

Comparing the consistency of expert land cover knowledge

Alexis Comber^{1#}, Peter Fisher^{1*}, Richard Wadsworth²

¹ Department of Geography,
University of Leicester,
Leicester, LE1 7RH, UK
* Email: pff1@le.ac.uk

² Centre for Ecology and Hydrology,
Monks Wood, Cambridgeshire,
PE28 2LS, UK
Email: rawad@ceh.ac.uk

corresponding author, Now at
ADAS Woodthorne,
Wergs Road, Wolverhampton,
WV6 8TQ, UK

Email alexis.comber@adas.co.uk
Tel. +44 (0) 1902 693129
Fax. +44 (0) 1902 743602

Abstract

Data integration can be hindered by differences in data semantics and meaning. The problem is that different data encapsulate different conceptual views of the world. Integration approaches have been developed based on modelling expert opinion of how dataset relate, rather than statistical descriptions of data correspondence. But different experts have different opinions and this is a problem in the interpretation of remotely sensed data as much as in other areas of endeavour. In work reported here, the opinions of three experts were used to examine the semantics of land cover information derived from satellite imagery. We examined the integration of two land cover datasets of the same area at different dates where the land cover mapping classes are very different, and apparently incompatible. The approach adopted involves expert opinion of how the two land cover datasets relate under a scenario of idealised relations. The work reported here compares the performance of three different experts in three different scenarios, and evaluates their performance at identifying areas of land cover change. The results show that overall they identify the same parcels as potential change areas but different experts are more reliable at identifying change in specific landscape types.

Key words: expert knowledge, consistency, change

1. Introduction

Compared with the myriad studies of the effectiveness of different computer algorithms in the differentiation of land cover types in satellite imagery (e.g. Jensen 1996; Tso and Mather 2001; Richards 1993), there are very few studies which report the confusion from operators of those algorithms. Indeed there are only a few which examine the inconsistent differentiation of classes in aerial photography. Middelkoop (1990) examined operator variation in the differentiation of land cover classes in manual interpretation of aerial photography, but was mostly concerned with how differences could be visualised. He used a group of motivated and informed postgraduate students. Pomeroy and Cline (1953) on the other hand asked professional soil surveyors to differentiate the soils of an area in a manual interpretation of aerial photography. They showed a minimum level of agreement between the two operators. Edwards and Lowell (1996) asked 9 expert foresters to differentiate forest types in simulated aerial photography, but were more concerned with the representation of the boundaries than the interpretation of the stands. McGwire (1992) extended this sort of analysis and compared operators of computer-assisted interpretation of satellite imagery of an area near the University of California, Santa Barbara. He found that certain analysts produced better results although the degree of their success was not related to greater knowledge of the study area or degree of training.

However, no one seems to have presented a systematic approach to evaluating the skill or the knowledge of an expert or compared the effectiveness of different experts. In this article we attempt to do some of this. Specifically, the objective of the work reported here was to be able to provide statements about the suitability of different experts and different types of expert knowledge (“scenarios” through the text) for specific problems. Some of the differences between experts that we are interested in addressing are:

- Does information from different experts identify different sets of results?
- It is possible to characterise each expert’s view of the interaction between the real world and its representation in the data?
- Are some experts’ “better” for specific landscape types?
- Do “good” experts have easily identifiable characteristics?

In addressing these questions the aim is to develop a basis for decision making for specific landscape questions.

The paper proceeds by introducing the background to this work: expert opinion in land cover mapping, and explores the example problem of land cover mapping in the UK (section 2). The paucity of formal approaches for comparing expert knowledge is reported in section 3 together with the approach adopted here. Section 4 describes the analytical methodology, and the results are presented in section 5 before some discussion (section 6) and conclusions (section 7).

2. UK land cover mapping

Land cover data in the UK provides a good illustration of this problem. Data exist for 1990 (LCM1990) and for 2000 (LCM2000) for the entire country and full descriptions can be found in Fuller et al. (1994; 2002). LCM1990 and LCM2000 record different land cover

features in very different ways. Table 1 describes a summary of their technical differences. Crucially the thematic classes are similar but subtly different.

Due to these differences and accuracy issues Fuller et al. (2002) issued a ‘health warning’ against comparing the datasets directly. Fuller et al. (2003) recommended that future analyses concentrated on approaches that use existing knowledge, employ the parcel structure of LCM2000 to interrogate LCMGB and generate parcel reliability information from the spectral heterogeneity meta-data. The analysis described in section 3 has these characteristics.

3. Expert approaches to relating discordant data

There are many examples in the literature of work that has integrated expert knowledge within a GIS. Typically these aim to achieve some degree of automation (e.g. Skelsey et al., 2003; Comber et al., in press b), to incorporate knowledge that may be difficult to derive through data mining (e.g. Yamada et al., 2003) to describe an approach to expert knowledge acquisition (e.g. Zhu, 1999) or to support a decision-making process with a specific objective (e.g. Alho and Kangas, 1997). In all cases the expert knowledge is compared with a known answer (truth) to the specific question being asked of the expert system.

The integration objectives and UK land cover maps present a different type of problem in that it is difficult to determine whether any *objective* geographic truth exists: it is impossible to say that one set of mapping methods and concepts (e.g. LCMGB) is better in some absolute sense than another (e.g. LCM2000). In previous work we have shown how knowledge from a single expert may be used to relate information from one dataset to another (Comber et al., 2003b, in press a, in press c) to identify inconsistency (i.e. change and error) and have shown the expert approach to be more reliable than standard statistical approaches such as discriminant analysis (Comber et al., submitted). However in the process of the research we collected information relating the LCMGB target classes and LCM2000 broad habitats from 2 other experts and asked each of the 3 experts to consider 3 different scenarios. The experts were all chosen to be firstly well informed individuals who are knowledgeable about land cover as a concept, about land cover mapping from remotely sensed data, and about the uses to which that information is put. They were chosen to be representative of individuals who work in the generation of landcover maps (the Producer), in the subsequent distribution of the information (the Distributor), and the uses of the information (the User). Note that these names are used deliberately to caricature the people concerned, and it is not to say that they share their opinions with all experts who might characterise themselves as users, producers or distributors. Each was asked to consider three different scenarios. They were asked to consider:

- ‘Semantic’ representing the expert understanding of the links between the semantics of the two classifications, without considering aspects such as common spectral confusions or known directions of change;
- ‘Technical’ describing the expert’s heuristic knowledge of where spectral confusion and other problems may occur;
- ‘Change’ to represent the expert’s opinion of the changes between pairs of LCMGB and LCM2000 classes, that one class in 1990 would change into another class in 2000.

In each scenario, they prepared a pair-wise comparison in a table, recording:

- +1 for those situations where they were sure that the LCM2000 and LCMGB classes would overlap (expected);
- 0 for those situations where they were uncertain as to whether the classes would overlap (uncertain); and
- -1 when they were sure that the classes would not overlap (unexpected).

This gives a matrix a total of 9 sets of expert knowledge stored conveniently as pair-wise comparisons in Look-Up Tables with LCM2000 broad habitats as rows and LCMGB target classes as columns.

In the current paper we wish to explore similarities and differences amongst the experts in terms of how they view the landscape and to assess their relative suitability for different landscape questions.

4. Methods

The differences amongst experts were analysed using 2 approaches. First by applying a hierarchical clustering algorithm to the expert pair-wise relations between LCM2000 broad habitats and LCMGB target classes. Second by analysing how well each expert predicts actual changes identified as part of a field survey.

4.1 Hierarchical clustering

The LCM2000 broad habitats were grouped using the expert scores (+1, 0 or -1) on the LCMGB classes with a hierarchical clustering algorithm identifying between-groups linkage and squared Euclidean distance. The objective was to characterise the expert's view of the landscape. In this context, the dendrograms produced by the cluster analysis illustrate the extent to which the expert believes that it is possible to use one of the classifications to say anything about the other.

4.2 Predicting actual change

The land cover data has been explored for a 100km by 100km area in the UK, Ordnance Survey tile SK which includes *Arable*, *Pastoral* and *Marginal Upland* zones of the 6 environmental zones derived from the ITE Land Classification Great Britain (Bunce et al., 1996).

Previous work developed and applied a single semantic LUT to identifying land cover change between LCMGB and LCM2000 (Comber et al., in press a, in press c). This LUT was constructed using a scenario of 'idealised semantic relations'. Having extracted the expert's opinion of the semantic relations, the method goes through the following stages:

1. All LCMGB pixels in a parcel are extracted and coded based on the expert's LUT. The sum for each code is calculated, characterising each parcel.
2. A similar process is repeated using parcel spectral heterogeneity attributes and the data producers description of spectral overlap, generating a second parcel characterisation.

3. Changes in the characterisations give a measure of parcel inconsistency due to data errors or to land cover change.
4. Inconsistency scores were normalised into a “belief” in change and were combined using Dempster-Shafer.
5. A sample of parcels was visited. Parcel class and whether change had occurred since 1990 (in the opinion of the surveyor) recorded.
6. The extent to which expert belief in change partitioned ‘change’ and ‘no change’ field data parameterised the expert reliability.

The headline result from the earlier work is that inconsistency between LCMGB and LCM2000 was identified in 100% of the parcels, with 41% of the inconsistency being attributable to change and 59% attributable to error in either dataset (Comber et al., in press a).

However, as noted in section 3 opinions of other experts were sought, and other LUTs exist. Using the LUTs of the different experts and scenarios it is possible to generate a series of beliefs in inconsistency for each parcel in the manner summarised above. These were combined using a standard Dempster-Shafer approach. The different beliefs and combinations of belief were analysed in terms of how well they partitioned parcels that had been visited in the field.

Some 345 parcels were visited and assessments were made of whether the land cover matched LCM2000 and it had changed since 1990. The data was assembled from the different sets used to assess earlier stages in the development of this method (see Comber et al., in press a; in press c) and are therefore not a random sample: approximately half were identified as possible locales of change because of their large vectors of inconsistency between LCMGB and LCM2000 and half were selected at random. The expert belief in change for the field data parcels was analysed sequentially for each expert: *Semantic*, *Semantic and Technical* and *Semantic, Technical and Change*. The belief was combined using Dempster-Shafer as described in Comber et al. (in press a). A threshold of combined belief greater than 0.95 was used to select candidate change parcels.

5. Results

5.1 Hierarchical clustering

The dendrograms for the semantic relations are shown in detail in Figure 1, the technical relations in Figure 2 and change relations in Figure 3.

The expert LUT's encapsulate how the expert thinks any pair of concepts (one from each classification system) relate to each other, however, it can also be interrogated to provide information about how well classes within one system can be distinguished using the other system. If any two rows on the LUT had an identical sequence of values this shows that the expert does not think those two classes can be distinguished using the classification system represented in the columns. Consider a trivial example, one system has two classes deciduous and coniferous trees while the other system has two classes tall and short woody vegetation. In this case the expert might consider that coniferous and deciduous trees are semantically

distinct but that the second classification (tall and short woody vegetation) does not provide any information to distinguish between them. As the LUT has more than 20 rows and columns it can be difficult to compare and then visualise how similar different rows are. Dendrograms from hierarchical clustering are a convenient tool for visualising these tables.

The upper dendrogram in Figure 1 was constructed from the User's Semantic LUT. The dendrogram clearly shows that this expert does not believe that the LCMGB allows the three grassland classes of calcareous, acid and neutral grass (7.1, 8.1 & 6.1) to be distinguished. The User also sees it as being relatively difficult to distinguish water from sea (13.1 & 22.1) dense and open dwarf shrub heath (10.1 & 10.2) and suburban and urban developments (17.1 & 17.2). In contrast the Distributor, believes that the grasslands (6.1, 7.1, 8.1) can be distinguished from each other (using LCMGB) albeit not very strongly as they are closely clustered in the dendrogram, but, this expert does not believe that the three arable classes (4.1, 4.2 & 4.3) can be distinguished. Finally the Producer, agrees with the distributor that the arable classes are not distinguishable and with the user that the neutral and calcareous grasslands are indistinguishable.

It is less easy to make such general statements from the two other scenarios (Figures 2 and 3). In these cases, whilst some singletons are identified, generally the cluster analysis reveals perceptual hierarchies in the classification scheme, or confusion about what the classes mean. However, it is possible to explore differences amongst experts by examining whether some of the more common land classes have similar neighbours in each of the dendrograms.

Generally despite the differences overall in the dendrograms the experts are in reasonable agreements for the most common classes. For example for arable cereals (4.1) the User links it semantically with horticulture, technically with setaside grass (5.2) and with non-rotational arable (4.3) for change. The Supplier links it semantically and technically with the other 2 arable classes, but less strongly and with improved grassland (5.1) for change. The Producer links it semantically with the other 2 arable, with non-rotational arable technically and with the other 2 arable for change. In contrast a land cover like Broadleaved woodland (1.1) is consistently associated with coniferous woodland (2.1) by the User, but the Supplier associates it weakly with lowland classes semantically, with heaths technically and with coniferous woodland for change. The Producer links it with nothing semantically, with horticulture and suburban development (17.1) technically (and relatively strongly) and with a group of upland and semi-natural classes for change.

In short, there are a number of points of similarity in the tables, but there are also important areas of disagreement.

5.2 Predicting actual change

The extent to which the different expert beliefs in inconsistency partition the change data is shown in Table 2 for the SK area and for 3 zones within that area. In the entire study area 345 parcels were visited and 57 changes were found. The Arable zone included 132 of the visited parcels with 18 changes, the *Pastoral* zone 107 parcels with 26 changes and the *Marginal Upland* zone had 104 visited parcels; with 13 changes.

Overall for the whole SK test area and for the Pastoral and Arable landscapes, the results improve as more evidence is introduced, that is more actual changes are partitioned from the field data. Whilst the difference between the experts is small, the Producer is the expert who best partitions the visited parcel information. In the Marginal Upland landscape, however results deteriorate as more evidence is introduced with only the User maintaining consistent levels of reliability. In more heterogeneous landscapes it may be preferable to use evidence from just the Semantic relations. The impact of heterogeneous landscapes in such semi-natural areas on the results and the cost of additional evidence (increasing the errors of commission) are discussed below.

The experts appear to identify the components of inconsistency (such as change) with equal reliability. Therefore it is important to determine whether the same sets of change parcels or subsets related to different broad habitats or other landscape processes were being identified. The parcels identified as 'change' by the combined evidence of each expert were compared. The overall findings are summarised in Figure 4 which shows a Venn diagram of the extent to which beliefs from the different experts correspond. The agreement amongst the experts is 84% ($(129 + 162) / 345$). Thus in the vast majority of situations the different experts are identifying the same parcels.

6. Discussion

The algorithm used to create the dendrogram is constrained to eventually link all cases however distinct they are. Despite this there are noticeable qualitative differences between the dendrograms, the User has a much more clearly defined hierarchy of relationships than the other two experts. In particular the Producer has a very "bushy" dendrograms with many classes converging at a similar level. Interestingly with two perfect classifications (each class perfectly distinct and with no overlaps) then the dendrogram would be a perfect bush (all classes would join at the same level of similarity). It is possible to test the significance of the apparent clusters within the dendrogram by examining the partition coefficient or entropy when using a fuzzy clustering algorithm such as k-mean. Such a procedure confirms the visual impression that all the User's LUT's partition the 200 broad habitats in terms of the 1990 target classes effectively into a few (5 to 8) classes with only the occasional "singleton" (a class with a single example), while for the Producer the optimum classification (maximum partition, minimum entropy) occurs with many more classes, many of which are singletons. The dendrograms can therefore be seen as underpinning the roles of the different experts: at one extreme the producer (a remote sensing expert) has a particular view of the process of classification (distinct, non-overlapping, and complete) that reflects their view of the landscape; at the other, the user may be more aware of some of the operational problems involved in the two datasets, and some of the ambiguities in either classification.

It is, of course, also possible to ask an Expert how they think concepts within a single classification relate to each other; that is produce a LUT where the rows and the columns hold the same classes. In a "perfect" classifications the LUT would simply have "1's" on the leading diagonal and "-1's" everywhere else. In fact this rarely happens and Figure 5 illustrates the LUT produced by the User for the LCMGB only. On this you can see that the

User thinks that the scrub and deciduous classes are effectively identical and that the urban and suburban classes are very closely related. In general, although the User sees all the upland classes as being closely related, there is a strong hierarchical structure in their view of the classification.

Comparing the different proportions of the field data and success in the use of the tables in predicting change (Table 2), showed the Producer and the Distributor to be more reliable than the User overall, and specifically in Arable and in Pastoral landscapes: they identified more change parcels with greater reliability than the User in the homogenous landscapes. The User identified more change areas in the ecologically heterogeneous, semi-natural Marginal Upland landscapes. In this instance, the greater uncertainty indicated by the User reflected their view of the definitional and conceptual uncertainty present in many of the Upland land cover class types.

It is noteworthy that in the Marginal Upland zone, the additional information provided by the Technical and Change LUTs for each expert does not improve the identification of change. In part this is due to the nature of the Dempster-Shafer theory of evidence. Additional evidence pushes the belief in a hypothesis towards the tails of the distribution. That is, small and on their own insignificant pieces of information, gain in significance towards the endorsement or refutation of a hypothesis, when combined with other evidence. Essentially Dempster-Shafer assesses the belief in a hypothesis as expressed by the probability that the proposition A is *provable* given the evidence (Comber et al., 2004). However, the inherent heterogeneity of the landscape is problematic for land cover mapping either using a by-pixel classification (LCMGB) or a segmentation approach (LCM2000). In the former there may be too much short range variation in the mapped land cover, whilst the latter may represent an over-simplification of the landscape.

7. Conclusions

The questions we posed at the beginning of this work were not concerned with whether the proposed methodology was robust. Rather the issue was to determine how dependent the results were on the quality of the expert and differences between them. The results allow the following conclusions to be made:

- In this instance, dendrograms provide a useful tool to visualise where experts believe that class-to-class similarities and overlap exist and gives an overview of the way they think reality is partitioned by different data;
- Expert LUTs incorporate a wider landscape view than a confusion matrix between the data elements of LCMGB and LCM2000;
- Overall, the experts are very nearly as reliable as each other and the agreement between them is 84%;
- However, the information provided by different experts is more reliable for analysis of different landscape zones;


The implications of these conclusions are that expert selection is not important if the study area is large with a range of different landscape types. However for specific questions pertaining to specific landscape types, experts have been shown not to be equally knowledgeable. Therefore, as they have different opinions, expert selection is important in

specific landscapes. It is likely that this is especially important when the cover types are more heterogeneous where the distinctions between different land covers are difficult to determine spectrally, botanically and in terms of bio-geographic niche.

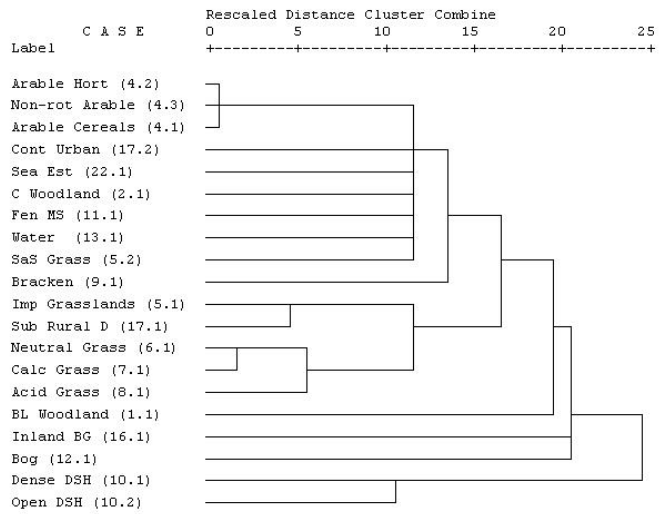
Finally, this work has benefited from the considerable object level metadata that is available for LCM2000, but such metadata is rare in other data products. Therefore, as part of communicating data meaning in a wider sense, current metadata reporting should be extended to include expert mappings of how the concepts within a single dataset relate to each other.

Acknowledgements

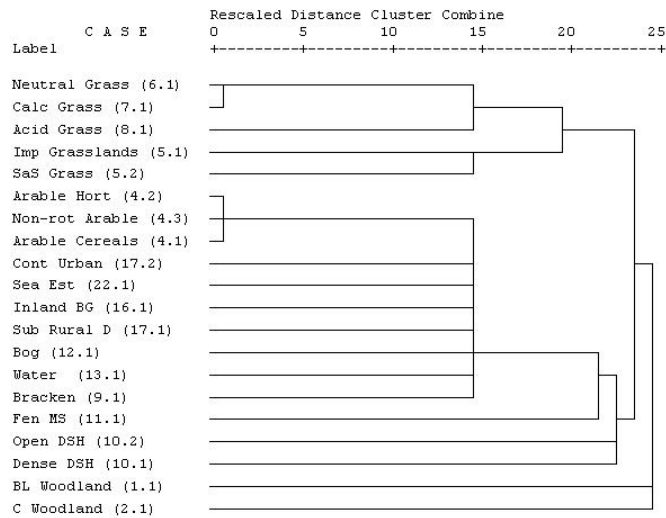
Ideas behind this paper were presented at ISSDQ 2004, Brück, Austria and describes work done within the REVIGIS project funded by the European Commission, Project Number IST-1999-14189. We would like to thank our collaborators, especially Andrew Frank, Robert Jeansoulin, Geoff Smith, Alfred Stein, Mike Worboys and Barry Wyatt.



User



Distributor



Producer

Figure 1. Dendrograms of expert *Semantic* relations showing clusters of LCM2000 broad habitat classes grouped on expert relations with LCMGB target classes.




Figure 2. Dendrograms of expert *Technical* relations showing clusters of LCM2000 broad habitat classes grouped on expert relations with LCMGB target classes.




Figure 3. Dendrograms of expert *Change* relations showing clusters of LCM2000 broad habitat classes grouped on expert relations with LCMGB target classes.

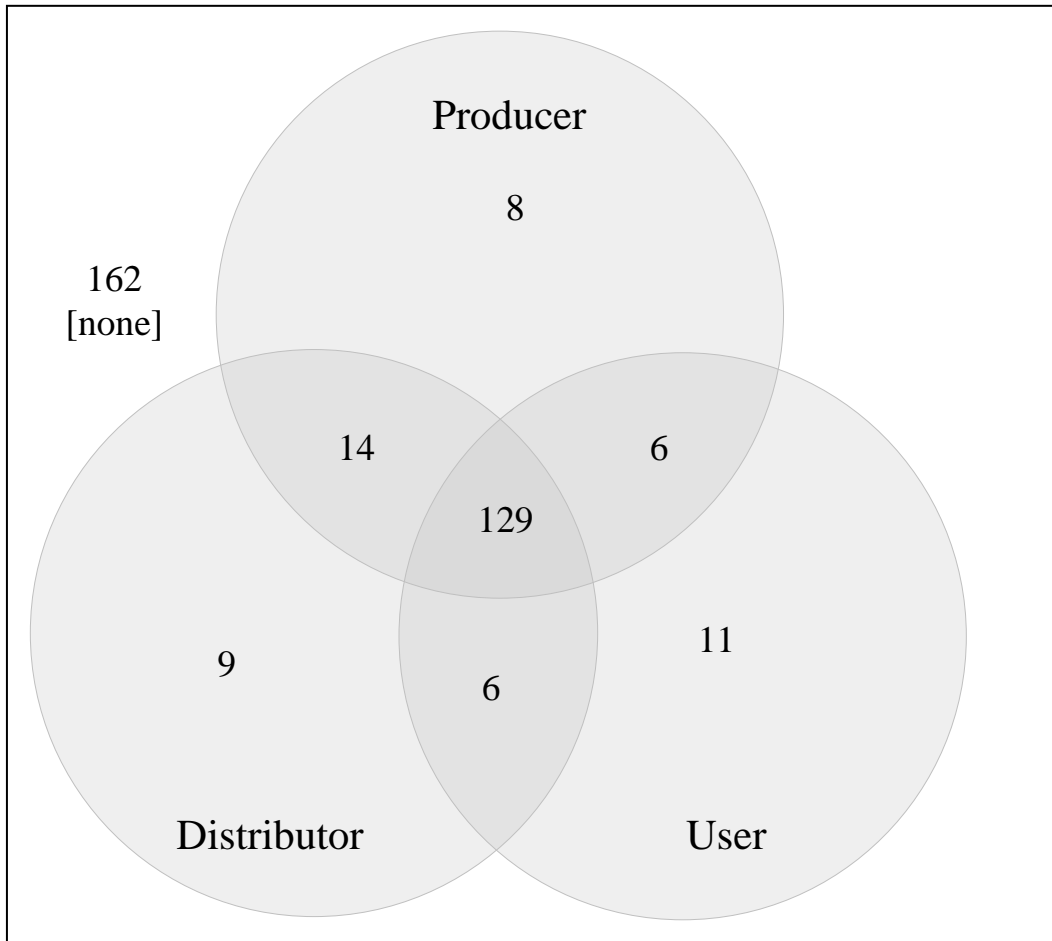


Figure 4. A Venn diagram of the subsets parcels with a combined belief ≥ 0.95 as identified by the different experts.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
1	1	1	0	0	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
2	1	1	0	0	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	1	-1
3	0	0	1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	0	-1	-1
4	0	0	1	1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	0	-1	-1	-1	-1	-1	-1	-1	0	-1
5	-1	-1	-1	-1	1	0	0	1	0	0	0	0	1	-1	-1	-1	0	-1	-1	-1	-1	-1	-1	0	1
6	-1	-1	-1	-1	0	1	1	0	0	0	0	-1	0	-1	-1	-1	-1	-1	0	-1	-1	-1	-1	-1	0
7	-1	-1	-1	-1	0	1	1	1	1	0	0	-1	0	-1	-1	-1	0	0	0	-1	-1	0	0	0	0
8	-1	-1	-1	1	1	0	1	1	1	0	0	0	0	-1	-1	-1	0	-1	0	-1	-1	0	0	1	0
9	-1	-1	-1	-1	0	0	1	1	1	1	1	0	1	-1	-1	0	1	-1	-1	-1	-1	-1	-1	1	1
10	-1	-1	-1	-1	0	0	0	0	1	1	1	0	1	-1	-1	0	0	-1	-1	-1	-1	-1	-1	0	1
11	-1	-1	-1	-1	0	0	0	0	1	1	1	1	1	-1	-1	-1	1	-1	-1	-1	-1	0	-1	1	1
12	-1	-1	-1	-1	0	-1	-1	0	0	0	1	1	1	0	0	0	0	-1	-1	-1	-1	-1	-1	0	0
13	-1	-1	-1	-1	1	0	0	0	1	1	1	1	1	-1	-1	-1	0	-1	-1	-1	-1	-1	-1	0	0
14	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	0	-1	1	1	0	-1	-1	-1	0	-1	0	-1	-1	-1
15	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	0	-1	1	1	0	-1	-1	-1	0	-1	0	-1	-1	-1
16	-1	-1	-1	-1	-1	-1	-1	-1	0	0	-1	0	-1	0	0	1	0	-1	-1	-1	-1	1	-1	-1	-1
17	-1	1	-1	0	0	-1	0	0	1	0	1	0	0	-1	-1	0	1	-1	-1	-1	-1	-1	-1	1	-1
18	-1	-1	-1	-1	-1	-1	0	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	1	-1	-1	-1	-1	-1	-1
19	-1	-1	-1	-1	-1	0	0	0	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	1	1	1	1	1	-1	-1
20	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	0	0	-1	-1	-1	1	1	1	-1	1	-1	-1
21	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	1	1	-1	1	-1	-1
22	-1	-1	-1	-1	-1	-1	0	0	-1	-1	0	-1	-1	0	0	1	-1	-1	1	-1	-1	1	-1	-1	-1
23	-1	-1	0	-1	-1	-1	0	0	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	1	1	-1	1	-1	-1
24	-1	1	-1	0	0	-1	0	1	1	0	1	0	0	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	1	1
25	-1	-1	-1	-1	1	0	0	0	1	1	1	0	0	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	1

Legend

Sea	1	Mown / Grazed Turf	6	Dense Shrub Moor	11	Conifer	16	Urban Development	21
Inland Water	2	Pasture / meadow / amenity grass	7	Bracken	12	Upland bog	17	Recently felled	22
Coastal bare	3	Rough / Marsh Grass	8	Dense Shrub Heath	13	Tilled Land	18	Inland Bare Ground	23
Saltmarsh	4	Moorland Grass	9	Scrub / Orchard	14	Ruderal Weed	19	Lowland bog	24
Grass Heath	5	Open Shrub Moor	10	Deciduous Woodland	15	Suburban / Rural Development	20	Open Shrub Heath	25

Figure 5. User opinion of how LCMGB classes relate in terms of being expected (1) , uncertain (0) and unexpected (-1).

	LCM1990	LCM2000
Input Data	Multi-date composite Landsat TM images	
Pre-Processing		Image sharpening, Cloud detection, Atmospheric correction, Topographic correction, Image segmentation
Classification	Per-pixel	Per-parcel and per-pixel
Post-Processing	Simple knowledge-based corrections using masks	Knowledge-based corrections using within-and between-parcel context and ancillary data
Outputs	Class per-pixel	Class per-parcel Top 5 classes Per-pixel classes
Objectives	Demonstration of Satellite Imagery	Address post-Rio policy objectives (e.g. Habitats Directives, BAP's)

Table 1. A comparison of key characteristics of the Land Cover Maps 1990 and 2000 of the United Kingdom

	Overall			Marginal Upland			Pastoral			Arable		
	D	P	U	D	P	U	D	P	U	D	P	U
Semantic												
Changes	0.61	0.61	0.56	0.54	0.54	0.62	0.77	0.77	0.69	0.44	0.44	0.33
Omission	0.39	0.39	0.44	0.46	0.46	0.38	0.23	0.23	0.31	0.56	0.56	0.67
Commission	0.29	0.29	0.28	0.19	0.21	0.22	0.38	0.37	0.33	0.30	0.30	0.30
Semantic and Technical												
Changes	0.65	0.68	0.65	0.54	0.54	0.54	0.81	0.81	0.81	0.50	0.61	0.61
Omission	0.35	0.32	0.35	0.46	0.46	0.46	0.19	0.19	0.19	0.50	0.39	0.39
Commission	0.38	0.38	0.36	0.25	0.31	0.31	0.46	0.43	0.43	0.41	0.40	0.40
Semantic, Technical and Change												
Changes	0.68	0.74	0.70	0.46	0.54	0.62	0.81	0.85	0.77	0.67	0.72	0.67
Omission	0.32	0.26	0.30	0.54	0.46	0.38	0.19	0.15	0.23	0.33	0.28	0.33
Commission	0.41	0.40	0.39	0.31	0.31	0.34	0.49	0.46	0.43	0.44	0.43	0.39

Table 2. The proportions of parcels identified as ‘change’ and ‘no change’ in the field, correctly partitioned by combinations of evidence from the Distributor (D), Producer (P) and the User (U). The ‘best’ expert results (i.e. most change identified) are highlighted in bold.