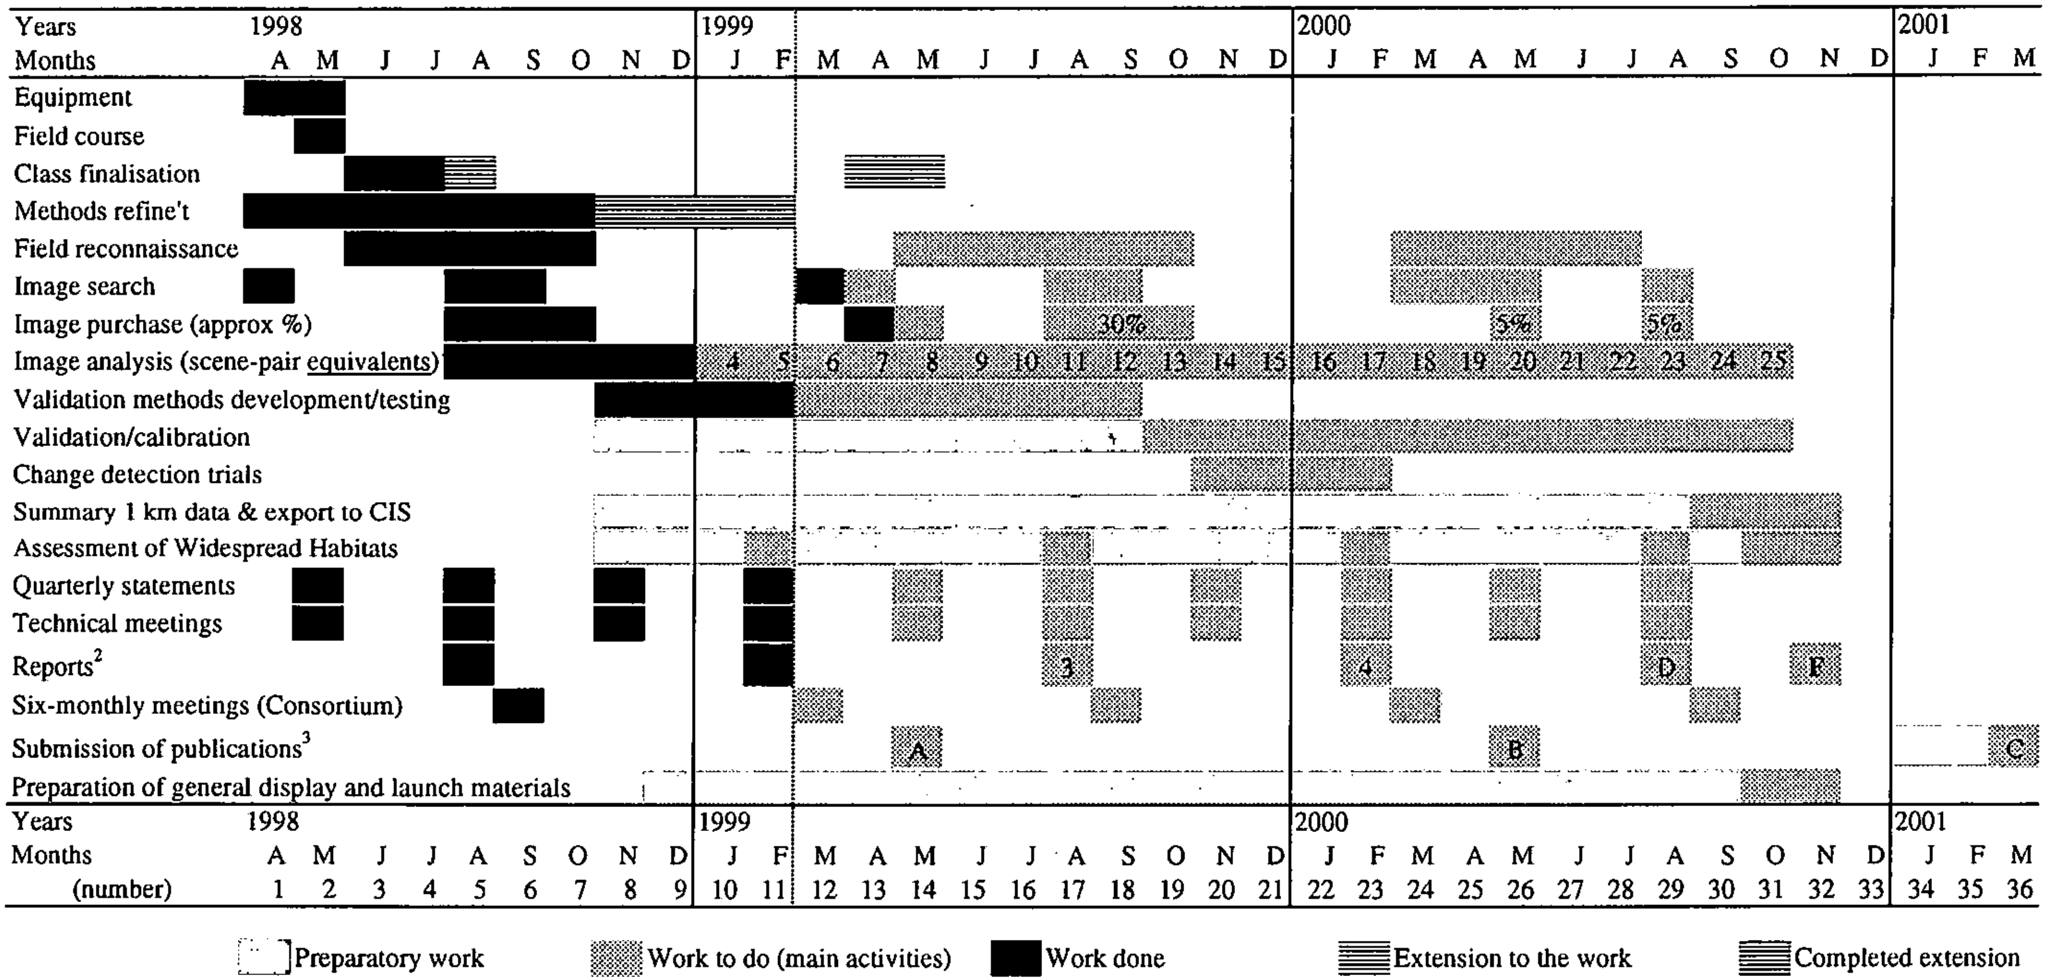

MAIN REPORT

1. INTRODUCTION

Countryside Survey 2000 (CS2000) is a major national audit of the habitats, plants, landscape features and land types of the British countryside at the end of the Millennium. The Survey is a jointly funded research programme involving several Government Departments, Agencies and the Natural Environment Research Council (NERC). Land Cover Map 2000 (LCM2000) is part of CS2000. With the recent inclusion of Northern Ireland, LCM2000 also now has links to the Northern Ireland Countryside Survey.

LCM2000 will provide a census of the countryside of the United Kingdom. LCM2000 outputs will be in the form of digital maps and databases, plus a range of derived products, held in a geographical information system (GIS). The Consortium of Departments and Agencies who are funding LCM2000 is listed in Table 1.

Table 1. The Consortium of funding agencies for Land Cover Map 2000.

Countryside Council for Wales
Department of the Environment, Transport and the Regions
Environment Agency
Ministry of Agriculture, Fisheries and Foods
Natural Environment Research Council
Scottish Natural Heritage
Scottish Office
Welsh Office
Environment and Heritage Service (Northern Ireland)
Department of Agriculture for Northern Ireland
Ordnance Survey of Northern Ireland

Most of this Report pre-dates the final agreement on inclusion of Northern Ireland in LCM2000; it therefore covers work done in developing a strategy for mapping Great Britain. Extra details on the extension to Northern Ireland are given separately (Fuller *et al.* 1999a).

2. BACKGROUND

LCM2000 updates and upgrades the *Land Cover Map of Great Britain* (LCMGB), made in 1990-92. Refinements include:

- Improved accuracy of classification,
- Added thematic detail,
- Compatibility with other systems of environmental survey and evaluation,
- Closer integration between field and satellite data.

With these improvements in mind, ITE undertook a range of methodological developments. Most important of these is CLEVER-Mapping - the Classification of Environment with Vector- and Raster-Mapping (Smith *et al.* 1998).

3. AIMS

The aims for LCM2000 are:

- To undertake a census survey of the land cover / widespread broad habitats of the UK at the turn of the Millennium;
- To apply the most appropriate satellite imagery and automated image processing techniques to achieve a classification accuracy of 90% for target classes;
- To produce and make available, under licence, a range of geographically referenced data outputs on land cover characteristics, tailored to the needs of Consortium members;
- To calibrate and validate satellite-derived classifications against ground reference data, publish results of the correspondence analyses and provide a guide to their interpretation.

This Report is the second six-monthly *Interim Report* on LCM2000 (see Fuller *et al.* 1998c). It follows the earlier reports, i. to the Joint Management Team of CS2000 (Fuller *et al.* 1998d & 1998e), ii. the *First* and *Second Progress Reports* (see Fuller *et al.* 1998b & 1998c) to the LCM2000 Consortium and iii. inclusion in the *Countryside Survey 2000 Integrated Progress Reports* (Fuller *et al.* 1998f & 1999b). The present Report covers work done up to 19 February 1999.

4. IMAGES

4.1 Acquisitions in 1998

Image acquisitions were described in earlier reports (Fuller *et al.* 1998c, 1999a, 1999b). Image searches for winter-summer 1997-98 showed that, with substitute second scenes as necessary, there was 90-95% coverage by Landsat Thematic Mapper (TM) data for the target 1997-98 winter period. Coverage for the target 1998 summer period was more problematic requiring TM images and Indian IRS-1C, LISS III data to give 60%-70% coverage for the target 1998 summer period mid-May to late July. TM also imaged SE England in August, too late for arable crop mapping, but complementing the LISS data to add detail for semi-natural cover types. Overall, well over half of the UK was covered by both summer and winter imagery.

ITE purchased winter TM imagery and matching summer TM and IRS imagery of England and Wales. ITE and the Consortium agreed to delay work in Scotland until next year, when the availability of late summer 1998 and summer 1999 imagery will also be known.

4.2 Image searches winter 1998-99

Searches have been completed for about half the winter of 1998-9; further advance is now slowed by the rate at which Eurimage, the distributors, add quick-look images to the website archive. Total acquisitions for northern and western Britain and Northern Ireland, where coverage is still needed, 60 new scenes were added to the archive. Only two winter scenes show <10% cloud cover; three scenes have 11-30 %; and 14 scenes have 31-50 %. While this acquisition rate is not good and reflects the continuing wet and cloudy weather over the UK, the best period in any winter would be the late winter early spring when dates will most closely match summer 1999 acquisitions. Even if no further scenes are added, acquisitions in winter 1997-98 give near complete coverage: though such scenes might not be as suitable for mixing with summer 1999 images, they nonetheless all but guarantee full winter coverage; and, in the most problematic northern and western areas, the transient nature of arable cropping, which renders older scenes less valuable, is less problematic than elsewhere.

5. BROAD HABITATS

The Specification for LCM2000 sets out a plan to map, as far as possible, widespread examples of Broad Habitats, as defined under the Biodiversity Action Plan. A list of target cover classes has been defined which generally match widespread Broad Habitats, only excluding small scale features, those which incorporate land use or contextual characteristics and marine types. The classification was reported in detail in the First Interim Report (Fuller *et al.* 1998c) with amendments agreed at a subsequent Consortium Meeting. Further Subclasses and Variants will allow relation to NLUSS and other classifications.

Broad Habitats based upon the 'presence' rather than the 'dominance' of plant types (e.g. woodlands described as >25% trees, bogs based upon >25% cover of peatland species) may not be discernible on images and contradict the concept of a cover map: thus, woodland polygons might be mapped as their dominant cover per-parcel but with the mosaic of trees measured per-pixel; bogs and heaths will be the focus of especial attention, with expert advice, in the field reconnaissance surveys of 1999.

The LCM2000 team attended the field surveyors' Training Course to ensure application of the Broad Habitat classifications matching, as near as possible, that of field surveyors. A *Key to Vegetation and Land Cover Types*, given in the *Field Handbook*, which identifies vegetation and contextual indicators was used by the LCM2000 reconnaissance team during the collection of ground reference data. The Consortium Meeting of 9 September 1998 raised what members saw as residual problems relating to sites with as little as 25% tree cover, heather- or grass-dominated bog and heath mosaics.

The currently planned classification appears in Table 1, superseding that of the First Interim Report. Uncertainties still seem to remain, more widely, regarding upland bogs, heaths and moors and the LCM2000 team have undertaken to meet in the field with the SNH peatland survey team, the CCW Phase I team and with Northern Ireland Countryside Surveyors to clarify and possibly varying definitions and generate a consistent and applicable definition for remote sensing. Any further amendments to the classification will refine rather than radically altering plans: none will be allowed to cause inconsistencies in relation to earlier outputs. All will be put before Consortium Meetings. **It is important to remember that the final classification represents an aggregation of spectral subclasses tailored to match the target list: the subclass details will be permanently held in the GIS allowing alternative re-aggregation at any time.**

Ground reference data collection in 1998 consisted of 6 reconnaissance surveys, mostly in lowlands, structured to maximise diversity of cover recorded. A total of 163 image subsets were marked up, giving *c.* 10 000 items of land cover information for future training. The completion of reconnaissance surveys provides the ground reference data needed for classification of stock images until late spring 1999.

6. IMAGE ANALYSIS

Operational use of CLEVER-Mapping based upon image segmentation is (as has always been made clear) entirely new and unique to LCM2000. It has necessarily offered a steep learning curve in the early stages. Furthermore, refinements of CLEVER-Mapping were being made right up until production started, and a number of pre-operational procedures were being brought into operation during the early phase of LCM2000. These developments were designed to improve accuracies and thematic details, beyond the immediate demands of the original Specification, but to the benefit of the entire Consortium and other potential users.

The early stages of LCM2000 have also involved large amounts of procedural developments. Some have been in the form of software refinements; many of them have required the development of macros to automate, as far as possible, the processes of image analysis; but others simply formalise and make objective the complete sequence of activities needed to conclude satisfactorily the making of LCM2000. Methodological issues are still developing quickly; these are now small modifications, 'fine tuning' really, but significant advances in the efficiency and/or accuracy of the production process.

The following sections outline the large investment of time in the design of techniques and analytical 'tools' which now form the basis of the operational programme. While every effort has been made to present an up-to-date account, continuing refinements will undoubtedly soon render some statements out of date in this early production phase.

6.1 Pre-processing of image data

The Specification envisaged very limited pre-processing. Between the time when the proposal was written and the analytical work commenced, a number of very significant improvements in pre-processing were explored. These, it is believed, can substantially improve results. For example, correction of atmospheric haze would prevent misclassification due to the distortion of spectral signatures by the variable atmospheric contribution; removal of terrain illumination differences would best ensure that the segmentation of images was based upon cover-differences rather than sunlit and shaded facets of topography.

It was decided that ITE would incorporate these improvements into the processing stream. It was recognised that additional developments of procedures would delay the start of operational processing, but it was concluded that such a delay would be compensated by later time-savings, when classification-training would not need to compensate for distorted spectral signatures and when subsequent knowledge-based corrections would be far less necessary. These developments have been outlined in earlier reports, discussed at Consortium meetings and were presented in detail (and approved) at the Technical Advisory meeting of 20 January 1999.

Three possible pre-processing stages have been investigated: de-striping, atmospheric correction and illumination correction. In addition, questions regarding geo-registration of multi-sensor images and resampling algorithms for segmentation outputs have demanded a further investment of time in methodological developments. Other pre-processing facilities have included automated procedures to identify and mask out clouds and cloud-shadows.

The additional work concerning pre-processing was unscheduled in the original Specification. While these potential improvements have proved too useful to ignore, they demanded a more complex approach to image pre-processing. This has delayed the start of production mapping. However, the improved classification which will result, will offer consequent time savings at post-classification stages. The experiments, results and transfer to operations are described below.

De-striping

It had been hoped to employ the de-striping procedure demonstrated by Helder *et al.* (1997), but this procedure could not handle the format in which European TM data supplied. As TM data are not severely affecting by such striping, ITE decided that minor striping problems could be accommodated by virtue of the per-segment approach to classification. The IRS data are more badly affected by striping but unfortunately the suppliers have 'geo-corrected' the images, to compensate for Earth rotation, using interpolation algorithms. In so doing, and because the 'push-broom' scanner configuration (unlike TM's) works at right-angles to the Earth's rotational movement the suppliers have skewed the vertical stripes into diagonals on the image. This prevents re-alignment of the image data to the original structure, thereby completely precluding any sensible de-striping process. While this is a problem, geo-registration with cubic convolution resampling (see later) will smooth these affects; and per-segment classifications will ensure that residual striping is 'averaged out' within segments. So it is considered that the problem can be overcome by pragmatic means.

Atmospheric correction

Atmospheric haze both attenuates the amount of light reaching the surface and also scatters light which has not interacted with the surface into the sensor. These effects are most pronounced in shorter wavelengths and distort the information recorded by the sensor. Various algorithms are available which attempt to model the effects of the atmosphere on the light passing through it, based on the information recorded in the image. The algorithms generally try to identify areas in the image for which the true reflectance can be estimated and in this way assess the distorting characteristics of the atmosphere at the time of imaging. These characteristics are then used to model the atmosphere and remove its effects from the image.

Liang *et al.* (1997) developed such a method for TM data which automatically seeks out examples of cover types whose reflectances can be estimated and interpolates the 3-dimensional atmospheric characteristics between the examples to correct the image. The software has been made available for ITE use. It has been modified to accommodate differences in the of the European TM supply format. It has been made operational and tested with full TM scenes. The correction is highly effective, removing not only haze but even penetrating areas of thin cloud on the uncorrected image.

The correction of IRS LISS data cannot directly use this software which relies upon the presence of the greater number of bands of TM data. However, the same principles can apply: other bands, sensitive to atmospheric haze, can be used to estimate atmospheric characteristics and thereby compensate for its effects. These developments of the Liang software have been made by ITE and have, for the first time, made the procedure operable on LISS data.

The procedures for atmospherically correcting both TM and IRS LISS data are now built into the processing line.

Cloud and cloud-shadow removal

Cloud and shadow masking uses reflectance recorded in TM band-3 (red) and radiance in band-6 (thermal) and ratios of these. Red response is always high; because clouds are cold, thermal response is low; and the ratio gives a range which is nearly unique to cloud. By inspecting the images, thresholds can be set. A mask is created where the image falls within the pre-specified threshold in all three cases (red, thermal and ratio values). The mask is then grown, by 8 pixels to eliminate thin undetected cloud around edges. An offset for the cloud shadow is measured by inspection. The two masks, cloud and shadow, are added together: the result showing clear ground, cloud and/or shadow is used to include / exclude sections of image in subsequent processing. All these stages have been automated in an analytical model, with only brief inspections needed to pick thresholds prior to the analysis.

Geo-registration

Geo-registration allows the images, usually corrected by the suppliers to compensate for the Earth's rotation and curvature, to be registered to the British National Grid (BNG) with a suitable output pixel size. This correction involves collecting a sample of ground control points (GCPs) on an image and, using a digitising tablet, the equivalent points on an OS 1:50 000 paper map; if an image has been registered to a map, then images which overlap are better registered image-to-image.

The image analysis system calculates a transformation between image and map, calculating the differences in scale, rotation and offset. The normal form of transformation is a polynomial model, relating x-image to X-map and y-image to Y-map, with a cross-over term which means that x on the image can change in scale as y on the image changes and *vice versa*. The more GCPs that are chosen, then potentially the more complex the model which relates x-, y-image to X-, Y-map. A first order model assumes a simple linear scale difference plus rotation and offset; second order model allows the scale to vary, gradually, increasing or decreasing across the image; more complex algorithms can accommodate highly variable changes in scale, due to sensor movement (unlikely) or terrain (more usual). These issues were all addressed in LCMGB 1990 (Fuller *et al.* 1994) and the problems might be thought of as having been solved; but LCMGB always used directly overlapping TM scenes, viewed from the exact same position in space. Due to the difficulties of obtaining images in the poor summer of 1998, LCM2000 is using some adjoining TM scenes, viewed from paths about 90 km away; it is also combining TM and LISS images with different viewing geometries and pixel sizes.

A complex series of experiments have sought the best possible combination of GCP collection strategies (image-to-map and image-to-image, varying the distribution and number of GCPs) and transformation order (first, second etc. to about seventh order). In brief, the results have shown that:

- primary corrections should be image-to-map,
- in any areas where images overlap (including same season and contrasting seasons), further GCPs should be collected image-to-corrected-image,

- overly complex polynomial models merely serve to accommodate minor misplacements in GCPs, worsening geo-registration,
- there should always be an ample excess of GCPs for the chosen model (more GCPs are needed for higher order models) to ensure that minor errors in GCP location can be detected as residual errors in the fit of the model,
- this excess of GCPs better ensures that errors are not accommodated by the transform, and that such distortions do not get incorporated into an erroneous model and a poor rectification.

Resampling

When producing an output image with pixels located relative to the reference projection, it is necessary to chose a resampling algorithm which can calculate realistic reflectances for each pixel, based on the reflectances of the surrounding pixels in the unregistered image (the true grid will be offset from the arbitrary pixel grid). The nearest neighbour resampling algorithm takes, as the name implies, the nearest pixel from the unregistered image and assigns its reflectance values, unaltered, to the pixel in the output cell. Other methods interpolate between the surrounding pixels in the unregistered image: bilinear interpolation simply mixes surrounding reflectance values in proportion to their probable contribution to the output pixel; cubic convolution uses a mathematical weighting to improve the interpolation. These methods are written up in all standard remote sensing text books (e.g. Lillesand & Keiffer 1994, Schowengerdt 1997); they are mentioned here because the methods of GCP collection, the polynomial modelling and the resampling method can very substantially influence the geometric correspondence between winter and summer images. This in turn substantially affects the image segmentation, the residual proportion of edge pixels and the final classification outputs. The resampling algorithm can exaggerate problems or, better, can resolve them: so, for example, nearest neighbour resampling can enhance within-field heterogeneity and cause artificial over segmentation; but, on single-date images, it may 'smudge' mixed boundaries less than cubic convolution. Conversely, co-registration of summer and winter scenes, using nearest neighbour resampling, can result in apparent mis-registrations, because the nearest neighbour method, when extracting reflectance values from the unregistered image, is forced to round up or down its x- and y-location to take the nearest whole pixel; thus, co-registered scenes may exhibit the cumulative effects of having half-pixel displacements in opposite x- and y-directions, giving potentially a $\sqrt{2}$ pixel difference or 1.4 pixels displacement, one image relative to the other. The conclusion, after trials of the resampling methods, has been that cubic convolution produces better results cosmetically, which are also better for segmentation purposes. The nearest neighbour results will be retained, as these will possibly be more appropriate for classification purposes, as they contain reflectance values recorded in the original data.

Illumination correction

Landscape facets are illuminated differentially according to whether they are horizontal, face the sun, or face away from the sun. In the latter case, they may have a slope sufficient for them to be completely shaded from direct solar illumination. This differential illumination across the landscape, and its consequent effects on the radiation recorded by the sensor, can be modelled using a digital terrain model (DTM) and compensated for, thereby offering corrected data based upon a theoretical horizontal surface illuminated from directly overhead. Such correction is important if facets of land surface are not to have a highly significant and perhaps dominating effect upon the results of image segmentation.

ITE have contributed financially to the operationalisation of software which was developed by Cambridge University and made pre-operational in the CLEVER-Mapping programme. The software, now fully functional, offers the option of illumination correction, prior to segmentation. The software is written as a macro within the Erdas framework, allowing the correction to be applied without time-consuming file conversions.

In previous reports, there were concerns over the use of an Ordnance Survey (OS) DTM as data are prohibitively expensive - £20k nationally; the data also come with restrictions regarding onward use of outputs, and/or a commitment to pay royalties to the OS for Land Cover data sales (including royalties on the contractual arrangements with the Consortium). EDX Engineering Inc., a US company, has constructed an alternative dataset from OS 1:63 360 maps which have passed their copyright date. This DTM costs about half the price of the OS product, without restrictions on use, nor royalties to pay. The data comprise a 50 m grid-based DTM, with integer heights in metres. In comparisons with the OS equivalent, the data are shown to have a root mean square error averaging just 2.5 m.

Terrain corrections are now routinely being used for all images: this not only compensates for the terrain but also corrects the illumination to give results equivalent to a nadir viewing angle with the sun directly overhead: the reflectances are all normalised and standardised, allowing inter-comparison within and between images. From this corrected dataset, we can build a library of spectral signatures for target classes, throughout the UK and in varying seasons.

6.2 Band combinations for analyses

TM offers 6 different spectral bands, from visible blue to mid infrared, with a spatial resolution of 28.5 m and a thermal infrared band at a coarser 120 m spatial resolution. LISS offers 3 bands from green to near infrared with a spatial resolution of 23.5 m and middle infrared band with a spatial resolution of 70 m. A summer-winter composite of two TM scenes could give 12 usable bands; a TM-LISS composite gives 10 bands; where, as in SE England, an extra TM scene is added, there are 16 available bands. Along margins, scenes overlap and several composites can contribute data with perhaps tens of bands and some locations. All such data can be captured by IGIS to record spectral response per band per date. However, there may be disadvantages in using bands which confuse rather than inform the process. Moreover, many of the processes are limited, in principle or by design, to operate on a smaller bandset. Details are given under relevant sections of this report.

6.3 Segmentation

Segmentation procedures, written by CUGD, were only operable in prototype form when the project started. The software had been demonstrated, but was only capable of segmenting scenes of about 1000 pixels square. The process ran in the Microsoft Windows environment, preventing easy use with the other processing modules, such as Erdas and Laser-Scan IGIS, which run on Unix workstations. Laser-Scan has therefore implemented a fully operational version of the segmentation software for the Unix environment: it is in effect a 'beta-test' version, suitable for ITE operations, though not yet ready to market; it has, however, been fully tested and approved by ITE for LCM2000 operations.

The segmentation consists of two separate stages, i. edge-detection to identify boundary features and ii. region growing from seed points selected to avoid the edges identified in the first stage.

There have been a number of methodological issues to address during the early production phase of LCM2000 because the procedure had never been used operationally; and certainly it has not been tested in production with the target classification, subclasses and variants of LCM2000. The issues include:

- Band selection for edge-detection and segmentation
- Choice of edge-detection methodology
- Thresholds to identify edges,
- Settings the degree of segmentation and the thresholds to achieve that,
- Post-segmentation boundary rejection and generalisation

These issues are discussed in turn.

Band selection

It is only possible to use three bands for edge-detection / segmentation. As the classification will be derived from a combination of summer and winter images it is suggested that the winter image would contribute just one band, and summer image the other two: if there are two summer scenes, then each could supply one band.

In addition to the use of pure bands, mathematical combinations of bands within the same image could offer substitutes. Vegetation indices (e.g. Lillesand & Keiffer 1994) produce a one-band image with resulting pixel values closely related to vegetation cover: in tests made under LCM2000 developments, it was found that the index might fail to separate a dense grassland from an adjoining woodland, so this option was discounted. A principle component (PC) analysis (e.g. Lillesand & Keiffer 1994) can summarise the key characteristics of a multi-spectral image in fewer bands than the original image: thus, the first PC band might encapsulate as much as 80% of the variation across the image. The first PC may show more detail than does any single band, though the distinctions between red and infrared reflectance can be more useful than use of PC-1 and PC-2. Thus the following rules are applied in band selection:

- If an image contributes two bands, they are to be red and middle infrared (IR)
- For TM, middle IR is more useful,
- For LISS, near-IR must be used as LISS middle IR has a coarser spatial resolution than the near-IR).
- If a winter image contributes just one band then it would be near-IR, which is generally the brightest in winter,
- If two summer scenes each contribute a band, then use of PC-1 should be investigated for each scene.

Edge-detection

The segmentation procedure builds polygons around 'seedpoints' which have been selected as being within a segment or a land parcel; an edge detector is used to ensure that the appropriate seedpoints are selected away from parcel-edges. The CUGD software was based on a Sobel edge detection algorithm (e.g. Mather 1987) to define the mixed boundary pixels and to avoid heterogeneous within-parcel pixels. The Sobel method gave outputs which weighted the strength of apparent edges and mosaics, giving a grey-tone image, from white where edges are strongest to black where no edges are detected. The algorithm, as implemented by LSL, sometimes formed double-lines of pixels along each side of certain linear features, leaving a void of 'non-edge' pixels which might be selected erroneously to form seed points within the linear feature. Experiments with other edge detectors showed that a filter which measure local variance, thereby measuring local heterogeneity, was an alternative to the Sobel method.

For segmentation, the resultant edge images (Sobel and variance-based) were 'density-sliced' to 'threshold' out the key edges, leaving the remainder as 'within field' pixels. Thresholding to extract 'edge' and 'non-edge' pixels from continuously variable variance images was also the subject of experimentation: the best results were those which left 'voids' in urban areas, sufficient to form seed points, while identifying most key linear features in the wider countryside (e.g. field boundaries, roads and narrow rivers) as such. The software now offers advice on thresholds, based upon examination of the histogram of output values, defining a threshold which has been shown to achieve the stated objective.

Segmentation and thresholds

The level of subdivision resulting from the segmentation procedure is in the control of the operator. By defining very tight thresholds, where small differences in reflectance values would define the spectral boundaries between segments, many very small segments can arise. In principle, this could reach the point where each pixel is itself a segment. Conversely, by allowing very loose thresholds, the parcels can be large, ultimately to the extreme of covering a whole image. The aim is to ensure a field-by-field segmentation, also separating zones within urban and suburban areas and subdividing heterogeneous semi-natural zones into meaningful segments.

The edge-image could affect the outputs, for example:

- edges and linear features themselves formed polygons;
- linear features also linked areal features one to another (e.g. two suburban areas linked into one polygon by a connecting line of mixed boundary pixels).

A wide-ranging series of trials, tested Sobel and variance edge-detectors with a range of threshold values to reject or include pixels as 'real' edges or not. The Sobel algorithm, using the 'recommended' threshold gave generally the best results, with no consistent improvement, often the contrary, when over-riding the use of the recommended value. The variance image was density sliced at varying levels to provide a range of edge-inputs for segmentation. The Sobel detector, while visually no better, was nonetheless far superior in its definition of edges/seedpoints for the subsequent segmentation; this is not surprising in that it had been 'tuned' for that purpose by the software developer.

The segmenter appeared to be fairly robust in its tolerance of different edge thresholds, with substantial changes in threshold needed to induce commensurately large changes in output segments: the fine level of control which could thus be exerted was gratifying, suggesting that the segments were real entities, not ephemeral features governed by highly sensitive input parameters.

Unfortunately, the level of spectral distinctions between the various cover types can vary according to the type: the distinction between wheat and grass might be subtle, while a single suburban zone might comprise a multitude of spectral variations of no consequence for the target classification. It is necessary to chose a level of segmentation which is adequate to separate all the target classes (and ideally all the subclasses suggested in the earlier Report) while avoiding the risk of sub-segmenting entire features (e.g. substantially subdividing fields).

During the early phase of LCM2000 it has been necessary to learn how best to cope with these problems and to choose objective criteria to select the thresholds which i. achieve a field-by-field separation, ii. grow urban and suburban areas into segments, rather that recording complex mixtures of pixels, and iii. which adequately separate what should be individual segments but which are, potentially, erroneously linked by 'linear features' (e.g. suburban areas linked by the road between them). Indeed, it was necessary to avoid the artificial growth of lines of highly mixed pixels which could form complex linear polygons, comprising various types of vegetation, often with roads and water features, which tended to spread across large tracts of image, sometimes picking up adjacent segments of similar spectral response.

Thresholds based upon reflectance differences between polygons determine how fine or coarse the segmentation is. The selection of the optimal threshold is, as above observations imply, a process fraught with some difficulty. Three different values are required, one for each band in the segmentation. A series of experiments was undertaken to test varying thresholds and to examine the results. Variable thresholds were tested in sequence: it was decided that these should be different per band, reflecting the data range to be found in the band. Determination of the range included the visual examination of histograms showing reflectances per band: there was a need to take account of long tails in a histogram distribution, where a very few pixels distorted the measured range, clipping the bottom and top 5% or so of data values to assess the range occupied by 95% of the reflectance data. An objective way of achieving this, used means $\pm$ 2 standard deviations to gather similar information.

Different proportions of the potential range of thresholds were then tested for their effect on output segmentations. While wholly objective criteria for segment-assessments could have been used, these would have relied upon the use of scarce validation data (e.g. field survey squares), too few to ensure widespread testing and, anyway, needed as an independent source of validation-data. In this instance, it has been necessary to use semi-objective procedures, where rules have been defined (as above) but the success in their application is tested by inspection. Essentially, the aim was to ensure that, as far as possible, no complex (mixed) polygons would result. To achieve the desired objectives, requires an iterative process of trial, inspection, re-analysis with altered edge-image inputs and/or segregation thresholds, re-inspection etc. until an acceptable segmentation results: the process is no different in principle to the iterative and subjective processes of training and classification used in the conventional per-pixel processes (Kershaw & Fuller, 1992), in LCMGB 1990 (Fuller, 1994) and in many other mapping exercises.

By trial and observation, it was decided that the edge-threshold determined by the software could not discernibly be improved upon; and a threshold equal approximately to one-fifth the reflectance range was appropriate for segmenting each band. These levels allocated about 95% of all pixels to segments; they gave field-by-field segmentations though often sub-segmenting some fields; this over-segmentation was necessary to prevent growth of mixed chains of edge pixels; urban areas were adequately subdivided to at least match target classes. It was decided that any simplification of unecessary sub-segmentations would be better undertaken by more 'intelligent' generalisation methods or in vector analyses, post-classification.

Trials plus inspections will ensure that these rules apply equally well to other images in the processing chain; or, if not, that thresholds are adjusted to achieve the best possible effects.

Post-segmentation generalisation and boundary rejection
It is possible, having derived an agreed segmentation, to simplify the results using spatial generalisation procedures - for example, eliminating reject pixels (not part of any segments); also dissolving small polygons and growing adjoining regions into the dissolved areas, using intelligent spatial-contextual-analyses. Though these steps may incorporate odd pixels in a polygon, they are necessary to reduce the final vector dataset to a manageable number of polygons. The CLEVER-Mapping procedure, which captures polygon reflectances from core pixels only, takes sufficient account of edge pixels in a parcel to avoid difficulties. However, because removal of edges may potentially lose useful data about landscape patterns, the information on dissolved polygons is retained; and per-pixel analyses will be used to capture information on heterogeneity.

After a series of experiments, it has been decided that:

- non-segment edge pixels should be dissolved into adjoining parcels,
- single-pixel islands should be dissolved into their surroundings
- small segments of less than nine pixels should be attached to their nearest neighbouring segments, using a spectral minimum distance rule,
- segments < 16 pixels (<1 ha - the minimum mappable unit) might also be dissolved (a decision on this will be made by comparing classifications with and without 9-15 pixel polygons),
- because the segmenter works on absolute difference, brighter fields may be sub-segmented despite a proportional difference in mean reflectances per band being <5%; these sub-segments are aggregated prior to classification using proportional difference to identify polygons to be joined,
- the generalisation procedure will keep track of dissolved and aggregated polygons, should later analyses require inspection and classification at the detailed level.

6.4 Conversion to vector format

Once an acceptable segmentation of the images has been achieved then it is necessary to create vector versions of the segments and therefore build the GIS database. This is a simple procedure of raster-to-vector conversion where the boundaries between segments with different values in the raster images are represented by vector lines. The procedure attempts to create vector polygons giving the 'land parcels' required for classification. Therefore the vector lines representing the boundary of a segment are built into a polygons. Attributes to hold information required by the land parcel can then be attached.

6.5 Training

Training for per-parcel analyses operates in a similar way to that used in conventional maximum likelihood per-pixel classification. However the procedure is much more efficient in terms of time and, therefore, it is much easier to build up a representative sample of training areas, each with a large number of pixels. Furthermore, definition of training area outlines is objectively based on the land parcels. Land parcels can be identified individually by pointing to and highlighting them, without the need to painstakingly draw the outline. The process of training is probably 20-50 times as quick, allowing a much larger sample; it also allows the operator time to define check parcels used later in validation.

The field reconnaissance data are used to steer the training process, with land parcels which are recognisable on field-maps being identified as training or validation polygons.

Another refinement built into IGIS has been the opportunity to review training areas after collecting the chosen sample. A procedure has been written in LULL (Laser-Scan's macro language) which displays 'chips' representing the remotely sensed data extracted for each training area, in classes defined by the operator, side-by-side rather than dispersed across the image. The operator can then compare and contrast training areas, re-label them to a series of different spectral subclasses, reject odd examples and alter them from/to training/validation polygons. Additional information can tell the operator the size of each polygon to give some idea of its contribution to the final classification. The operator can review the training areas in any band-combination summer, winter or a composite, to ensure that the spectral subclasses are not mixed in any bandset.

The combination of automated training area delineation plus easy review, editing and sorting into subclasses represents a huge leap in the ease and effectiveness of processing and, in truth, probably will save 2-3 days from each 'training session' for each scene pair.

6.6 IGIS 'self-training'

A consequence of using object-oriented procedures in CLEVER-Mapping is that land parcels classified from one image can be examined on an overlapping scene: thus, an unclassified scene can draw upon an area of overlap, preferably in the same year, to locate near identical segments and pick up an *a priori* class label which might be used to extrapolate the classification spatially, from one scene to another. The spectral data are derived from the scene to be classified using the *a priori* class labels; these spectral data can then be used to train the maximum likelihood classifier and statistically extrapolate the training data. Rules can be used to guarantee a minimum acceptable overlap for label-acquisition, statistical thresholds can ensure that only those labels applied with say >95% confidence are transferred; and the user can add other training polygons to supplement the number or the range of classes identified. The review procedure (displaying image chips coloured to show the range of training data) can be operated as explained above, to ensure that the training set matches standards which an operator would apply. Polygons can be retained as training examples or made into validation polygons to test the 'self-training' procedure. This process will speed up production and it will also help achieve the best possible edge-matched join along image-boundaries.

6.7 Extraction of subclass statistics

Classification requires, first, that the training data are used to interrogate the image to derive statistical measures for reflectances in each chosen band and for each spectral subclass. CLEVER-Mapping uses a shrinking procedure, when extracting raster data for polygons, to avoid edge pixels and ensure the use of 'pure' core pixels of a cover type. The shrinkage is a dynamic process whereby the required amount of shrinking is applied and, if insufficient raster data is collected, the shrinkage amount is reduced and the raster extraction is repeated. This continues until enough raster data is extracted or the shrinkage amount reaches zero. The results of this procedure, in terms of pixels extracted and shrinkage achieved, are reported and stored on the land parcel for future reference.

6.8 Classification

The polygons will be used to 'cut through' any scenes covering a polygon: thus, in areas of overlapping scenes, multiple datasets might be drawn upon. However, in 1990, it was shown that red, near infrared and middle infrared bands were most useful in discriminating classes. These will be the automatic choice for most image analyses. Only where other bands needed to separate target classes will they be called upon.

The classification procedure uses the same maximum likelihood algorithm as does per-pixel classification. However, it applies the procedure to the polygon mean statistics to select the most likely class in statistical terms. CLEVER-Mapping in IGIS records not only the most likely class and its probability but also records the probabilities for all the other subclass options. Per-pixel classifications are also being made to record the natural heterogeneity associated with polygons.

6.9 Knowledge-based correction

A combination of external data plus the class probabilities will be used in knowledge-based correction to identify land parcels classified with low confidence and/or those with classes out of their natural context (e.g. urban areas in mountains). The basic procedure was developed for per-pixel use by Groom and Fuller (1996). The additional use of class-probabilities makes it a particularly subtle yet powerful tool for post-classification correction. The procedure might also be used, after calibration against CS2000 field survey data, to reduce biases by seeking out polygons of over-estimated cover types and checking that they are classified with high levels of confidence. Similarly under-estimated types might be reviewed in polygons where they came second in likelihood to see if there is good reason to correct the classification. It is expected that such procedures will be tested and demonstrated in the next quarter of LCM2000.

6.10 Validation

Validation will be a two-stage process:

- use of the check-polygons to validate the results derived from training polygons
- use of field survey data would be used to validate or, more correctly, calibrate the GIS database

The former will assess how far the extrapolation of training data results in the attachment of correct labels to areas of known cover. The latter is to be used to derive a method for 'translation' between detailed field classes and the more generalised target classes of LCM2000, providing assessments of bias in field estimation and/or LCM2000 classification, plus calibrated correspondence measures to allow inter-comparisons, integrated statistical uses and unified cover estimates at any scale of use; also improved extrapolation of field data in a spatial context and/or better focus of LCM2000 data in a thematic context.

The principles of the per-parcel validation have been tested in the Land Cover Map of Jersey (Smith and Fuller, 1998 & in prep.) while the basic principles of the calibration exercise are described by Fuller *et al.* (1998b). The practice of such comparisons has been discussed with Peter Rothery (Statistician at ITE Monks Wood) and continued use of these has been agreed as sensible approaches. A technical Advisory group session will be devoted to validation procedures probably in late summer / early autumn, once a representative selection of completed classifications become available to test procedures.

6.11 Production

Production has involved processing 5 scenes in parallel from initial atmospheric corrections through geo-registration, illumination correction, co-registration of summer-winter data and segmentation stages, rather than the procedures implied in the initial GANNT prepared for the *Specification*. The needs to ensure that geo-registration has worked in all the combinations of overlapping TM and LISS summer and winter scenes has made this the only sensible approach. The GANNT now records progress in scene-pair equivalents, using the scene-pair processing timetable to rate progress for each scene pair in turn and adding the elements together for each scene pair to establish progress in terms of the time allocated for each task and the tasks completed. The progress has also taken into account the development work on pre-processing and its expected contribution to overall progress - very conservatively, we have allowed that the refinements will save one day per scene-pair in subsequent processing. Certain tasks are a little ahead of schedule: image searches are half-completed for winter 1998-99, ahead of schedule. Image purchases cover 50% of Britain, against an original expectation of 40% coverage. On balance, it is estimated that overall progress is perhaps a month behind schedule. But increased throughput resulting from refined pre-processing procedures will easily make good that one-month delay. The situation in LCM2000 mirrors almost exactly that in 1990 when the LCMGB spent about 3 months to complete the first scene-pair and the task finished on time and within a similar timescale to that of LCM2000.

6.12 Assessment of widespread Broad Habitats

The target classification is similar to that used in 1990, but with the changes required by the need to map Broad Habitats. At this early stage in the classification, it is impossible to report in more than just general terms upon the Broad Habitats mapped in the first scene pair to be classified.

Essentially, the classification is close to that of the LCMGB 1990 and, as such, leaves little doubt as to the potential to identify the Broad Habitats. Where LCMGB failed to produce adequate results, this has been take into account by simplification of the LCM2000 classification. Classes with a lesser expectation of accuracy have been entered as subclasses and variants, without the promise of a 90% accuracy (though efforts will be made to achieve this if possible).

Initial classifications have shown that, in principle, the Broad Habitat classification can be achieved. Overall success will start to become more clearly apparent once finalised classifications result. Validation based on check polygons will provide an early and objective measure of success.

7. CONCLUSIONS

1. Winter TM images and matching summer TM and IRS images cover most of England and Wales. A digital terrain model (DTM) for Britain will contribute to improved analyses.
2. Image searches for half the winter of 1998-9 show a poor recent acquisition rate but, even if no further scenes are added, acquisitions in winter 1997-98 give near complete winter coverage of Britain.
3. Procedural developments have involved software refinements; macros have been developed to automate the processes of image analysis; other developments have formalised and made objective the sequence of processing activities.
4. Geo-registration and resampling algorithms have demanded methodological testing. Pre-processing facilities have included automated procedures to identify and mask out clouds and cloud-shadows. Improvements in the processing stream include newly operational atmospheric corrections and illumination corrections for terrain-induced effects.
5. The pre-processing improvements have delayed the start of production mapping but will produce time savings at post-classification stages.
6. The Unix version of CLEVER-Mapping segmentation is operational.
7. Bands have been examined and selected for edge-detection and segmentation; the Sobel edge-detector is to be used; thresholds have been set to identify edges and to control the degree of segmentation; procedures have been developed and tested for post-segmentation boundary rejection and generalisation.
8. After segmentation, raster-to-vector conversion draws boundaries between segments.
9. Field reconnaissance data have been used to train the classifier, with field-mapped land parcels objectively based on the segments. The process is probably 20-50 times quicker than per-pixel training.
10. A review of training data, using display image 'chips', allows the operator to compare training areas, define spectral subclasses, reject odd examples and flag training or validation polygons.
11. The interrogation of the image to derive statistical measures for reflectances in each chosen band and for each spectral subclass uses a dynamic shrinking procedure to ensure the use of 'pure' core pixels of each cover type.
12. The classification procedure applies a maximum likelihood algorithm to the polygon mean statistics to select the most likely class in statistical terms. However, and uniquely, CLEVER-Mapping in IGIS also records the probabilities for all other subclass options. Per-pixel classifications also record the natural heterogeneity associated with polygons.
13. Validation has been designed as a two-stage process: first, using check-polygons to validate the results per-polygon; second, using field survey data to fully calibrate results, offering a 'translation' between detailed field classes and the generalised target classes of LCM2000.
14. Production has involved processing 5 scenes in parallel, from initial atmospheric corrections through geo-registration, illumination correction, co-registration of summer-winter data and segmentation stages. Preliminary classifications of each will be produced in March 1999.
15. Initial classifications of widespread Broad Habitats have shown that, in principle, the Broad Habitat classification can be achieved. Overall success will only become clearly apparent once validation, based on check polygons, provides an early measure of success.
16. The GANNT (Figure 1) records progress in scene-pair equivalents. The assessment has taken account of the development work on pre-processing and its contribution to overall progress. Some tasks are a little ahead of schedule; however, because developments have

been extended to refine methods, automate procedures, and reduce post-classification analyses, on balance, production is perhaps a month behind schedule.

17. Increased throughput, resulting from refined pre-processing procedures, will quickly make good that delay.

8. REFERENCES

Barr, C.J. 1998. *Countryside Survey 2000: Field Handbook. 3rd Draft.* Institute of Terrestrial Ecology, unpublished.

Fuller, R.M., Barr, C.J. & Wyatt, B.K. 1998a. Countryside Survey from ground and space: different perspectives, complementary results. *J. Environmental Management,* **54**, 101-126.

Fuller R.M., Gerard, F.F., Hill, R.A., Smith, G.M., Thomson, A.G., 1998b. *Countryside Survey 2000 - Part III. Module 7. Land Cover Map 2000. First Progress Report.* Unpublished ITE report to the LCM2000 Consortium.

Fuller R.M., Gerard, F.F., Hill, R.A., Smith, G.M., Thomson, A.G., 1998c. *Countryside Survey 2000 - Part III. Module 7. Land Cover Map 2000. First Interim Report incorporating the Second Quarterly Progress Report. CSLCM/Interim1.* Unpublished ITE report to the LCM2000 Consortium.

Fuller R.M., Smith, G.M. & Hill, R.A. 1998d. Land Cover Map 2000. In: *Countryside Survey 2000 Report CSJMT 7/4. Part III. Module 7.* Unpublished report to the Joint Management team of Countryside Survey 2000. 1p.

Fuller, R.M., Smith, G.M. & Hill, R.A., 1999a. *Countryside Survey 2000. Land Cover Map 2000. The Land Cover Map Of Northern Ireland. An Extension To Land Cover Map 2000. Preliminary Report.* Unpublished ITE report to the Northern Ireland Members of LCM2000 Consortium (and provided to the full LCM2000 UK Consortium).

Fuller R.M., Smith, G.M. Hill, R.A. & Sanderson, J.M. 1999b. Land Cover Map 2000. In: Hornung, M. (ed.) *Countryside Survey 2000: Second Integrated Progress Report.* Unpublished ITE report to the DETR.

Fuller, R.M., Groom, G.B. & Jones, A.R. 1994. The Land Cover Map of Great Britain: an automated classification of Landsat Thematic Mapper data. *Photogrammetric Engineering & Remote Sensing.* **60**, 553-562.

Fuller, R.M., Smith, G.M. & Hill, R.A., 1998e. *Countryside Survey 2000 Report CSJMT 8/6. Land Cover Map 2000.* Unpublished report to the Joint Management team of Countryside Survey 2000. 2p.

Fuller, R.M., Smith, G.M., Hill, R.A., Thomson, A.G. & Gerard F.F. 1998f. Module 7. Land Cover Map 2000. In: Hornung, M. (ed.) *Countryside Survey 2000 First Integrated Progress Report.* Unpublished ITE report to the DETR.

Groom, GB. & Fuller, R.M. 1996. Contextual correction: techniques for improving land cover mapping from remotely sensed images. *International Journal of Remote Sensing.* **17**, 69-89.

Helder, D., Boncyk, W. & Morfitt, R. 1997. Landsat TM memory effect characterisation and correction. *Canadian Journal of Remote Sensing,* **23**, 299-308.

Kershaw, C.D. & Fuller, R.M. 1992. Statistical problems in the discrimination of land cover from satellite images: a case study in lowland Britain. *International Journal of Remote Sensing,* 13, 3085-3104.

Liang, S., Fallah-Adl, H., Kalluri, S., JaJa, J., Kaufman, Y.J. & Townshend, J.R.G. 1997. An operational atmospheric correction algorithm for Landsat Thematic Mapper imagery over the land. *Journal of Geophysical Research,* 102, (D14), 17173-17186.

Lillesand, T.M. & Keiffer, R.W. 1994. *Remote sensing and image interpretation.* Chichester: Wiley.

Mather, P. 1987. Computer processing of remotely sensed images – and introduction. Chichester: Wiley.

Schowengerdt, R.A. 1997. *Models and methods for image processing.* London, Academic Press.

Smith, G.M., & Fuller, R.M., 1998, *CLEVER-Mapping of land cover in Jersey. Final Report.* Institute of Terrestrial Ecology Report to the States of Jersey.

Smith, G.M., & Fuller, R.M., *In prep.*, Multi-sensor, high resolution, knowledge-based per-parcel classification of land cover: an example in the Island of Jersey.

Wyatt, B.K., Greatorex-Davies, N.G., Bunce, R.G.H. Fuller, R.M. & Hill, M.O. 1993, *The comparison of land cover definitions.* Countryside 1990 Series: Volume 3. Department of the Environment, London.

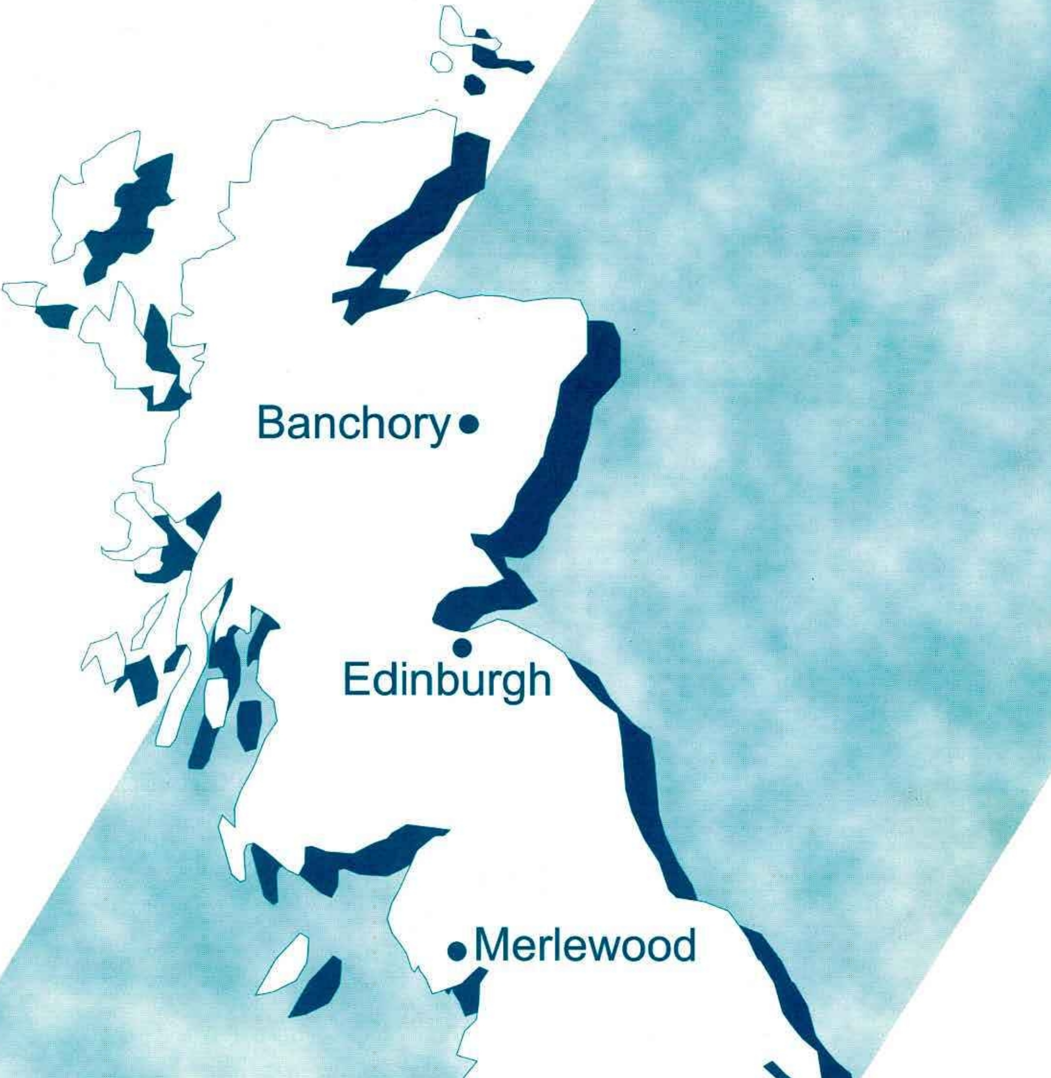

Natural Environment Research Council