Visual Exploration of Uncertainty in Remote-sensing Classification

Frans J.M. van der Wel
Utrecht University, Faculty of Geographical Sciences, Cartography Section
From April 1997: Royal Netherlands Meteorological Institute (KNMI)
P.O. Box 201, 3730 AE De Bilt, The Netherlands
e-mail: `welvdf@knmi.nl`

Linda C. van der Gaag
Utrecht University, Department of Computer Science,
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
e-mail: `linda@cs.ruu.nl`

Ben G.H. Gorte
International Institute for Aerospace Survey and Earth Sciences (ITC)
P.O. Box 6, 7500 AA Enschede, The Netherlands
e-mail: `ben@itc.nl`

Abstract

Exploratory analysis of remotely-sensed data aims at acquiring insight as to the stability of possible classifications of these data and their information value for specific applications. For this purpose, knowledge of the uncertainties underlying these classifications is imperative. In this paper, we introduce various measures that summarise for a classification, in a single number per pixel, the distribution and extent of the uncertainties involved. Since exploratory analysis needs effective ways of conveying information to the user, we in addition address various ways of cartographic visualisation of uncertainty.

1   Introduction

The ability of space- and airborne instruments to measure the amount of electromagnetic radiation reflected and emitted by the Earth's surface has proved to be valuable for the understanding of our environment, as it provides for an overwhelming flow of data on the appearance and condition of our planet. The data yielded by remote sensing can be subjected to various types of computer-assisted manipulation, to arrive at derived data sets tailored to different types of application. Computer-assisted classification of remotely-sensed data into qualitative classes, for example, is useful for extracting information that can be exploited for cartographic purposes, such as in the generation of thematic maps of land cover types. Denègre (1994) elaborates on the cartographic meaning of remotely-sensed data acquired by Earth observation satellites such as the Landsat series and SPOT series.

For a proper cartographic application, the fitness for use of a set of remotely-sensed data needs be assessed. The practicability of the data and their classification can be established by means of an *accuracy assessment* procedure. An error matrix is created for the classification by matching a random sample and its counterpart from a reference data set representing the actual environment. From such a matrix, various statistics are derived (Janssen and Van der Wel, 1994). Accuracy

assessment based on an error matrix, however, has several drawbacks. Among these is the non-spatial and general character of a global statement like 95% accuracy for an entire classification. Also, the random selection of sample data disregards the occurrence of differences in stability of class assignments. Moreover, accuracy assessment is a time-consuming and cost-intensive process. As a consequence, it is easily omitted which, of course, is undesirable and may lead to the use of data that are unfit for the application at hand.

For assessing the fitness for use of a set of remotely-sensed data, however, accuracy is not the only consideration. More in general, the phrase *data quality* is used to refer to the extent to which the characteristics of the data meet the requirements of the application aimed at by the user. A high quality indicates a relatively high information value for the considered application or, stated in another way, a good fitness for use. *Uncertainty* is a key issue in quality assessment and, therefore, in the assessment of fitness for use of a data set. Unwin (1995) presents an overview of the issue of uncertainty in classification.

For investigating uncertainty, effective measures need to be designed. To this end, it is relevant to consider the purpose to which these measures are to be employed. In this paper, we focus on an exploratory perspective. *Exploratory analysis* of a set of remotely-sensed data aims at acquiring insight in the stability of various possible classifications of these data. For this purpose, knowledge of the uncertainties underlying these classifications is imperative. As in exploratory analysis, classification is an iterative process, it needs not only measures for assessing the uncertainty in a classification but also effective ways to convey this information to the user. Visualisation is generally considered a useful means of communication of potentially relevant information. In this paper, we present a class of measures of uncertainty tailored to the purpose of exploratory analysis of remotely-sensed data and address various ways of cartographic visualisation of uncertainty.

The paper is organised as follows. In the next section, we discuss the issue of uncertainty in the classification of remotely-sensed data. Then, the current state-of-the-art in cartographic visualisation of uncertainty is reviewed. Subsequently, a class of highly informative measures building on the notion of weighted uncertainty is introduced. These measures allow for quite convincing visualisations, as is shown in the last section using examples from the CAMOTIUS[1] project.

2   Uncertainty in remote-sensing classifications

Discrete remote-sensing classification, by its very nature, neglects the fuzzy character of our environment and, as a consequence, introduces *uncertainty* in class assignments. This uncertainty is typically further propagated throughout subsequent stages of post-classification editing, intended to produce a cartographically attractive product: post-classification editing may, for example, induce an unpredictable smoothing of potentially significant phenomena. Lunetta *et al.* (1991) give an overview of the sources of errors and uncertainties in remote-sensing classification.

The uncertainty that is introduced during classification of a set of remotely-sensed data is characterised by the *probability vectors* that are yielded as a byproduct by most probabilistic classification procedures (Goodchild *et al.*, 1992). In a *maximum a posteriori classification*, for every pixel in the data, a vector of probabilities is calculated that specifies for each recognised class its probability of being the true class. If a pixel is viewed as a statistical variable C that has one of the classes C_i, $i = 1, \ldots, n$, for its value, then the vector computed for this pixel specifies the posterior probabilities:

$$\Pr(C = C_i \mid x)$$

for $i = 1, \ldots, n$, where x denotes the available data. The probability vectors yielded by classification reflect the differences in uncertainty in the resulting classification and may be considered indicative of dubious classifications, of mixed pixels, of heterogeneous classes, or of fuzzy boundaries between classes (Goodchild *et al.*, 1992).

Building on probability vectors, various *measures of uncertainty* may be used. An example of such a measure is the maximum probability appearing in a probability vector. The value of this maximum probability is an expression of the strength of the class assignment and of possible confusion with other classes. A high maximum probability, such as 0.95, raises few doubts about the classification of the pixel at hand. A lower maximum probability, such as 0.50, however, could very well indicate a dubious class assignment. The difference between the maximum and second ranking probabilities in the vector may then serve to clarify this doubiousness as it provides information about the certainty of the most probable class. Note that these measures of uncertainty do not provide the user with a global indication of the practicability of a set of remotely-sensed data and their classification as accuracy measures do.

3   Visualisation of uncertainty

Cartographers have traditionally been concerned with quality, as their integrity is reflected in the appearance of their maps. This implicit "stamp of approval", as Morrison (Guptill and Morrison, 1995) calls it, has gradually lost its significance with the advent of computer technology and the dissemination of mapping abilities to application specialists rather than cartographers. On-screen maps are more transitory than paper maps. Also, several different on-screen maps are easily generated as different visualisations of the same data, which further contributes to their ephemeral character. Traditional visualisation by — possibly insignificant — colours for a remote-sensing classification may not be sufficient for the application aimed at by the user. Some visualisation of data quality is called for. In fact, graphical representation of the quality of remotely-sensed data has recently evolved as a challenging cartographic activity: several interesting representation methods are being explored, building on the possibilities offered by computer technology (Beard *et al.*, 1991).

Using the visualisation model of DiBiase (DiBiase *et al.*, 1992), we review several techniques for the visualisation of uncertainty, tailored to the support of exploratory analysis of remotely-sensed data. A distinction is made between static and dynamic visualisation. In static visualisation graphic variables are applied for capturing uncertainty into a spatial pattern; dynamic visualisation adds time as an extra dimension to the visualisation. Conveying uncertainty by means of graphic variables in a static map is firmly rooted in cartographic tradition and can, therefore, benefit from the existence of a sound theoretical framework (Bertin, 1983). Dynamic cartography, on the other hand, is still evolving: its — sometimes — immature concepts have yet to deal with considerations of feasibility, both with regard to technology and perception. The review in this section is by no means exhaustive, but is intended to give an impression of some of the more feasible ways to convey uncertainty.

3.1   Static visualisation

A starting point for the visualisation of uncertainty by graphic variables is provided by the theory of semiology of graphics of Bertin (1983). In his theory, Bertin distinguishes six graphic variables: colour hue, orientation, shape, size, texture or grain, and value, each of which can be attached to points, lines, and areas in a map. In addition to these six variables, other graphic variables have been suggested in recent years, such as structure or pattern arrangement (Muehrcke and Muehrcke, 1992) and, with respect to the representation of uncertainty, colour saturation (MacEachren, 1992). The graphic variables that are most suitable for conveying the quantitative character of uncertainty are value (grey-scale maps), colour saturation (bivariate maps), colour hue (associative ranking maps) and a combination of the latter two variables (dichotomy maps) (Van der Wel *et al.*, 1994).

Probably the simplest visualisation of the uncertainties underlying a remote-sensing classification is in a *grey-scale map* depicting the maximum probability per pixel as a colour value. The colour values used, typically range from black for a low maximum probability to white for a high one. An example grey-scale map is shown in Figure 1. It will be evident that also grey-scale
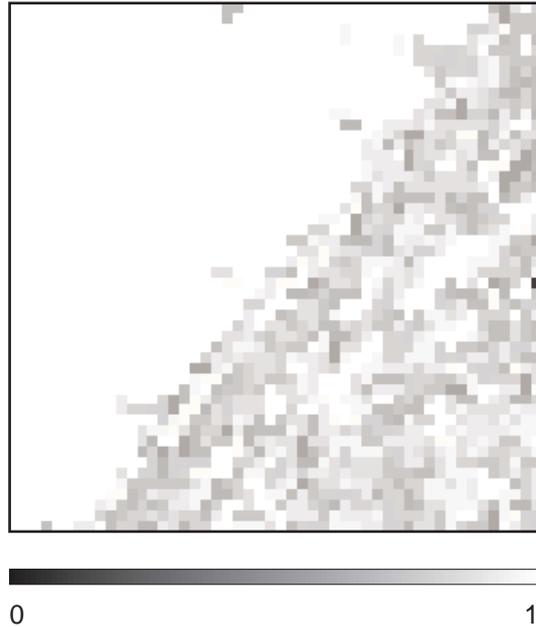

Figure 1: Maximum probabilities depicted in a grey-scale map.

maps can be generated that represent for each single qualitative class this class' probability per pixel. The information content of such maps, however, is rather restricted. This can be improved to some extent by the simultaneous representation of the classification itself in a separate map (Foody and Curran, 1994; MacEachren, 1992; Van der Wel *et al.*, 1994). Relating uncertainty to classes by means of a pair of static maps, however, is quite cumbersome.

Computer technology enables straightforward integration of the representations of qualitative classes and quantitative uncertainty in a classification by means of colour hue and colour saturation respectively. The resulting *bivariate maps* are highly informative, reflecting the strengths with which the classes have been assigned to the various pixels. The high information content, however, produces a complex interpretation key. Moreover, colours with low saturation may be difficult to distinguish, thus contributing to confusion in the interpretation (Brown and Van Elzakker, 1993).

By nature, colour hue fails to give an impression of order as required for conveying the quantitative character of uncertainty. It can nevertheless be applied for this purpose by taking advantage of the subjective associations by which people deal with a sequence of hues. An example is the *traffic light principle* illustrated in Figure 2. Red, yellow, and green convey prohibition, alertness, and permission, respectively. From the probability vectors yielded by classification, an *associative ranking* map can be generated in which pixels having a high maximum probability in their vector and a large difference of this probability with the second ranking one are displayed in green and pixels with a low maximum probability and a small difference with the second ranking probability are displayed in red. Yellow spots in the map indicate ambiguous pixels, that is, pixels with a moderately high maximum probability in their vector but a hardly convincing difference value.

The associative capability of colour hue can also be exploited to convey deviation from a user-defined threshold value. Suppose that a minimal probability of 0.50 for class assignment to a pixel is considered critical for some application. Pixels with a maximum probability in their vector lower than or equal to 0.50 then are characterised by a blue colour hue, whereas pixels with a maximum probability above the threshold value are depicted in red, with the colours blue and red
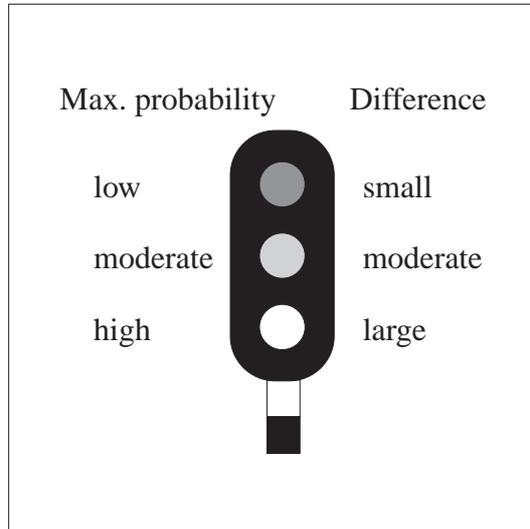
Figure 2: The principle of an associative ranking map.

chosen so as to refer to an association with temperature ranges. The amount of deviation from the threshold value given by the user can in addition be represented by colour saturation. Figure 3 illustrates the basic idea of such a *dichotomy map*. The usefulness of these maps strongly depends
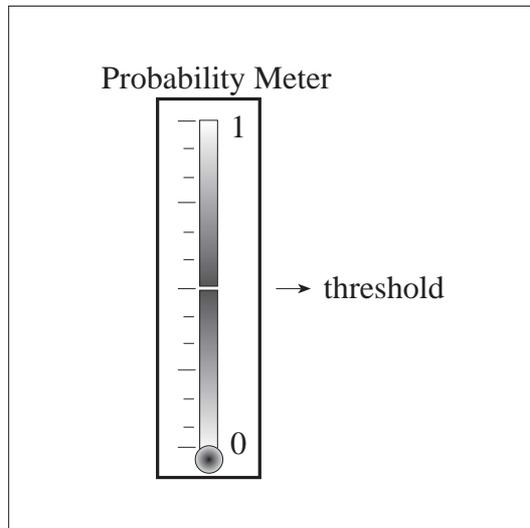

Figure 3: The basic idea of a dichotomy map.

on the acquaintance of the user with the selected association.

In addition to using graphic variables in a two-dimensional map, uncertainty may be represented by introducing a third dimension. In an uncertainty landscape, peaks express high probabilities, valleys lower probabilities. An uncertainty landscape may be draped with a many-coloured class layer.

3.2   Dynamic visualisation

Dynamic visualisation, in which time is added as an extra dimension, is rapidly gaining ground. MacEachren (1994) distinguishes between the use of temporal graphic variables to animate static

maps and their use to depict dynamic processes in time. Because the uncertainties underlying a remote-sensing classification basically lack meaning in terms of evolution in time, we focus on the use of temporal variables to make static maps more readily accessible.

A *clickable map* basically is a traditional many-coloured map of a classification of remotely-sensed data with activatable extra information. The information concerning the uncertainty underlying the classification is not directly visible in the map but may be activated by pressing the mouse button at a pixel. Upon pressing the button the pixel's vector of probabilities is shown as a separate graphic. Alternatively, pressing the mouse button can reveal the second most likely class for the pixel, directly in the image, by changing the colour assigned to the pixel. By transecting the image with the mouse, an uncertainty profile can be created, displaying for every pixel visited, the associated probability vector.

Instead of graphically visualising the uncertainty underlying a remote-sensing classification, Fisher (1994a) proposes sound as a sensory variable: a strident noise may be used to alarm a user whose cursor is entering a less reliable part of the map, relating uncertainty proportionally to tone. Sound can facilitate the visualisation of uncertainty while preserving the static graphic variables for the thematic information of the classification at hand. The perceptual consequences of exploiting sound for conveying information have, however, hardly been explored (Krygier, 1994).

A traditional many-coloured map of a remote-sensing classification and its associated maximum probability grey-scale map may be toggled to result in a combined sensation of classes and uncertainty in a *dynamic map pair*. The time lapse between the alternating display of maps can generally be selected interactively. Caution, however, is called for to avoid annoying effects of only slowly interchanging images. Such an alternating display of maps typically results in a phenomenon known as *chromatic adaptation*: the display of the many-coloured map causes an after-image that is projected onto the maximum probability map.

The sequential display of different many-coloured thematic maps and associated uncertainty maps derived from the same set of remotely-sensed data may make up an *animation*. For example, the probabilities from the probability vectors yielded by classification can be visualised successively according to their ranking, thereby enabling the representation of all probabilities underlying the classification. The animation can follow different schemes, such as sequential, progressive, cyclic, and back-and-forth. Fisher (1994b) proposes an animation where the display of probabilities is based on the principle of duration: the higher a class' probability for a pixel, the longer the period of time the class colour is displayed. Animations can also aid in exploring a three-dimensional uncertainty landscape: skimming over its peaks and valleys can have a dramatic impact on the understanding of the uncertainty involved.

4   Weighted uncertainty

The vectors of probabilities computed by a classification procedure for a set of remotely-sensed data provide useful information about the quality of the resulting classification in terms of the uncertainty involved. In the foregoing, we briefly reviewed some simple quality measures that are in use for extracting this type of information. Unfortunately, as these measures do not address the entire probability vectors, they fail to capture the distribution and extent of the uncertainties underlying a classification. Overall knowledge of the uncertainties involved, however, is imperative for acquiring in-depth insight in the quality of the classification. To exploit fully the information content of a probability vector for this purpose, additional measures of uncertainty are called for.

For a detailed exploration of the uncertainties underlying a remote-sensing classification, we propose using a measure from among a class of measures that build on the notion of *weighted uncertainty*. A well-known example of such a measure is the *entropy* measure originating from information theory (Shannon, 1948; Kullback, 1954). The measure pertains to a statistical variable and the uncertainties in its various possible values, expressing in a single number the distribution and extent of these uncertainties. We would like to note that the use of the entropy measure for capturing the uncertainty underlying a remote-sensing classification has already been briefly hinted at by Goodchild *et al.* (1994).
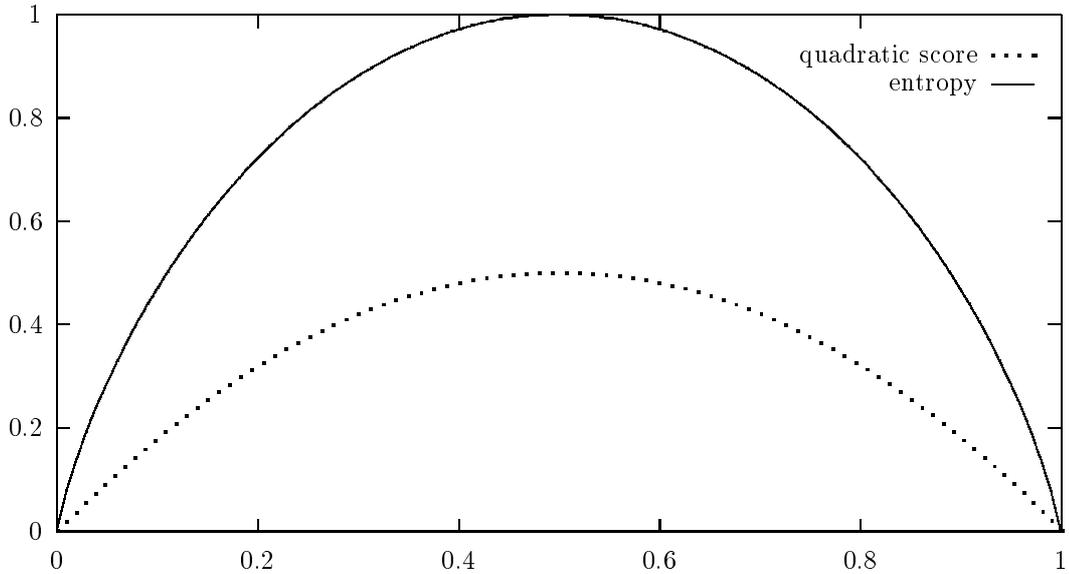
Figure 4: The entropy and quadratic score of a pixel.

In the entropy measure, the uncertainty in a single value of a statistical variable is defined as the *information content* of a piece of information that would reveal this value with perfect accuracy. For a pixel in a remote-sensing classification, viewed as a statistical variable C, the uncertainty in class C_i is defined as:

$$- \log_2 \Pr(C = C_i \mid x)$$

for $i = 1, \ldots, n$, where x denotes the available data; the uncertainty is measured in units of bits of information. Generally, the true class of the pixel is not known and, as a consequence, the amount of information required to reveal the pixel's class is unknown. The entropy of the pixel is therefore defined as the expected information content of a piece of information that would reveal its true class. To this end, the entropy measure combines the uncertainties in the various classes of the pixel by weighting them by their probabilities:

$$- \sum_{i=1,\ldots,n} \Pr(C = C_i \mid x) \cdot \log_2 \Pr(C = C_i \mid x)$$

The pixel's entropy is minimal if the uncertainty as to its true class has been resolved. Thus, if $\Pr(C = C_i \mid x) = 1$ for some class C_i, $1 \leq i \leq n$, that is, if class C_i has been established with perfect accuracy, then the entropy equals zero and there is no further information required to reveal the pixel's true class. Conversely, the entropy is maximal if none of the classes is preferred, that is, if there is utter ignorance as to the pixel's true class. So, if the probabilities $\Pr(C = C_i \mid x)$, $i = 1, \ldots, n$, are uniformly distributed, that is, if for all classes C_i we have that $\Pr(C = C_i \mid x) = \frac{1}{n}$, then the entropy is at maximum. Figure 4 depicts the entropy of a pixel that can have one of two classes; the x-axis represents the probability of one of these two classes and the y-axis represents the pixel's entropy.

The main advantage of the entropy measure over the simple quality measures mentioned before, is its ability to summarise all the information contained in a vector of probabilities in a single number. The entropy measure is, however, not the only measure that exhibits this property. In essence, any measure that builds on the notion of weighted uncertainty will do the same thing.

As another measure of weighted uncertainty, we briefly discuss the *quadratic score* (Glasziou and Hilden, 1989). The quadratic score builds on the notion of confirmation. The uncertainty in

a single class for a pixel is the amount of *probability* required to establish this class with complete accuracy. The uncertainty in class C_i is defined as:

$$1 - \Pr(C = C_i \mid x)$$

for $i = 1, \ldots, n$, where x once more denotes the available data. The quadratic score of the pixel is then:

$$\sum_{i=1,\ldots,n} \Pr(C = C_i \mid x) \cdot (1 - \Pr(C = C_i \mid x))$$

This measure exhibits the same behaviour in its minimum and maximum values as does the entropy measure. The two measures differ, however, in their slopes as is shown in Figure 4. The slope of the entropy measure is steeper than the slope of the quadratic score. As a result, the entropy measure for example more strongly weights small deviations from probabilities equal to zero or one than the quadratic score.

5 Visualisation of weighted uncertainty

The class of measures of uncertainty that have been introduced in the previous section allow for capturing overall knowledge of the uncertainties underlying a remote-sensing classification. As these measures summarise the distribution and extent of the uncertainties involved in a single number per pixel, they provide for conveying uncertainty information in a straightforward way. Entropy can, for example, simply be visualised in a grey-scale map. The practicality of an entropy map for exploring uncertainty in remote-sensing classifications is illustrated using an example taken from the CAMOTIUS-package. The CAMOTIUS-package is a PC-application, written in Visual Basic, that provides a user-friendly approach to data exploration. It offers, for example, a number of probabilistic classification procedures, each yielding vectors of posterior probabilities as described before.

Our example addresses a data set selected from a Landsat Thematic Mapper image of the Westland area of the Netherlands. In classifying these data, six land cover classes were distinguished: water, grass, arable land, forest, bare soil, and built-up area. The class of built-up area has been further subdivided in three additional land use classes as urban, industrial, and greenhouse, the latter being characteristic for the Westland study area. Figure 5 shows some examples of grey-scale maps of the uncertainties underlying our classification, as offered by the CAMOTIUS-package. These maps, and especially the entropy map, reveal interesting information. The distinction between the agricultural areas and built-up areas is clearly reflected by the entropy map while the other maps fail to reveal this information unambiguously. The fuzzy character of urban areas results in a high entropy (light grey to white in the map) as opposed to the low entropy (dark grey to black) for pastures that are easily distinguished by the classifier. This is especially well illustrated for the area around Maasland, a village that is delineated by a greyish, north-south orientated spot in the lower right corner of the entropy map. Also, it appears that employing the entropy measure provides an edge detector. Transitions between homogeneous areas, such as pastures, are characterised by a high entropy.

In addition to grey-scale maps, the CAMOTIUS-package provides bivariate maps for the simultaneous visualisation of qualitative classes and entropy values. Bivariate maps may further be extended to clickable maps in which pressing the mouse button at a pixel shows the entire vector of probabilities. Figure 6 illustrates visualisation of the probability vector as an array and Figure 7 shows it as an uncertainty profile.
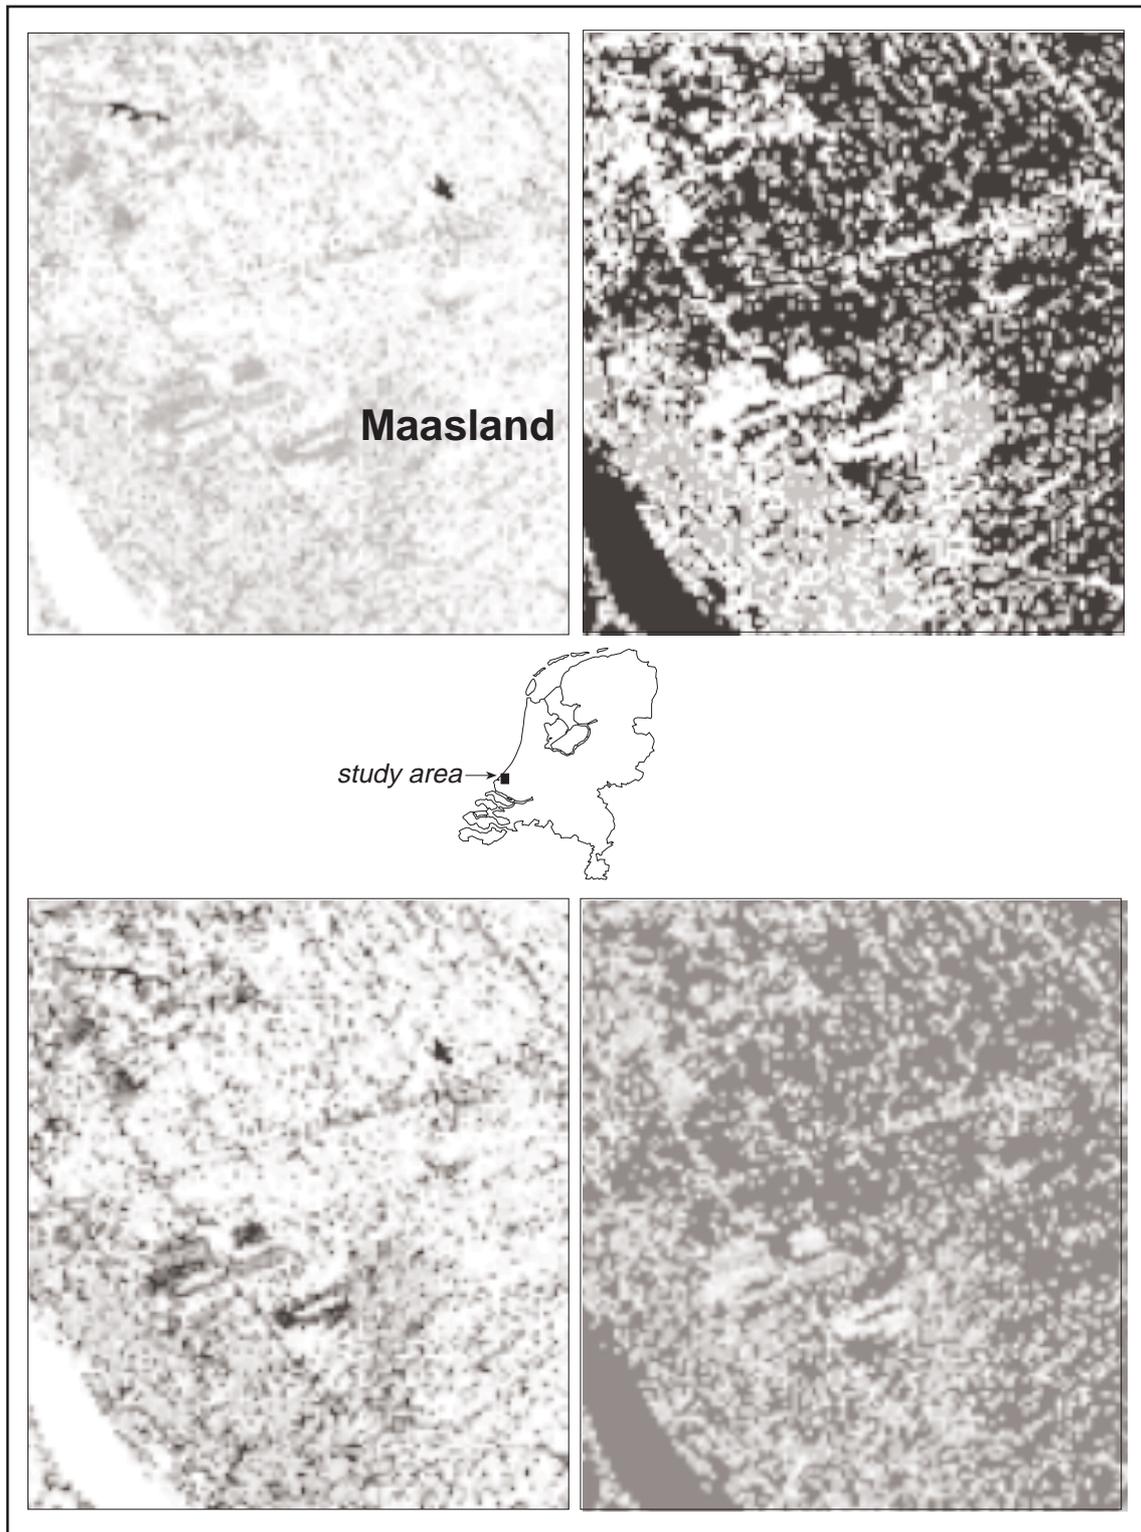
Figure 5: Examples of grey-scale maps as provided by the CAMOTIUS package: clockwise from top left are the maximum probability, entropy, difference value, and second ranking probability maps.
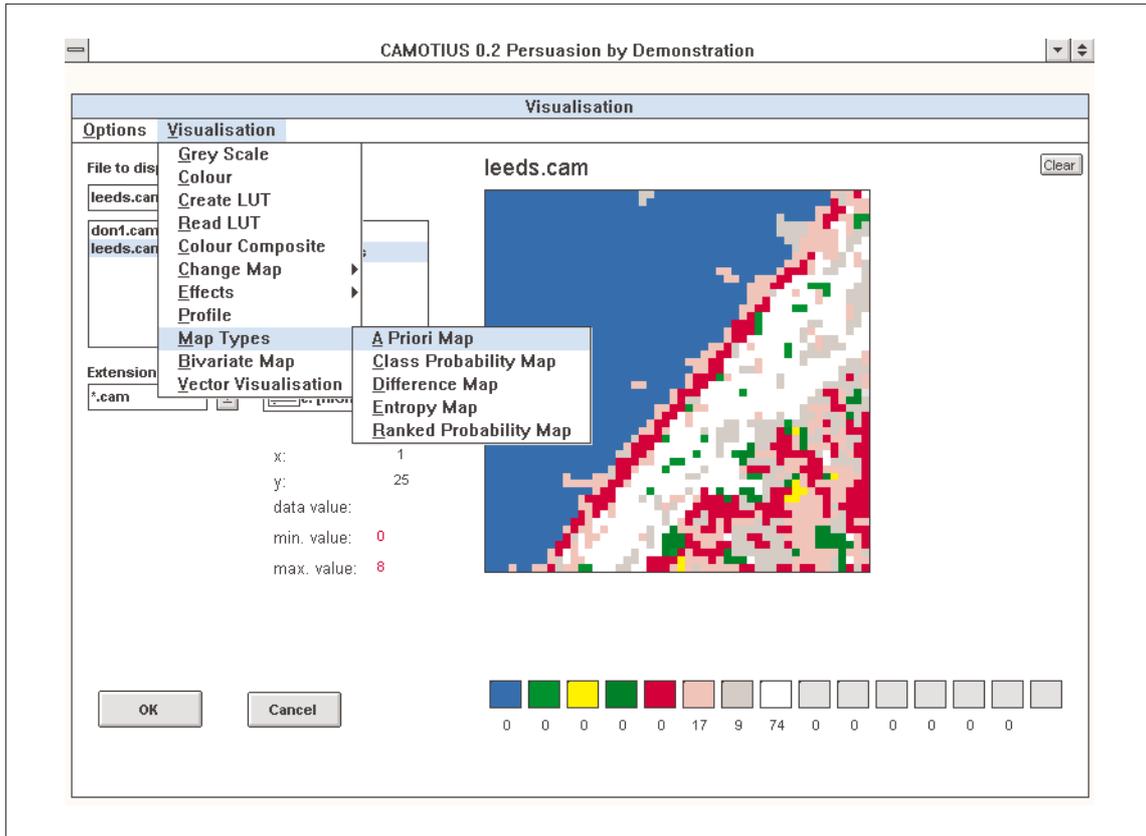
Figure 6: Representation of a probability vector as an array.

6 Conclusions

The advent of computer technology facilitates an exploratory analysis of various classifications of a set of remotely-sensed data to arrive at attractive and cartographically sound products. To support exploratory analysis, a user needs to be provided with informative measures of uncertainty and with effective ways of visualising uncertainty. Various measures are already in use for extracting uncertainty information from the vectors of probabilities yielded by most classification procedures. These measures, however, fail to capture the distribution and extent of the uncertainties underlying a classification and, as a consequence, may hide relevant information. We have proposed supplementing them by measures that build on the notion of weighted uncertainty, addressing the entire probability vectors. Since these measures summarise all information contained in a vector of probabilities in a single number per pixel, they allow for straightforward visualisation. We have addressed the visualisation of entropy, as an example measure of weighted uncertainty, in the context of the CAMOTIUS-package.

Providing a user with informative measures of uncertainty and with ways of cartographic visualisation of uncertainty is of key importance to ensure the use of data that are fit for the application aimed at by the user. We hope that the measures and visualisation techniques introduced in this paper may contribute to an effective assessment of fitness for use of a set of remotely-sensed data.
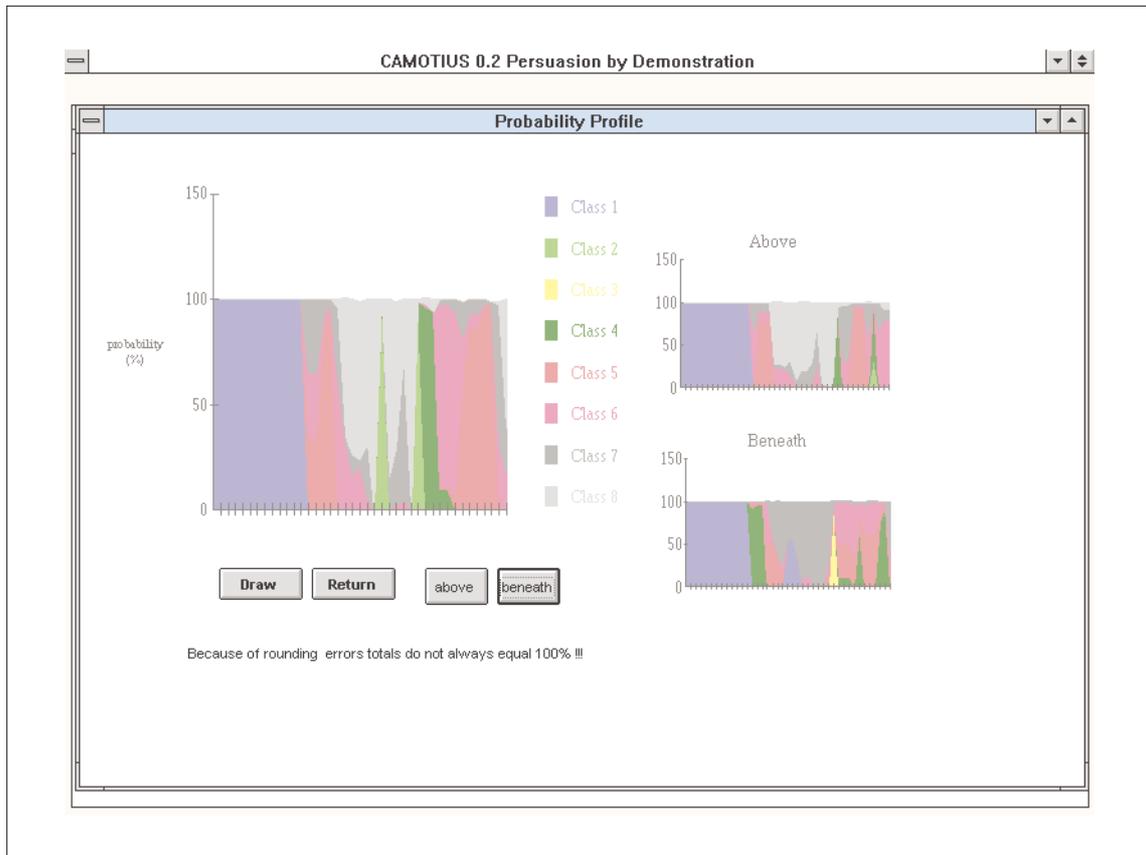
Figure 7: Representation of a probability vector as a profile.

References

M.K. Beard and B.P. Buttenfield (1991). *Visualising the Quality of Spatial Information: Scientific Report of the Specialist Meeting*, NCGIA Technical Report 91-26.

J. Bertin (1983). *Semiology of Graphics* (translated by W.J. Berg), University of Wisconsin Press, Madison.

A. Brown and C.P.J.M. van Elzakker (1993). The use of colour in the cartographic representation of information quality generated by a GIS. *Proceedings of the 16th International Cartographic Conference*, vol. 2, pp. 707 − 720.

J. Denègre, Editor (1994). *Thematic Mapping from Satellite Imagery. A Guidebook*, Elsevier Science Ltd., Oxford.

D. DiBiase, A.M. MacEachren, J.B. Krygier, and C. Reeves (1992). Animation and the role of map design in scientific visualisation. *Cartography and GIS*, vol. 19, pp. 201 − 214.

P.F. Fisher (1994a). Hearing the reliability in classified remotely sensed images. *Cartography and GIS*, vol.21, pp. 31 − 36.

P.F. Fisher (1994b). Visualisation of the reliability in classified remotely sensed images. *Photogrammetric Engineering & Remote Sensing*, vol. 60, pp. 905 − 910.

G. Foody and P. Curran (1994). *Environmental Remote Sensing from Regional to Local Scales*, Wiley & Sons, Chichester.

P. Glasziou and J. Hilden (1989). Test selection measures. *Medical Decision Making*, vol. 9, pp. $133 - 141$.

M.F. Goodchild, Sun Guoqing, and Yang Shiren (1992). Development and test of an error model for categorical data. *International Journal of Geographical Information Systems*, vol. 6, pp. $87 - 104$.

M.F. Goodchild, L. Chih-Chang, and Y. Leung (1994). Visualizing fuzzy maps. In: H.M. Hearnshaw and D.J. Unwin, Editors. *Visualisation in Geographical Information Systems*, Wiley & Sons, Chichester, pp. $158 - 167$.

S.C. Guptill and J.L. Morrison, Editors (1995). *Elements of Spatial Data Quality*, Elsevier Science Ltd., Oxford.

L.L.F. Janssen and F.J.M. van der Wel (1994). Accuracy assessment of satellite derived land-cover data: a review. *Photogrammetric Engineering & Remote Sensing*, vol. 60, pp. $419 - 426$.

J.B. Krygier (1994). Sound and geographic visualisation. In: A.M. MacEachren and D.R.F. Taylor, Editors. *Visualisation in Modern Cartography*, Elsevier Science Ltd., Oxford, pp. $149 - 166$.

S. Kullback (1954). *Information Theory and Statistics*, Wiley & Sons, New York.

R.S. Lunetta, R.G. Congalton, L.K. Fenstermaker, J.R. Jensen, K.C. McGwire, and L.R. Tinney (1991). Remote sensing and geographic information system data integration: error sources and research issues. *Photogrammetric Engineering & Remote Sensing*, vol. 57, pp. $677 - 687$.

A.M. MacEachren (1992). Visualizing uncertain information. *Cartographic Perspectives*, no.13, pp. $10 - 19$.

A.M. MacEachren (1994). Time as a cartographic variable. In: H.M. Hearnshaw and D.J. Unwin, Editors. *Visualisation in Geographical Information Systems*, Wiley & Sons, Chichester, pp. $115 - 130$.

P.C. Muehrcke and J.O. Muehrcke (1992). *Map Use. Reading, Analysis, Interpretation*, 3rd edition, J.P. Publications, Madison.

C.E. Shannon (1948). The mathematical theory of communication. *Bell Systems Technical Journal*, vol. 27, pp. $379 - 423$.

D.J. Unwin (1995). Geographical Information Systems and the problem of "error and uncertainty". *Progress in Human Geography*, vol. 19, pp. $549 - 558$.

F.J.M. van der Wel, R.M. Hootsmans, and F.J. Ormeling (1994). Visualisation of data quality. In: A.M. MacEachren and D.R.F. Taylor, Editors. *Visualisation in Modern Cartography*, Elsevier Science Ltd., Oxford, pp. $313 - 331$.