The Continuum of Classification Fuzziness in Thematic Mapping

Giles M. Foody

Abstract

Thematic mapping from remotely sensed data is generally achieved through the application of a supervised image classification. Although this is one of the most common applications of remote sensing, the maps derived are often of insufficient accuracy for operational use. Consequently, considerable research effort has been directed at increasing the accuracy of thematic mapping, particularly through the development of classification techniques that make fuller use of the information content of remotely sensed data. One major set of problems limiting the accuracy of thematic maps derived from remotely sensed data relates to conceptual issues associated with the use of classification techniques as the tool for mapping. It is shown that the conventional "hard" classifications may be less appropriate than fuzzy classifications and that a continuum of classification fuzziness can be defined. The potential for classification at any point along this continuum, from completely crisp to fully fuzzy, is discussed and may provide a framework for realizing more fully the potential of remote sensing as a source of thematic map data. The implications of the continuum on spatial data standards and reporting are also briefly discussed.

Introduction

Remotely sensed data may be used to map, monitor, and estimate the properties of environmental features. Of these three main areas of application, mapping, and thematic mapping in particular, is one of the most widespread and may be a necessary precursor to the others. Thematic mapping from remotely sensed data is found in numerous fields-of-study from forestry through climatology to geology. It also has a relatively long history, with maps derived initially from interpretation of aerial photography but with emphasis increasingly placed on computer-based analyses of digital imagery. Traditionally, classification techniques have been used as the tool for thematic mapping, whether based on visual or digital analyses of the remotely sensed data. Thus, for instance, classification is the basis of mapping through aerial photograph interpretation as well as from computer based analyses of digital sensor data (Lo, 1986; Campbell, 1996). Numerous classification approaches have been used with varying degrees of success. Despite the considerable developments made recently, the accuracy with which thematic maps may be derived from remotely sensed data is, however, often still judged to be too low for operational use (Townshend, 1992), and the problem appears to be largely independent of classifier type (Wilkinson, 1996). Thus, despite thematic mapping being one of the most common applications of remote sensing, with a history extending back over several decades, we have not yet reached the stage at which accurate land-cover maps can be derived from remotely sensed data on an operational basis (Townshend, 1992).

A range of reasons may be cited for the failure to realize the enormous potential of remote sensing as a source of land-cover and other thematic data. These include issues such as the nature of the classes, the spectral and radiometric resolutions of the remotely sensed data, and the methods used in mapping. Here attention is focused only on the latter issue. In particular, this paper aims to discuss briefly the conventional "hard" classification-based approaches to mapping and highlight their inappropriateness relative to "fuzzy" methods, especially for imagery containing a high proportion of mixed pixels. A major aim, however, is to show that widely used "fuzzy" classifiers are only a partial solution to the mixed-pixel problem as they do not accommodate fuzziness throughout the classification process. A continuum of classification fuzziness will be presented and methods shown that may be used at points along the continuum to assist mapping investigations and thereby provide a greater realization of the potential of remote sensing as a source of thematic data. Some of the main implications for spatial data standards and reporting are also discussed.

Classification for Mapping

Of the many approaches available, per-pixel supervised image classifications are commonly used in thematic mapping from remotely sensed data. Such classifications are the sole focus of this paper, although much of the discussion will be of wider applicability. A variety of classification algorithms have been used in remote sensing, notably conventional statistical procedures such as the maximum-likelihood algorithm and, recently, approaches based on artificial intelligence such as neural networks (Schowengerdt, 1997). Whatever the specific algorithm adopted, the classification process may be broken down into three basic stages. First, the training stage in which class descriptors are generated. These are used by the classification algorithm in the second, class allocation, stage to allocate each pixel of unknown class membership to the class with which it has greatest similarity. Third, the testing stage, in which the accuracy of the classification is assessed. The approaches used in each stage are generally intuitively appealing, widely available, and, when their application is appropriate, can be used to classify data accurately. There are, however, problems in their use. At the outset it should be stressed that classification is a subjective process and the quality of the final classification is, therefore, in part a function of the analyst's skill and judgement. The analyst has, for instance, control or influence over many factors

Department of Geography, University of Southampton, Highfield, Southampton SO17 1BJ, United Kingdom (g.m.foody@soton.ac.uk).

Photogrammetric Engineering & Remote Sensing,
Vol. 65, No. 4, April 1999, pp. 443–451.

0099-1112/99/6504–443$3.00/0

which affect classification accuracy. This includes fundamental issues such as the selection of the training sample size and sampling design, the definition of the classes, and the determination of numerous parameters linked to the selected classification algorithm such as those defining the learning rate in a neural network and prior probabilities in probabilistic classifiers (Foody, 1999). Furthermore, many of the classification techniques used make assumptions about the data which are often not realized. The widely used maximum-likelihood classification, for example, assumes that the data are normally distributed. This algorithm may, of course, still be used when its assumptions are not satisfied. Although small deviations from the assumed conditions may be unimportant, others may be a source of significant misclassification. Perhaps more importantly, and the focus of this paper, the appropriateness of classification as the tool for mapping is itself debatable.

The classification techniques commonly used are "hard," with each image pixel allocated to a single class. A pixel therefore displays full and complete membership to a single class. Such approaches are only appropriate for the mapping of classes that are discrete and mutually exclusive and assume the data can be represented in crisp sets. On many occasions this will not be the case. Many land-cover classes, for instance, are continuous and so intergrade. This problem is partly a consequence of class definition. For example, how much tree cover is required for a patch of land to be classed as forest? Furthermore, classification is only appropriate if the basic spatial unit used, typically the pixel, is pure. This is rarely the case, with many pixels of mixed land-cover composition contained within a remotely sensed image (Crapper, 1984; Campbell, 1996). These mixed pixels may occur whatever the nature of the classes. For instance, with continuous classes, mixed pixels will occur frequently in the inter-class transition zones where the classes co-exist spatially whereas, for discrete classes, the area represented by a pixel will often enclose or straddle class boundaries. The exact proportion of mixed pixels in an image will vary with a range of factors, notably the land-cover mosaic on the ground and the sensor's spatial resolution, but is often very large (Crapper, 1984; Campbell, 1996). The proportion of mixed pixels will also generally increase with a coarsening of the sensor's spatial resolution. Thus, for coarse spatial resolution data sets used in mapping land cover at regional to global scales, and where remote sensing has perhaps its greatest potential role in thematic mapping, mixed pixels may vastly dominate imagery (Foody *et al.*, 1997). The mixed-pixel problem is not, however, confined to coarse spatial resolution data. At fine spatial resolutions mixing still occurs. At this scale much of the mixing arises from intra-class variability. With an agricultural crop class, for instance, the concern may be mixing due to the variable proportions of its components such as sunlit leaves, shadow, and soil within the sensor's field-of-view. Spectral mixing problems will, therefore, be evident in fine spatial resolution data sets, particularly for heterogeneous classes such as urban areas, and classification quality will remain a function of spatial resolution (Townshend, 1981; Cushnie 1987). Thus, the data from a range of proposed fine spatial resolution sensors (Barnsley and Hobson, 1996; Aplin *et al.*, 1997) may be as inappropriate for "hard" classification analyses as coarse spatial resolution data sets. If the full potential of remote sensing as a source of land-cover data is to be realized, alternative approaches to "hard" classification for thematic mapping may be required. These must enable a pixel to possess multiple and partial class membership, a characteristic feature of mixed pixels. As such, fuzzy rather than crisp sets provide a more appropriate framework for the analysis of remotely sensed data sets.

The inability of conventional "hard" classification outputs to appropriately represent themes such as land cover from data containing a significant proportion of mixed pixels has been a major driving force in the development of alternative approaches (Wang, 1990a). These have focused on fuzzy classification techniques in which the full class membership of each pixel is partitioned between all classes. In such a classification, a pixel can display any possible membership scenario, from full membership to one class through to having its membership divided, in any permutation, between all classes. Fuzzy classifications may be derived in a number of ways. There are, for example, a range of fuzzy classifiers (e.g. Cannon *et al.*, 1986). Alternatively, a fuzzy classification may be achieved by "softening" the output of a "hard" classification. This is typically possible because the conventional "hard" classification approaches generate significant class membership information for the determination of a class label which may be output and mapped. A probability vector containing the probability of membership a pixel has to each defined class could, for example, be output from a probabilistic technique such as the maximum-likelihood classification. There are many related ways of softening other classifications (Foody, 1996; Schowengerdt, 1997; Warner and Shank, 1997). In general, however, measures of the strength of class membership calculated for each pixel in the determination of its class label, which are normally discarded, may be output. This type of output makes fuller use of the information on class membership generated in the classification and may be considered to be fuzzy, as an imprecise allocation may be made and a pixel can display membership to all classes. A number of studies have shown that measures of the strength of class membership derived from a variety of classification algorithms may be used to indicate the land-cover composition of mixed pixels (Fisher and Pathirana, 1990; Bernard *et al.*, 1996; Foody, 1996; Foody and Arora, 1996; Schowengerdt, 1996; Veregin, 1996; Foschi and Smith, 1997).

Unlike their "hard" counterparts, fuzzy classifications can therefore provide a fuzzy class allocation that forms the basis of an appropriate representation of themes that may be considered to be fuzzy. However, they share with the "hard" classifications the three stages of training, allocation, and testing in their production and evaluation. The adoption of a fuzzy classification algorithm or softening of a "hard" classification output in recognition of the need for a class assignment that allows for multiple and partial class membership does not in itself make any account for mixed pixels in the training and testing stages of the classification. In many of the fuzzy classifications reported in the literature, therefore, fuzziness is only accommodated in the second, class allocation, stage of the classification process. The algorithm used to derive the fuzzy classification still requires training and its output testing, and, for these stages of the analysis, approaches designed for conventional "hard" classifications, which assume pure pixels, are typically employed. This is inappropriate if the theme of interest may be considered fuzzy. The fundamental problem is that the pixel, as the basic spatial unit of the remotely sensed data and derived classification, constrains the scale at which the ground data may be integrated into the analysis in the training and testing stages. Just as a mixed pixel cannot be appropriately allocated by a "hard" classification which assumes pure pixels, it cannot be used appropriately in conventional training and testing procedures which also assume pure pixels. For example, the composite spectral response of a mixed pixel does not ideally represent any of its classes, and any "hard" class allocation of a mixed pixel must to some extent be erroneous. Consequently, the use of a fuzzy classification technique may be considered to be only a partial solution to the mixed pixel problem as it only provides a means of appropriately representing a theme such as land cover that may be considered fuzzy at the scale of the pixel. Fuzziness should be recognized as an issue in the training and testing stages of the classification as much as in the allocation stage.

There is a need to recognize that fuzziness may, depending on the specific circumstances of an investigation, require accommodation in some or all three stages of the classification. Research has, however, focused mainly on the derivation of a fuzzy class allocation. Relatively little attention has been directed at the accommodation of mixed pixels in the training and testing stages of a supervised classification. Nonetheless, a variety of ways of accommodating mixed pixels in these stages of a supervised classification may be identified. Their effect could, for instance, be ignored and the analysis proceed as if mixed pixels were not present, making an implicit assumption that the pixels are in fact pure. Despite the mixed-pixel problem being well known, many analysts do exactly this, paying little more than lip service to the issue (Fisher, 1997). It must be recognized that erroneously assuming pure pixels will degrade all three stages of the classification; mixed pixels will degrade the class responses derived in training, cannot be appropriately allocated by a "hard" classification, and may vary in the accuracy of their allocation. As an alternative approach to handling the mixed-pixel problem, the analyst could harden the data. With this approach, each pixel is given the code of the dominant class. Although recognizing the existence of mixed pixels, no real accommodation for their presence is made and the classification will be degraded in the same way as if the mixed pixels had been ignored. An extreme solution to the problem would be to exclude mixed pixels from the analysis, although their identification may be difficult and this could result in the loss of most of the data set. This approach does, however, at least ensure that the data set used is appropriate for analysis by conventional techniques, which assume pure pixels, if desired.

Frequently, the only way the fuzziness of the land cover on the ground is accommodated in training and testing a classification is by deliberately avoiding mixed pixels (Metzler and Cicone, 1983). Thus, although an image may be dominated by mixed pixels, only pure pixels are selected for training. This typically involves selecting training sites from only very large homogeneous regions of each class in order to avoid contamination of training sites by other classes (Campbell, 1996). Moreover, research on refining training sets has often focused on removing potentially mixed pixels from the training set (Arai, 1992). With these approaches, it may be difficult to acquire a training set of an appropriate size. Moreover, the training statistics defined may not be fully representative of the classes and so provide a poor base for the remainder of the analysis.

In testing the classification, pure pixels only are again generally selected or assumed. Conventional measures of accuracy assessment, which were designed for application to nominal level outputs derived with a "hard" classification algorithm, are generally used. As the majority of pixels may be mixed, however, failure to include them in the accuracy assessment may result in an inappropriate and unrealistic estimation of classification accuracy. The accuracy statement may, for instance, only indicate the quality of the classification for "pure" pixels and provide a poorer expression of the quality for the remainder, possibly the vast majority, of the data set. There has recently been some investigation into the accommodation of fuzziness into the testing stage. Often this has been driven by the need to recognize that the ground data used in evaluating the accuracy of a classification are not error-free. Indeed many studies note that disagreements between the classification and ground data often arise through error or uncertainty in the ground data (Harris and Ventura, 1995; Abrams *et al.*, 1996; Bowers and Rowan, 1996). Alternatively, attempts have been made to accommodate fuzziness in the class allocation derived from a fuzzy classification (Maselli *et al.*, 1994; van der Wel *et al.*, 1998). Thus, techniques to accommodate fuzziness in the ground data or the remotely sensed data and derived classification have been suggested but little attention has focused on accommodating for fuzziness in both. Additionally the various methods suggested differ in the extent to which fuzziness is accommodated. Some approaches, for example, consider only the primary and secondary class labels while others utilize the full class-membership information contained within the ground data and classification.

Because a large proportion of image pixels may be mixed, it is important that they be accommodated throughout the classification. The degree to which fuzziness is accommodated will be a function not only of the nature of the data sets but also of practical constraints faced by the analyst. By recognizing that there are various degrees to which fuzziness may be incorporated in each stage, a continuum of classification fuzziness may be defined. This may help select appropriate techniques for the analysis of a particular data set. The aim of the remainder of this article is to describe briefly this continuum and discuss some of its implications to studies based on thematic mapping from remotely sensed data, specifically the assessment and reporting of data quality.

Fuzziness in Supervised Classification

Recognition of the mixed-pixel problem was a major driving force in the development of approaches for fuzzy classification (Wang, 1990a). As fuzziness may be a characteristic feature of *both* the remotely sensed data and the ground data, the use of a fuzzy classification algorithm alone may be insufficient for the resolution of the mixed-pixel problem (Foody, 1995; Foody and Arora, 1996). A more appropriate approach would be, within the practical and logistical constraints pertaining to the mapping investigation, to recognize the existence of mixed pixels and use techniques that accommodate for their presence throughout the mapping process (Foody, 1995; Foody, 1997). The background to this and some suggestions as to how it may be achieved follow. For clarity, the stages of the classification process will be reviewed briefly and possible means of accommodating fuzziness in each highlighted. For simplicity, reference will be made only to per-pixel classifications based solely on spectral information; but the discussion is applicable to other classification approaches (e.g., per-parcel classification using spectral and spatial discriminatory variables).

Training Stage

A fundamental issue in any supervised classification is the quality of the data used in training. The aim of the training stage is to derive an accurate and representative description of each class. Provided the remotely sensed data have been preprocessed appropriately, it may be possible to use library spectra. However, given the vast array of problems in achieving the desired accuracy of calibration and the range of episodic and external variables that could influence the spectral response of a class in the image to be classified, it is more common and appropriate to derive the training data from the image itself. Typically, this is achieved by extracting data from sites of known class membership in the image and statistically characterizing them in terms of the discriminating variables to be used in the classification. There is a large literature that helps define an appropriate training set for a classification. This specifies a range of considerations such as the sample size and design, the effect of spatial autocorrelation, and intra-image variations in the sensor's viewing geometry (Mather, 1987; Campbell, 1996). It is often impractical to collect training data that satisfy all the requirements. Indeed, training data are typically derived from a relatively small number of sites that are particularly pure or exemplars of the classes. This may not always be appropriate or even possible, particularly if the image is dominated by mixed pixels.

Fuzziness may, however, be accommodated into this stage of the classification. Wang (1990b) shows how data on the fuzzy membership properties of training samples may be used to derive refined class descriptors for use in a the class allocation stage. With this and other approaches, it may be possible to refine the training statistics derived from a training set in which some, or even all, of the pixels are mixed to obtain those statistics that would have been derived if the pixels had in fact been pure. These rectified training statistics could then be used in a conventional classification. Foody and Arora (1996) show how this may be achieved by using a linear mixture model in reverse to predict pure class responses from those derived from a sample of mixed pixels. They also illustrate an alternative approach, the use of a classification technique which enables the mixed-class membership of the training samples to be handled directly in the classification. The technique used was based on a feedforward neural network in which the class composition of each training pixel was defined as the target output during training.

Even if fuzziness is recognized, the analyst is often constrained to use only pure pixels in an attempt to ensure the correct use of a conventional classification algorithm. For this, the analyst typically aims to select training sites from large homogeneous areas of the classes and remove atypical data, even if the classes may be recognized to be continuous and, therefore, fuzzy (e.g., Joria and Jorgenson, 1996).

Allocation Stage

The derived training statistics are used to guide the allocation of pixels of unknown class membership to a class in this, the second, stage of the classification. To achieve this, the spectral response (and any other discriminating data) of a pixel is compared against that of each class in the training statistics, and the pixel is then allocated to the class with which it has the greatest similarity in accordance with some predetermined decision rule. The pixel may be allocated, for instance, to the class to which it is spectrally closest. There are many classifiers that may be used, and typically the output is wasteful of the class membership information derived as only the label of the class with which the pixel has the greatest strength of membership is provided. Recognition of the presence of mixed pixels leads to a desire to allow a pixel to display membership in more than one class. The means of accommodating mixed pixels is, therefore, to allow a pixel to have membership in more than one class. This is evident in studies that seek to apply secondary labels to class predictions and in fuzzier or softer classifications in which membership to all defined classes is possible. There are many approaches by which this may be achieved, notably by softening the output of conventional "hard" classifications. These include modifying conventional approaches such as the maximum-likelihood classification and neural networks as well as through the adoption of a fuzzy classification algorithm (Foody, 1996). The classification output may vary in its fuzziness from the single class label characteristic of a "hard" classification, through one depicting a primary and secondary label, to the provision of class memberships for each class. Which is used will depend on the application or desired means of representing the classification. Because many mixed pixels are usually composed of only a few classes and as it may be cumbersome to store all the class memberships for each pixel, it may be preferable to use only one or two secondary labels. This output may also be easier to represent cartographically than all the fuzzy membership information that could be derived for each class. Such an approach may also be preferable if the degree of correspondence between the fuzzy membership values and class composition is poor. This may, for example, occur if the training data are not fully representative of the classes (Canters, 1997).

Testing Stage

The utility of a thematic map is largely dependent on its quality. Typically, this is expressed in terms of the classification accuracy. The significance of this stage of the classification process is evident in the large literature devoted to it, particularly on the methods of accuracy assessment and comparison. Typically, the accuracy of an image classification is assessed using a global measure such as the percentage correct allocation or kappa coefficient of agreement, which provide a valuable summary of the overall quality of a classification. There has been considerable research undertaken on this stage of the classification to help the analyst derive an appropriate index of classification quality. Studies have focused on a range of issues beyond the index of accuracy itself to include factors such as the testing set composition and sampling design (Congalton, 1988; Campbell, 1996). There are many problems, however, with this stage, some of which are related to the fuzziness associated with mixed pixels.

Fuzziness may be accommodated in this stage and used to resolve some major problems in evaluating the quality of a classification (e.g., Gopal and Woodcock, 1994). For example, the ground data can rarely be assumed to be error-free. For this reason, the terms such as "ground truth" and even "accuracy" should be used and interpreted carefully. Testing a classification is rarely a matter of quantifying the accuracy of the class allocation with reference to a set of actual class labels. It is instead typically an evaluation of the level of agreement or correspondence between two sets of class allocations, both of which have their own error characteristics. Furthermore, in recognition of mixed pixels and the ability to extract more than just one class label for a pixel, approaches to accommodate secondary class labels have been investigated (Woodcock and Gopal, 1995; Woodcock *et al.*, 1996). These are required because the "hard" allocation for a mixed pixel must to some extent be erroneous. Because the magnitude of error may vary from negligible to complete error, approaches which can allow for the continuous nature of the class allocation to all classes may be most preferable. By using techniques that account for fuzziness in both the classification output and ground data, a more appropriate measure of classification quality may be derived (Foody, 1996).

A major problem with the testing stage of many classifications is that the accuracy measures used are typically global, with little if any information derived for individual locations. These measures, therefore, provide little information on the spatial structure of uncertainty within the derived classification (Canters, 1997). Often the user of a thematic map would benefit from the inclusion of per-pixel uncertainty information which can be derived. With a fuzzy classification output, for example, the magnitudes of the class memberships for a pixel provide the basis for the measurement and representation of features such as the confidence that may be placed with an individual class allocation (Corves and Place, 1994). A range of measures may be derived. For instance, the magnitudes of the measure of class membership a pixel displays to all classes could be used to derive indices to represent the uncertainty (Maselli *et al.*, 1994; Gong *et al.*, 1996; van der Wel *et al.*, 1997). Although this provides a useful indication of per-pixel classification quality, a global measure of accuracy is generally required as a summary statistic. The methods generally used, such as the percentage correct allocation or kappa coefficient of agreement, are not without problems. It is, for instance, typically assumed that the classes are unordered (van Deusen, 1996). In many instances, however, the classes may be ordered yet the conventional accuracy assessment techniques are used (e.g., Joria and Jorgenson, 1996). In such circumstances, the use of a conventional accuracy assessment measure, designed for application to nominal level data, will not make full use of the information content of the data and may not provide an appropriate index of classification quality. It may sometimes be preferable to use approaches that can accommodate the ordinal nature of the data set such as the weighted kappa coefficient (Cohen, 1968; Foody *et al.*, 1996). The analyst may at times be constrained to use a conventional measure of accuracy, for example, if evaluating the results against previously published work to maintain consistency, but the possibility that the measure does not fully use the data should be recognized and its implications considered.

The Continuum of Classification Fuzziness

Through a consideration of fuzziness applying to all three stages of the supervised classification, it is possible to define a continuum of classification fuzziness. With this continuum, it is apparent that there is not a simple distinction between "hard" and fuzzy classifications, but, rather, a range of classification approaches of variable fuzziness (Figure 1). To avoid confusion in terminology and because, for example, a "hard" classification may be applied (incorrectly) to fuzzy data, the continuum may be defined upon the nature of the data sets in terms of belonging to crisp or fuzzy sets (Klir and Folger, 1988).

Crisp sets are assumed throughout a "hard" classification. Thus, at the "hard" end of the continuum are what may, for the want of a better expression, be termed completely crisp classifications. These are based on the conventional approach to classification in which a pixel is associated with a single class at each stage of the classification. Most supervised classifications of remotely sensed data adopt or assume this approach, but its application may be inappropriate due to the presence of mixed pixels. At the other extremity of the continuum are fully fuzzy classifications. In these, fuzziness is accommodated in all three stages of the classification. This type of approach may provide a more realistic and accurate representation of the land cover of a site and use more fully the information content of the remotely sensed data. Between these extremes lie classifications of varying fuzziness. This includes those generally referred to in the literature as fuzzy classifications, in which only the class allocation stage actually accommodates fuzziness.

Although more difficult to derive than a completely crisp classification, a number of approaches exist for production of classifications of variable fuzziness. For instance, the conventional approach to maximum-likelihood classification can be modified to accommodate fuzziness in any or all three stages of the analysis (Foody and Arora, 1996). Neural networks are, however, particularly attractive, especially as they may accommodate the fuzziness of the training data directly into the classification (Foody, 1995; Foody *et al.*, 1997); the accommodation of fuzziness into the testing stage of the classification is effectively independent of the approach used to generate the fuzzy classification. Neural networks are also capable of providing a classification at any point along the continuum of classification fuzziness. They can, therefore, be used in the derivation of classifications varying from the completely crisp (Kanellopoulos *et al.*, 1992), through classifications of intermediate fuzziness (Foody, 1996), to fully fuzzy (Foody, 1995; Foody, 1997).

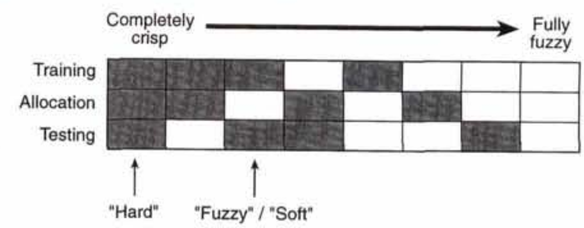

Figure 1. The continuum of classification fuzziness. Stages using crisp data are shown in black and those using fuzzy data in white; the distinction is less apparent in reality as the fuzziness of each stage is actually variable (see text for discussion). Because each stage of the classification may be based upon crisp or fuzzy data sets, a series of classifications of variable fuzziness may be defined ranging from those in which crisp data are used throughout (completely crisp classification) to those in which fuzzy data are used throughout (fully fuzzy classification). The traditional "hard" approaches to mapping are represented here at the completely crisp end of the continuum. Note also that the classifications termed as being "fuzzy" or "soft" in the literature generally lie between these extremes as typically only the class allocation stage accommodates fuzziness.

Because the fuzziness of each stage of the classification is itself variable, the classifications are arranged along a continuum of fuzziness and are not a small number of classifications that may be listed in order of fuzziness defined by the number of crisp and fuzzy stages. This can be illustrated briefly with reference to each stage. In training, the data may be crisp if the pixels are pure but for mixed pixels the training stage may vary in fuzziness. For instance, the training data used may vary from the code of the dominant class to individual class proportions. In the class allocation stage, a pixel could be allocated to a single class or the output might also contain the second most likely class through to the provision of membership grades to all classes. Lastly, the data used in testing the classification may, like the training data, vary from the code of the dominant class to individual class proportions and techniques which utilize the available information to variable extents adopted. In each stage there are, therefore, different ways of accommodating the fuzziness, which may use the information to differing extents. For example, in evaluating the accuracy of the output from a fuzzy classification, the analyst could degrade the data to enable the calculation of a conventional measure of classification accuracy (Foody, 1996), use, for instance, the two most likely classes of membership (Woodcock and Gopal, 1995; Zhang and Foody, 1998), or use all the data on class membership derived in comparison against the ground data set (Kent and Mardia, 1988; Foody, 1996). Therefore, two classifications that may accommodate for fuzziness in the same number of stages of the classification may still differ in their respective fuzziness.

Lastly, it may initially seem that fuzzy classifications, particularly the fully fuzzy ones, require more precise ground data for their application than would a conventional "hard" classification. This is not, however, the case. In a "hard" classification, the actual class of membership for each training and testing pixel must be determined to the same degree as in a fuzzy classification. The only difference is that in a "hard" classification each pixel should represent an area characterized by homogeneous cover of a single class whereas in the fuzzy classification each pixel may represent an area comprising any number of classes. In both, the analyst has to know the class composition of the entire area represented by each pixel; in any classification, the ground data for a pixel are supposed to represent the actual class cover observed on the ground. Thus, the fundamental requirement of ground data precision is *exactly* the same for classifications at any point along the continuum of classification fuzziness. So, while fuzzy classifications do require precise ground data for their application, it must be noted that the same level of ground data detail or precision is actually required for the correct application of a "hard" classification. The analyst may, of course, use less precise ground data in training or testing a classification but must accept that this could degrade the real and apparent accuracy of the classification. Using imprecise ground data in training a "hard" classification is, for instance, possible but, as this will cause the class training statistics to be contaminated by other classes, it may be a source of misclassification. It may also indicate that the analyst is prepared, or constrained, to ignore the presence of mixed pixels and so accept error in ground data when using a "hard" classification.

An Illustrative Case Study

The ability to classify a data set at various positions along the continuum of classification fuzziness can be illustrated with a small case study. This study focuses on the classification of land cover from airborne thematic mapper (ATM) data of suburban Swansea, UK. The test site is small and well defined, limiting complicating factors. For the study, detailed ground data were required on the class composition of all the pixels used in the analysis. There are two commonly used approaches to estimating the class composition of image pixels in the absence of detailed ground data. In both, the class composition of the pixels in the image to be classified is derived from a classification of imagery with a substantially finer spatial resolution in which the pixels may more reasonably be assumed to be pure. The first approach uses imagery acquired at two spatial resolutions. For instance, Landsat TM data may be used to derive the reference or ground data on class composition for an analysis of NOAA AVHRR data (e.g., Hlavka and Spanner, 1995; Foody *et al.*, 1997). A major problem with this approach is that even minor coregistration errors could result in considerable error in the land-cover composition estimates. The alternative approach is based on spatial degradation of a single image. The original undegraded image is used to derive the class composition of the spatially degraded image pixels. With this approach, the two data sets are coregistered, but it must be recognized that the method does not provide, particularly in terms of radiometry, an accurate simulation of a coarse spatial resolution image. This is not usually a problem, unless the analyst is trying to simulate data from a specific sensor when a more carefully defined degradation approach may be required. Although both approaches are not ideal, they do provide a means to estimate the class composition of pixels and thereby evaluate the quality of products such as fuzzy classifications. Here the spatial degradation approach was used in mapping three classes, trees, grass and asphalt. The data acquired in the 605- to 625-, 695- to 750-, and 1550- to 1750-nm spectral wavebands, which provide a high degree of inter-class separability, were used. These data were then degraded spatially with an 11 by 11 low-pass (mean) filter. From the resulting coarse resolution image, a sample of 50 pixels was extracted. This sample, while small, was restricted to pixels in which the class composition could be estimated with confidence (i.e., areas contaminated by other classes were avoided, etc.) from a visual classification of the

TABLE 1. ACCURACIES OF THE CLASSIFICATIONS DERIVED AT VARIOUS POINTS ALONG THE CONTINUUM OF CLASSIFICATION FUZZINESS. ACCURACY IS EXPRESSED AS THE RMS ERROR BETWEEN THE ESTIMATED AND ACTUAL CLASS COMPOSITION OF THE 35 PIXELS IN THE TESTING SET DEFINED; IT IS NOT THE ERROR OF NETWORK LEARNING. A LOW RMS ERROR INDICATES A CLOSE CORRESPONDENCE BETWEEN THE CLASSIFICATION OUTPUT AND GROUND DATA AND HENCE AN ACCURATE CLASSIFICATION. THE AVERAGE RMS VALUE INDICATES THE OVERALL ACCURACY OF THE CLASSIFICATION WHILE THE NUMBER OF CASES IN THE TESTING SET WITH LOW (<0.2) AND HIGH (>0.5) RMS ERRORS INDICATE THE MAGNITUDE OF INDIVIDUAL ERRORS WITHIN THE TESTING SET. IN THE TRAINING AND TESTING STAGES, THE FUZZY DATA WERE IN THE FORM OF CLASS PROPORTIONS WHEREAS THE CRISP DATA COMPRISED ONLY THE CODE OF THE DOMINANT CLASS. SIMILARLY, THE STRENGTH OF CLASS MEMBERSHIP, EXPRESSED BY THE MAGNITUDE OF THE OUTPUT UNIT ACTIVATION LEVELS, REPRESENTS THE FUZZY CLASS ALLOCATION WHEREAS THE CODE OF THE CLASS ASSOCIATED WITH THE MOST ACTIVATED OUTPUT UNIT FORMED THE CRISP CLASS ALLOCATION. INTERPRETATION AND COMPARISON OF THE ACCURACY STATEMENTS IS DIFFICULT. FOR EXAMPLE, BECAUSE THE TESTING SET CONTAINED MIXED PIXELS, THE ACCURACY ASSESSMENTS DERIVED USING CRISP DATA IN THE TESTING STAGE MUST TO SOME EXTENT BE FALSE. NOTE, THEREFORE, THAT THE APPARENT ACCURACY MAY DIFFER FROM THE ACTUAL. THE KEY ISSUE, HOWEVER, IS THAT THE ACCURACY STATEMENTS SHOWN WERE ALL EFFECTIVELY DERIVED FROM THE SAME DATA SETS BUT ARE DIFFERENT.

Data Used in Classification Stage			Classification Accuracy		
			Average	Number of Cases	
Training	Allocation	Testing	RMS	<0.2	>0.5
Fuzzy	Fuzzy	Fuzzy	0.150	29	1
Crisp	Fuzzy	Fuzzy	0.108	25	1
Fuzzy	Fuzzy	Crisp	0.204	27	5
Crisp	Fuzzy	Crisp	0.100	29	3
Fuzzy	Crisp	Fuzzy	0.150	22	4
Crisp	Crisp	Fuzzy	0.135	22	2
Fuzzy	Crisp	Crisp	0.116	30	5
Crisp	Crisp	Crisp	0.070	32	3

spatially undegraded data. This data set contained pixels of varying class membership properties, from pixels containing all three classes to pure pixels of each class. For the purposes of this investigation, 15 mixed pixels were selected to form a training set. Because the analyst would in practice generally attempt to use pure pixels as far as possible, the 15 training pixels were selected in such a fashion that there were five pixels that were strongly dominated by each class (>50 percent cover). The remaining 35 pixels formed the testing set upon which classification accuracy was assessed; further details on the data are given in Foody (1995; 1997).

A series of classifications was performed using a feedforward neural network. For these classifications, the data used in each stage of the classification were either in the form of fuzzy (e.g., class proportions) or crisp (e.g., dominant class label) sets, with a variety of classification scenarios, in terms of the three stages, being performed. In each classification, the network architecture (three input units, four hidden units, and three output units) together with the learning algorithm and its parameters were the same. To allow classification accuracies to be compared directly, a basic measure of accuracy was employed that may be used at any point along the continuum of classification fuzziness. This was the RMS error between the estimated and actual class composition of the pixels in the testing set. Thus, in each classification, exactly the same data and methods have been used to train and test the classifications; all that differed was whether the data used were fuzzy or had been hardened into crisp sets.

The results are presented in Table 1, which reveals a large variation in the measured accuracy. These accuracy statements are, however, difficult to evaluate and compare; markedly different results also may be obtained for a testing data set comprising different proportions of pure and mixed pixels and/or with a different degrees of mixing (Foody, 1996). The interpretation of the classification accuracy statements is also difficult as the quoted or apparent accuracy may be a poor indicator of the actual quality of the classification. This is particularly apparent when comparing results derived from classifications that were only dissimilar in the testing stage; here the accuracy statements derived with the fuzzy data may generally be considered a fairer indication of the real accuracy than the hardened data because the crisp class labels may hide error. The concern of this investigation, however, was not with issues such as the absolute accuracy but simply the variability between the derived classifications. The key point is that vastly different accuracies were derived from different points along the continuum. The ability to extract information, as indicated by the classification accuracy, varies along the continuum. Visual comparisons of hard and fuzzy classifications also reveal differences in the information on class membership and its distribution that may be derived (Figure 2).

The methods that may be used in each stage of the classification vary in their appropriateness along the continuum. This has numerous important implications for later users of the classification. Their use of the classification may, for instance, be enhanced by careful reporting of the whole classification process. This will help gain an appreciation of the real as opposed to apparent accuracy and standards of the data derived. Users will therefore, for instance, require a detailed documentation of the quality of the classification. This must go well beyond the provision of a conventional statement of classification accuracy or even a classification confusion matrix, as both may be highly inappropriate for the data. Instead, users will require information on the nature of the data sets and methods used to evaluate the classification and associated accuracy statement. There are also implications to the setting of and evaluation relative to data standards. As a basic example, how useful is a target standard for a classification of a kappa coefficient of 0.85 if the data comprise mixed pixels and the classes are ordered? This "hard" index of classification accuracy designed for nominal classifications cannot provide a universally appropriate basis for the evaluation of classification accuracy.

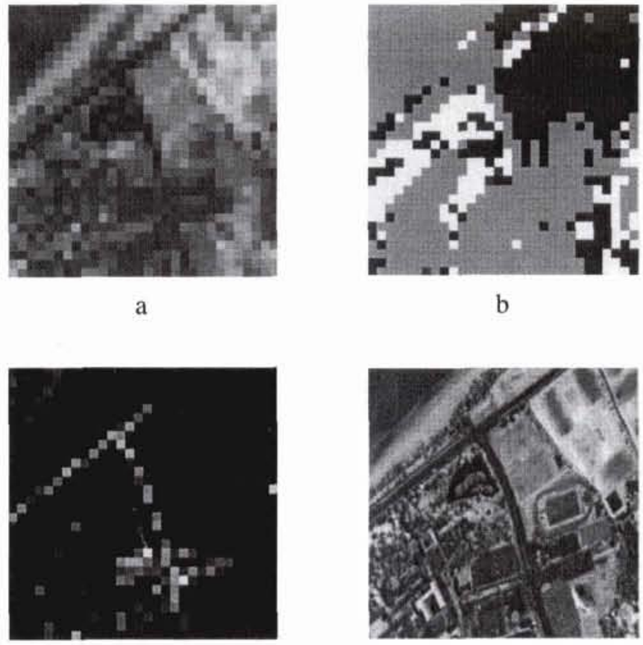

a b

c d

Figure 2. Outputs from hard and fuzzy classifications. (a) An extract of spatially degraded ATM data in the 695- to 750-nm waveband. (b) A conventional hard classification derived in which the light, medium, and dark tones represent the trees, asphalt, and grass classes, respectively. (c) A fuzzy classification output for the asphalt class in which the lightness of tone is positively related to the strength of membership to asphalt. (d) The original (undegraded) ATM data in the 695- to 750-nm waveband. Note that in the hard classification large areas have been commissioned into the asphalt class (e.g., the top left corner which actually comprises untrained classes (e.g., sandy beach)) which is shown to have low membership to asphalt in the fuzzy classification. Furthermore, the fuzzy classification indicates clearly the main road network and large car parks which are apparent in the fine spatial resolution image (d).

Discussion and Conclusions

Image classification, like many analyses of digital remote sensor data, is based generally on an assumption that the image pixels are pure. This assumption underlies the whole classification process even though pure pixels may comprise only a small proportion of the image. Despite this, conventional "hard" classification methods are often applied to remotely sensed data. The application of such methods can be a source of numerous problems. At a very basic level, using techniques that assume pure pixels on imagery dominated by mixed pixels may be very inappropriate and erroneous.

Mixed pixels cannot be accommodated sensibly in any of the three stages of a conventional supervised classification. Much of the effort directed at resolving this problem has focused on the second stage of the classification process with aim of deriving fuzzy or soft classifications which may represent the fuzzy theme more appropriately than would a "hard" classification. The realization of the effect and significance of mixed pixels was the driving force for the derivation and application of a range of fuzzy classifications. However, the term "fuzzy classification" has often been used in the literature to describe an analysis in which a fuzzy or soft classification output is derived, but no explicit reference is made usually to the training and testing stages. Indeed, the training data are often selected from large homogeneous regions in order to maintain training sample purity. The accuracy of a classification is also generally assessed with conventional accuracy assessment procedures which require pure pixels. In such an analysis, therefore, only one of the three stages of the classification is fuzzy; the training and testing stages may be considered to be "hard." The use of such a fuzzy classification does not, therefore, fully resolve the mixed-pixel problem, because their effects on the training and testing stages have not been considered. The accommodation of mixed pixels outside the class allocation stage is, however, required because the mixed pixels will contaminate the training data, and cannot be appropriately represented or evaluated by conventional methods.

The key issue is that the whole classification process, and not just the class allocation stage, is constrained by the scale of the pixel. In the training and testing stages, the ground data on class membership are, for example, generally related to the remotely sensed data at the scale of the pixel and so may be fuzzy. As such, if mixed pixels are abundant and the analyst requires the use of a fuzzy classification algorithm to obtain a realistic representation of the classes, then mixed pixels may be a major problem in the training and testing stages. Put another way, it is highly unlikely that mixed pixels present a problem in the class allocation stage only, particularly as the training and testing sets are supposedly representative of the classes over the whole image area. Although mixed pixels are undesirable in both the training and testing stages of a conventional supervised classification, they may be unavoidable and so their effects must accommodated if a theme such as land cover is to be mapped accurately and the map evaluated appropriately. The inclusion of mixed pixels, deliberately or not, into any of the three stages of the conventional "hard" supervised classification is inappropriate and likely to be a major source of error. Similarly, the inclusion of mixed pixels in the training and/or testing stages of a fuzzy classification algorithm may be inappropriate (Foody, 1999).

A variety of approaches may be used to reduce the effect of mixed pixels. In training and testing the classification, for instance, the analyst may deliberately use pixels located in large homogeneous blocks of a class to exclude as far as practicable mixed pixels, or at least the most mixed of pixels. But such pure, or at least near-pure, pixels may be unrepresentative of the rest of the image and be a poor base for training and testing a classification.

Recognizing that fuzziness may require accommodation in the training and testing stages of the classification, as well as in the allocation stage, enables a continuum of classification fuzziness, from completely crisp to fully fuzzy, to be defined. The analyst should be able to use this continuum to help identify techniques appropriate for the data sets and investigation in-hand. A range of possible approaches for classification along this continuum have been indicated; note also that techniques may be applied inappropriately at positions other than those to which they are suited. Techniques to accommodate fuzziness range from rectifying data (e.g., training statistics) to ensure that they are appropriate for use in a conventional methods, to accommodating fuzziness directly into the whole classification process.

The case study illustrated the potential to classify data at various points along the continuum. It used a neural network to derive a series of classifications of varying fuzziness. Resulting accuracy assessments varied markedly even though based on the same data sets and classification approach. The results emphasize the need to recognize the existence of the continuum and for the selection of techniques appropriate to the point along it, within the scope and constraints of a particular investigation.

Finally, the implications to users of the classification, particularly within a GIS, warrant brief attention. Remote sensing is a major source of environmental data for many GISs, yet "hard" techniques are commonly used in extracting the derived information. With much concern in GIS addressing issues such as uncertainty and error propagation (Canters, 1997), the requirement for accurate class labels is evident. Although the mixed-pixel problem has been recognized for a long period, relatively little work has been attempted to address it despite its significance in remote sensing and in linking spatial data sets (Fisher, 1997). To realize the full potential of remote sensing, greater explicit action on this problem is required. Perhaps by embracing more explicitly the issues within a fuzzy geographic paradigm (Openshaw, 1996), strides to realizing more fully the potential of remote sensing may be made. For thematic mapping, fuzzy classification approaches may be an important contribution. The extent to which fuzziness is accommodated, however, can vary and this may affect the real and apparent accuracy of products derived from remotely sensed data.

Acknowledgments

I am grateful to the NERC for provision of the ATM data as part of its 1990 airborne campaign, and for the constructive comments of the referees that helped improve the paper. The neural networks were constructed with the NCS NeuralDesk package.

References

Abrams, M., R. Bianchi, and D. Pieri, 1996. Revised mapping of lava flows on Mount Etna, Sicily, *Photogrammetric Engineering & Remote Sensing*, 62:1353–1359.

Aplin, P., P.M. Atkinson, and P.J. Curran, 1997. Fine spatial resolution satellite sensors for the next decade, *International Journal of Remote Sensing*, 18:1387–1381.

Arai, K., 1992. A supervised Thematic Mapper classification with a purification of training samples, *International Journal of Remote Sensing*, 13:2039–2049.

Barnsley, M., and P. Hobson, 1996. Making sense of sensors, *GIS Europe*, 5(5):34–36.

Bernard, A.C., I. Kanellopoulos, and G.G. Wilkinson, 1996. Neural network classification of mixtures, *Soft Computing in Remote Sensing Data Analysis* (E. Binaghi, P.A. Brivio, and A. Rampini, editors), World Scientific, Singapore, pp. 53–63.

Bowers, T.L., and L.C. Rowan, 1996. Remote mineralogic and lithologic mapping of the Ice River alkaline complex, British Columbia, Canada, using AVIRIS data, *Photogrammetric Engineering & Remote Sensing*, 62:1379–1385.

Campbell, J.B., 1996. *Introduction to Remote Sensing, Second Edition*, Taylor and Francis, London.

Cannon, R.L., J.V. Dave, J.C. Bezdek, and M.M. Trivedi, 1986. Segmentation of a thematic mapper image using the fuzzy c-means clustering algorithm, *IEEE Transactions on Geoscience and Remote Sensing*, 24:400–408.

Canters, F., 1997. Evaluating the uncertainty of area estimates derived from fuzzy land-cover classification, *Photogrammetric Engineering & Remote Sensing*, 63:403–414.

Cohen, J., 1968. Weighted kappa, *Psychological Bulletin*, 70:213–220.

Congalton, R.G., 1988. A comparison of sampling schemes used in generating error matrices for assessing the accuracy of maps generated from remotely sensed data, *Photogrammetric Engineering & Remote Sensing*, 54:593–600.

Corves, C., and C.J. Place, 1994. Mapping the reliability of satellite-derived landcover maps - An example from central Brazilian Amazon Basin, *International Journal of Remote Sensing*, 15: 1283–1294.

Crapper, P.F., 1984. An estimate of the number of boundary cells in a mapped landscape coded to grid cells, *Photogrammetric Engineering & Remote Sensing*, 50:1497–1503.

Cushnie, J.L., 1987. The interactive effect of spatial resolution and degree of internal variability within land-cover types on classification accuracies, *International Journal of Remote Sensing*, 8: 15–29.

Fisher, P.F., 1997. The pixel: A snare and a delusion, *International Journal of Remote Sensing*, 18:679–685.

Fisher, P.F., and S. Pathirana, 1990. The evaluation of fuzzy membership of land cover classes in the suburban zone, *Remote Sensing of Environment*, 34:121–132.

Foody, G.M., 1995. Fully fuzzy supervised image classification, *RSS'95: Remote Sensing in Action*, Remote Sensing Society, Nottingham, pp. 1187–1194.

———, 1996. Approaches for the production and evaluation of fuzzy land cover classifications from remotely sensed data, *International Journal of Remote Sensing*, 17:1317–1340.

———, 1997. Fully fuzzy supervised classification of land cover from remotely sensed imagery with an artificial neural network, *Neural Computing and Applications*, 5:238–247.

———, 1999. Image classification with a neural network: From completely-crisp to fully fuzzy situations, *Advances in Remote Sensing and GIS Analysis* (P.M. Atkinson and N. Tate, editors), Wiley, Chichester (in press).

Foody, G.M., and M.K. Arora, 1996. Incorporating mixed pixels in the training, allocation and testing stages of supervised classifications, *Pattern Recognition Letters*, 17:1389–1398.

Foody, G.M., G. Palubinskas, R.M. Lucas, P.J. Curran, and M. Honzak, 1996. Identifying terrestrial carbon sinks: Classification of successional stages in regenerating tropical forest from Landsat TM data, *Remote Sensing of Environment*, 55:205–216.

Foody, G.M., R.M. Lucas, P.J. Curran, and M. Honzak, 1997. Non-linear mixture modelling without end-members using an artificial neural network, *International Journal of Remote Sensing*, 18: 937–953.

Foschi, P.G., and D.K. Smith, 1997. Detecting subpixel woody vegetation in digital imagery using two artificial intelligence approaches, *Photogrammetric Engineering & Remote Sensing*, 63: 493–500.

Gong, P., R. Pu, and J. Chen, 1996. Mapping ecological land systems and classification uncertainties from digital elevation and forest-cover data using neural networks, *Photogrammetric Engineering & Remote Sensing*, 62:1249–1260.

Gopal, S., and C. Woodcock, 1994. Theory and methods for accuracy assessment of thematic maps using fuzzy sets, *Photogrammetric Engineering & Remote Sensing*, 60:181–188.

Harris, P.M., and S.J. Ventura, 1995. The integration of geographic data with remotely sensed imagery to improve classification in an urban area, *Photogrammetric Engineering & Remote Sensing*, 61:993–998.

Hlavka, C.A., and M.A. Spanner, 1995. Unmixing AVHRR imagery to assess clearcuts and forest regrowth in Oregon, *IEEE Transactions on Geoscience and Remote Sensing*, 33:788–795.

Joria, P.E., and J.C. Jorgenson, 1996. Comparison of three methods for mapping Tundra with Landsat digital data, *Photogrammetric Engineering & Remote Sensing*, 62:163–169.

Kanellopoulos, I., A. Varfis, G.G. Wilkinson, and J. Megier, 1992. Land-cover discrimination in SPOT HRV imagery using an artificial neural network — A 20-class experiment, *International Journal of Remote Sensing*, 13:917–924.

Kent, J.T., and K.V. Mardia, 1988. Spatial classification using fuzzy membership models, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10:659–671.

Klir, G.J., and T.A. Folger, 1988. *Fuzzy Sets, Uncertainty and Information*, Prentice-Hall International, London.

Lo, C.P., 1986. *Applied Remote Sensing*, Longman, Harlow.

Maselli, F., C. Conese, and L. Petkov, 1994. Use of probability entropy for the estimation and graphical representation of the accuracy of maximum likelihood classifications, *ISPRS Journal of Photogrammetry and Remote Sensing*, 49:13–20.

Mather, P.M., 1987. *Computer Processing of Remotely-Sensed Images*, Wiley, Chichester.

Metzler, M.D., and R.C. Cicone, 1983. Assessment of technologies for classification of mixed pixels, *Proceedings of the Seventeenth International Symposium on Remote Sensing of Environment*, ERIM, Ann Arbor, Michigan, 3:1015–1021.

Openshaw, S., 1996. Fuzzy logic as a new scientific paradigm for doing geography, *Environment and Planning A*, 28:761–768.

Schowengerdt, R.A., 1996. Soft classification and spatial-spectral mixing, *Soft Computing in Remote Sensing Data Analysis* (E. Binaghi, P.A. Brivio, and A. Rampini, editors), World Scientific, Singapore, pp. 1–6.

———, 1997. *Remote Sensing. Models and Methods for Image Processing, Second Edition*, Academic Press, San Diego.

Townshend, J.R.G., 1981. The spatial resolving power of Earth resources satellites, *Progress in Physical Geography*, 5:32–55.

———, 1992. Land cover, *International Journal of Remote Sensing*, 13:1319–1328.

van der Wel, F.J.M., L.C. van der Gaag, and B.G.H. Gorte, 1998. Visual exploration of uncertainty in remote sensing classification, *Computers and Geosciences*, 24:335–343.

van Deusen, P.C., 1996. Unbiased estimates of class proportions from thematic maps, *Photogrammetric Engineering & Remote Sensing*, 62:409–412.

Veregin, H., 1996. Error propagation through the buffer operation for probability surfaces, *Photogrammetric Engineering & Remote Sensing*, 62:419–428.

Wang, F., 1990a. Improving remote sensing image analysis through fuzzy information representation, *Photogrammetric Engineering & Remote Sensing*, 56:1163–1169.

———, 1990b. Fuzzy supervised classification of remote sensing images, *IEEE Transactions on Geoscience and Remote Sensing*, 28: 194–201.

Warner, T.A., and M. Shank, 1997. An evaluation of the potential for fuzzy classification of multispectral data using artificial neural networks, *Photogrammetric Engineering & Remote Sensing*, 63: 1285–1294.

Wilkinson, G.G., 1996. Classification algorithms - Where next? *Soft Computing in Remote Sensing Data Analysis* (E. Binaghi, P.A. Brivio, and A. Rampini, editors), World Scientific, Singapore, pp. 93–99.

Woodcock, C.E., and S. Gopal, 1995. Remote sensing of forests: New data layers for GIS, *Proceedings 1995 ACSM/ASPRS Annual Convention and Exposition*, ACSM/ASPRS, Bethesda, Maryland, 2:420–428.

Woodcock, C.E., S. Gopal, and W. Albert, 1996. Evaluation of the potential for providing secondary labels in vegetation maps, *Photogrammetric Engineering & Remote Sensing*, 62:393–399.

Zhang, J., and G.M. Foody, 1998. A fuzzy classification of sub-urban land cover from remotely sensed imagery, *International Journal of Remote Sensing* 19:2721–2738.

(Received 04 June 1997; revised and accepted 19 February 1998; revised 09 July 1998)