Chemometrics and Intelligent Laboratory Systems, 7 (1989) 67–86
Elsevier Science Publishers B.V., Amsterdam — Printed in The Netherlands

Strategy of Multivariate Image Analysis (MIA)

KIM ESBENSEN *
Norwegian Computing Center, P.O.B. 114, Blindern, N-0314 Oslo 3 (Norway)

PAUL GELADI
Research Group for Chemometrics, University of Umeå, S-901 87 Umeå (Sweden)

(Received 18 November 1988; accepted 25 April 1989)

ABSTRACT

Esbensen, K. and Geladi, P., 1989. Strategy of multivariate image analysis (MIA). *Chemometrics and Intelligent Laboratory Systems* 7: 67–86.

Bilinear decomposition (soft modelling using principal component analysis) of multivariate imagery results in: score and loading plots, score images, classification projections and residual images in the scene space. Feature space score plots are used as a starting point for pixel class delineations, followed by iterative scene space evaluation. This is a reversal of traditional image processing practice, which selects training samples in the scene space. The present feature space class definitions can be shown to have certain optimality characteristics with respect to traditional scene space delineations.

After problem-dependent relevant pixel class delineations have been obtained, one can compute corresponding local class PC-models that serve as an alternative basis for problem-dependent classification and sequential segmentation. Multivariate image analysis (MIA) allows interactive exploration and classification of most types of technical multivariate imagery. We present a general strategy for multivariate image analysis, illustrated by a remote sensing showcase.

INTRODUCTION

Most contemporary image processing (IP) imposes constraints on the types of image processing strategies possible. This is intimately related to the type of IP hardware architecture and software that has traditionally been employed. IP makes extensive use of three or more image and bit planes for image arithmetics, overlay displays etc. There is also a strong tradition for primary consideration of the image, or scene space. While this is indeed central to image processing, this has also — and probably largely unwittingly — resulted in a dominant emphasis on image plane operations, which has somewhat restrained the full potential of the application of unsupervised data analysis in IP. Unsupervised classification and exploratory data analytical methods are very useful in situations characterised by sparsity, or even complete absence, of domain-specific knowledge, such as, for example, when analysing a LANDSAT scene without previous 'ground truth' knowledge.

0169-7439/89/$03.50

An important distinction here is that between image processing (IP), i.e. mere image manipulations — however complex — and image analysis (IA), relating to a user-defined objective. Based on more than 10 years' accumulated chemometric experience with multivariate statistics and data analysis for a number of problem types, we show how an analogous strategy can be devised for the case of explorative multivariate image analysis, MIA. We show how these methods also allow supervised results, e.g. classification and pattern recognition.

EXAMPLE

Below we make extensive use of one particular LANDSAT remotely sensed multivariate image. This is intended as a Kuhnian exemplar for purposes of illustration only: We trust that our readers will willingly relate this work to their own subject matter. We specifically wish to demonstrate a completely general methodology, not just another specific image analytical application.

MULTIVARIATE IMAGERY

A graylevel image is comprised of two geometrical dimensions (x and y) for indexing individual pixels, each of which is characterised by an intensity level. This constitutes the simplest digital image unit. Alternatively one could call this a univariate image. The number of pixels is usually very large: 512×512, 1024×1024 and higher. Only univariate statistical methods are required to give a complete description of the intensity distribution. All the remaining information in this type of image resides in the spatial correlation, or contextual correlation.

A multivariate image may conveniently be viewed as a stack of such univariate images, each plane now representing one such intensity variable [1,2]. Alternatively, and more useful for data analytical image decomposition, one may view such an image as an array of pixels (e.g. with two geometrical dimensions), each associated with a p-dimensional vector of variables (p univariate image planes). In the case of remote sensing each such variable would correspond to a radiometric wavelength (or wavelength band). The possibilities for multivariate image analysis hinge critically upon correlations between these p variables. In this case, multivariate methods will be required in order to characterize the covariance, or correlation structure between these p variables. A very important distinction is that between spatial correlation (inter-pixel correlation) and the statistical correlation (inter-variable correlation). An image is thus characterised by a complementary set of spatial and statistical correlations; multivariate image decomposition addresses the interplay between these two types of 'spaces'.

If in addition there are more than, say, four or five such image planes (more variables than four or five for each pixel), MIA will be appropriate for image decomposition. As in any data analytical situation, there is a strong obligation upon the user to pose the scientific problem at hand in a relevant manner. We discuss in depth those requirements that are especially relevant for the case of multivariate image analysis in ref. 3. There is a complete progression from the simplest two-dimensional multivariate image unit upwards toward more complex higher-order data arrays, with higher meta-dimensionalities in both variable directions as well as in pixel directions [3]. An overview of methods of multi-way data analysis is given in ref. 4. The strategies for multivariate image decomposition developed below will also be relevant to these higher-order applications; the present paper will mainly address the basic type of two-dimensional imagery, however.

Many types of scientific and industrial endeavors produce output of a two-dimensional image nature; any such array can be viewed as an image, if its size is large enough [2]. Multivariate image analysis is an important part of multi-way analysis with some special properties that make it particularly user friendly.

EXPERIMENTAL

The example used in the present paper for illustration is a LANDSAT scene in 7 wavelength

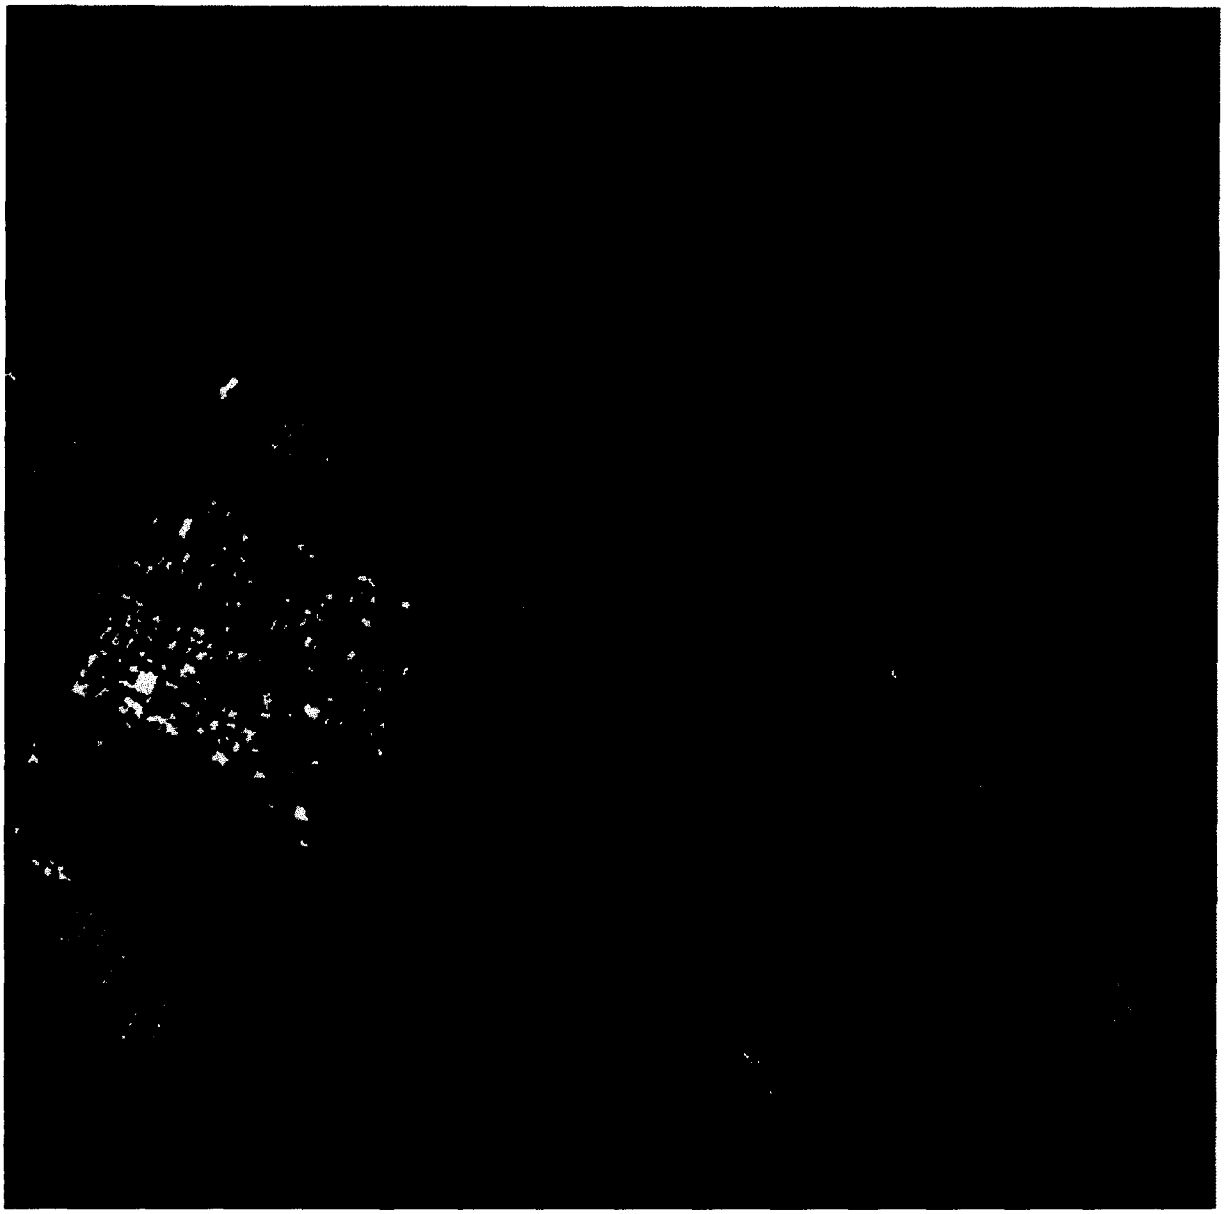

Fig. 1. One channel image (Riyadh scene) for reference to scene space displays below (here used as the red channel of the RGB display). Developed urban areas in left center; low elevation areas running diagonally NW–SE (cf. Figs. 4 and 5); high elevation shield areas in upper right part. Scene-space class of Figs. 6 and 7 is seen here as white overlay.

bands of the city of Riyadh and parts of the surrounding desert in Saudi-Arabia (Fig. 1). The size of the scene is 512×512. The analysis was done one a Compaq 386 20 MHz microcomputer with math coprocessor and EGA screen. The microcomputer was enhanced with a REVOLUTION 'Number Nine' [5] hardware card for graphics processing, giving RGB output for the video monitor. A high resolution Mitsubishi colour monitor was used for image display. Hard copy dumps of the images were made on a Tektronix 4693 D colour plotter; all figures in this paper (except Fig. 2) were created in this fashion.

Data bookkeeping and image processing was done using the ERDAS software package [6]. The MIA module that fits together with this package was written by us in FLEX, a particularly useful FORTRAN dialect. This FLEX is especially developed by ERDAS for the MS-DOS operating system in a package called TOOLKIT [6].

MULTIVARIATE IMAGE ANALYSIS (MIA) APPROACH

We present below our approach towards MIA. What may at first look like just another application of the well-known principal component (PC) transform in the area of image prosessing, will turn out to encompass a complete reversal of its traditional use. We want to be specific here: the MIA approach is unequivocally distinct from the mainstream image processing tradition. MIA constitutes a much more versatile and flexible tool for the creation of one's own image analysis strategy.

The MIA approach contains the following operations:

- calculating principal component scores (score images) and loadings (vectors);
- scatter plots of scores (or loadings) against each other;
- selection of classes on the score plots by user-defined masking; joystick or keyboard arrows;
- brushing of classes in multiple score plot split-sceen display;
- projection (transfer) of the pixels in the feature-space classes to the corresponding scene space locations;
- calculation of local PC-models, as determined in the score plots;
- calculation of residual images with respect to such local models;
- auxiliary functions for overlay masking, overlay toggle, colour slicing etc.

The essential capabilities of the MIA approach are shown in the illustrations below.

It has been a prime objective of the development of MIA that it should be able to run on all available hardware that is dedicated to IP/IA. We have focused in particular on the possibilities for carrying out IA on IBM PC/AT compatible hardware in order to develop an 'image analysis Volkswagen', if possible. In spite of the large amounts of data treated in image processing and the intensity of some of the more intricate, necessary calculations, MIA has reasonable processing times, even on modest microcomputers. For a 7-channel image of size 512×512, it takes less than 15 minutes to go from the raw data to score and loading plots on an IBM AT; this constitutes the heaviest calculation demand in MIA. These processing times are at least halved on the 386 series computers. (If mini or mainframe computers are available, processing time is an irrelevant issue.)

We present the workings of MIA as a top-down illustration of how one uses the implemented software; this article thus both details the theory behind, and the practical use of, the MIA approach.

BILINEAR DATA ANALYSIS OF MULTIVARIATE IMAGERY

The starting point for MIA is a traditional principal component analysis of the image to be analysed [7]. The PC scores are used to construct 2-D scatter plots (Fig. 2). The score plots allow the user to inspect the inherent data structure in the image. Fig. 2 shows two examples of such a PC-plot, here component 2 against component 1 (termed PC12). Obviously, the data analytical interest will, in many problem formulations, not necessarily be addressed to these two first components alone, but will also require higher-order components, e.g. PC13, PC14, PC23.

It is better for visual assessment to use a colour-slicing scheme for these score plots. We use a self-explanatory succession of colours to represent an increasing density of pixels with identical score-pairs. A grading from cold colours at the margins towards warm colours at the centers of the score distributions reveals modes and 'back-

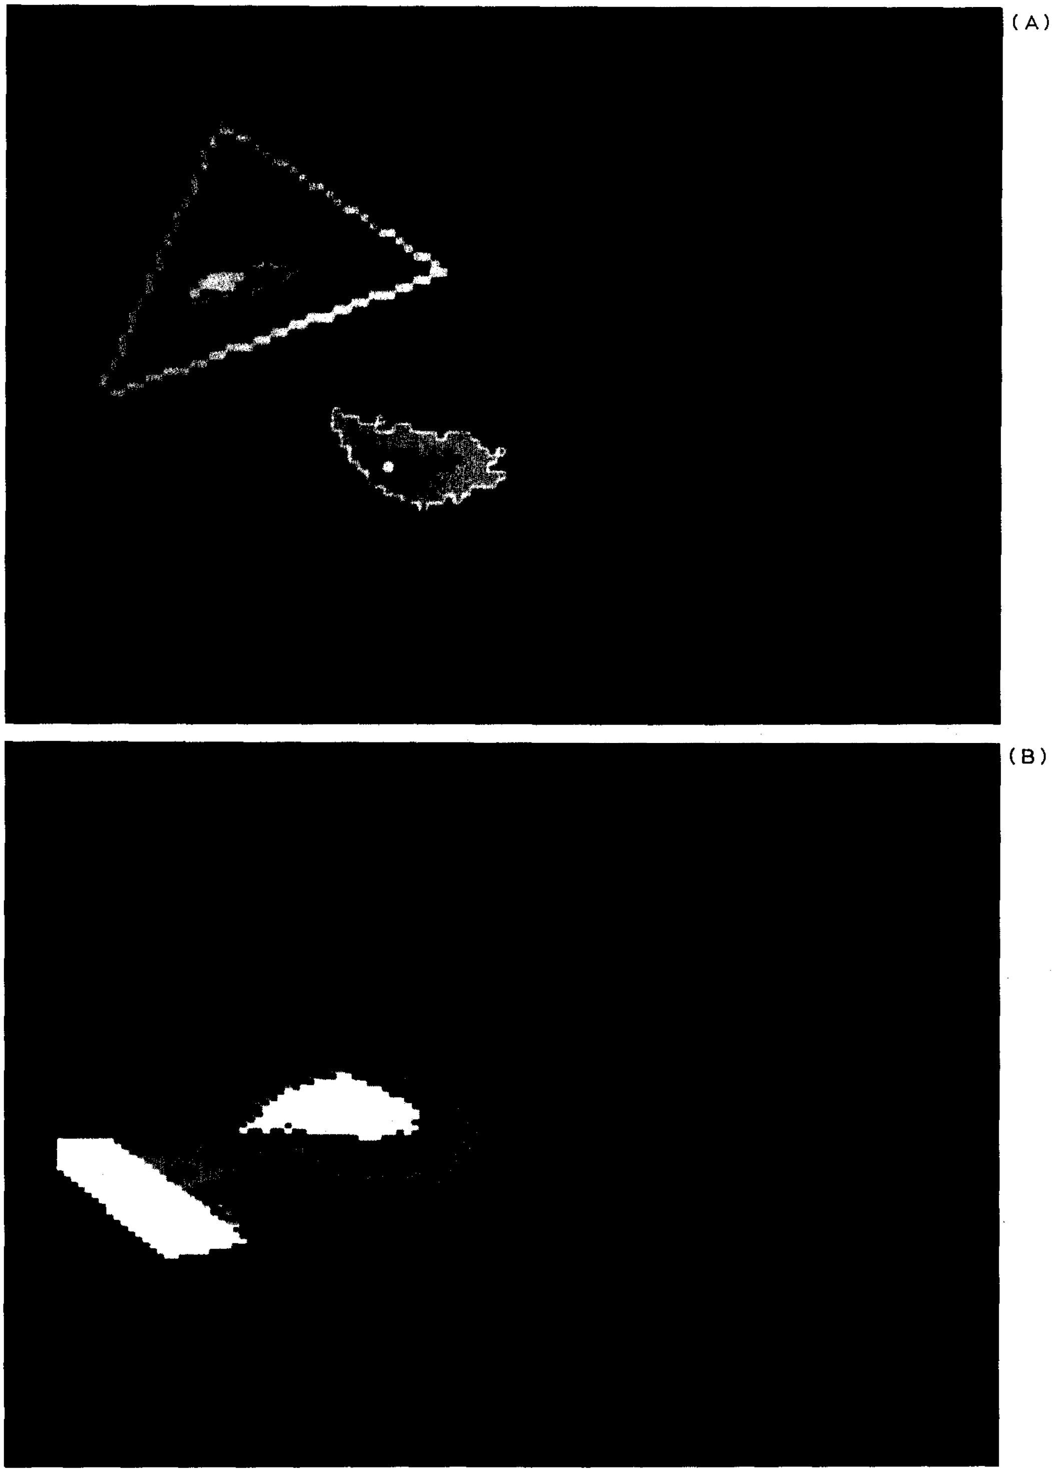

Fig. 2. Scatter plots of PC2 vs. PC1 from a standard LANDSAT scene. Figs. 2A and B represent PC12 plots of two different LANDSAT scenes with two different bimodal distributions. The most coherent modes have been designated as a pixel class (triangular envelope in Fig. 2A) by the image analyst; in Fig. 2B another class has been masked completely. These three-dimensional histograms have been colour-sliced to contour modes and trends of distributions; see text for details.

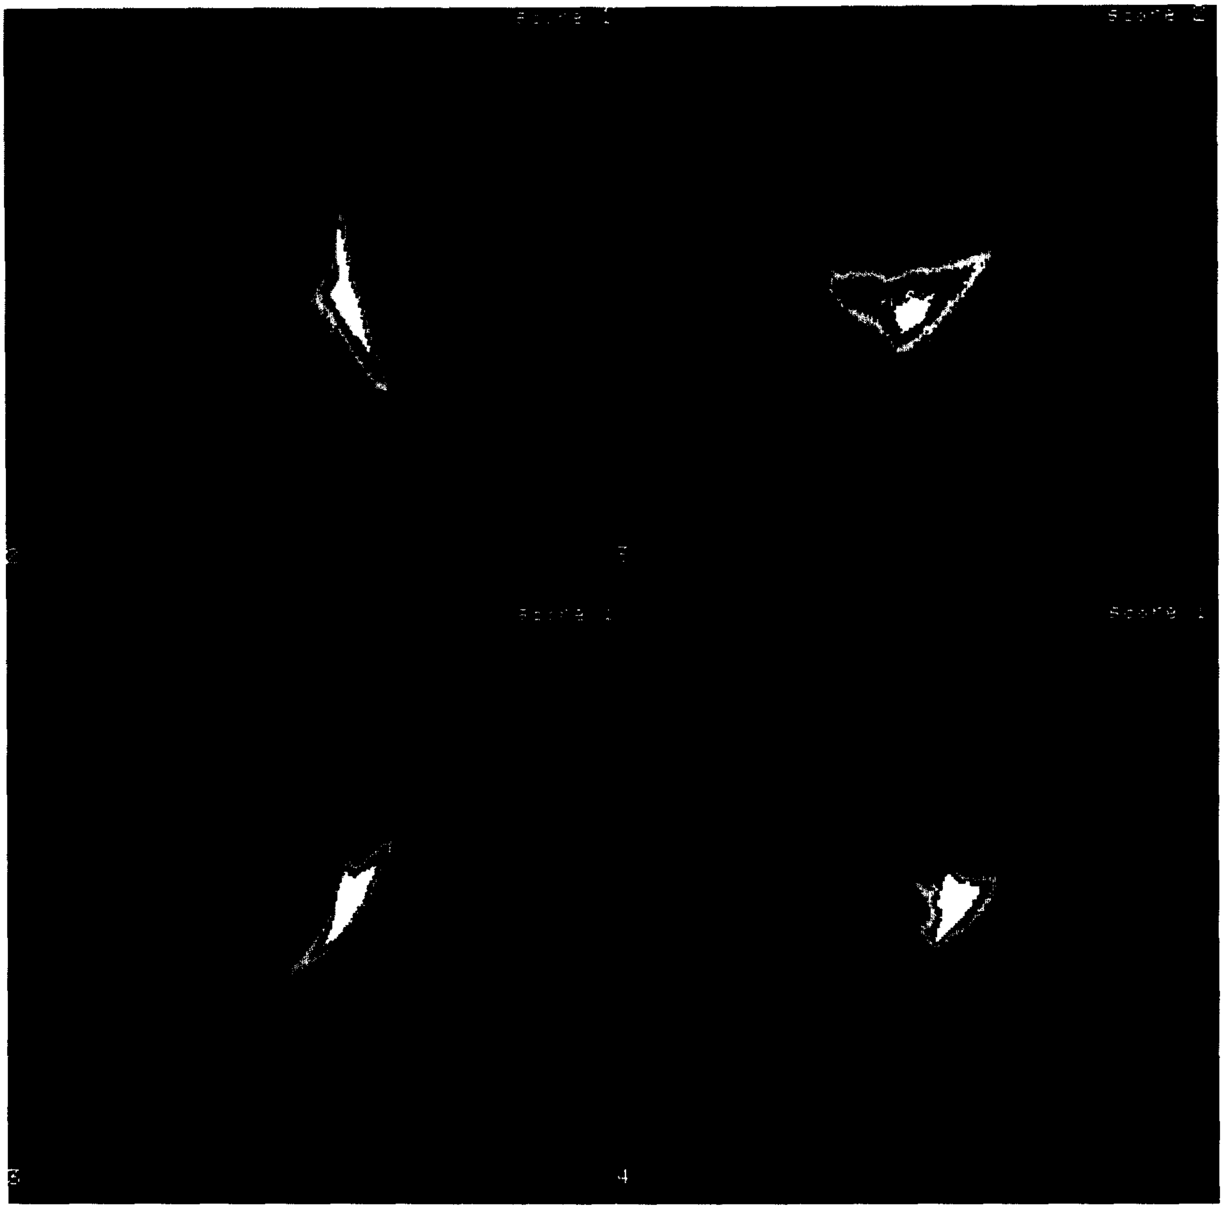

Fig. 3. Split-screen option of four matching PC-plots, cf. Fig. 2. The IP software used to construct this display layout makes for a downward orientation of the Y-axis. PC-axes have been numbered at their positive margins, respectively. This example, Riyadh scene cf. text and Fig. 1, is used for the remaining examples below. Note the detailed statistical data structure in this PC-score space, which forms the basis for delineating data classes.

bones' in this type of three-dimensional histogram. This colour-slicing is illustrated in Fig. 2, where, in this particular example, one observes a typical bimodal distribution, i.e. water vs. land-cover in LANDSAT scenes. (Note: Fig. 2 is the only plot not from the Riyadh scene, but is used for its particularly excellent illustrative characteristics.)

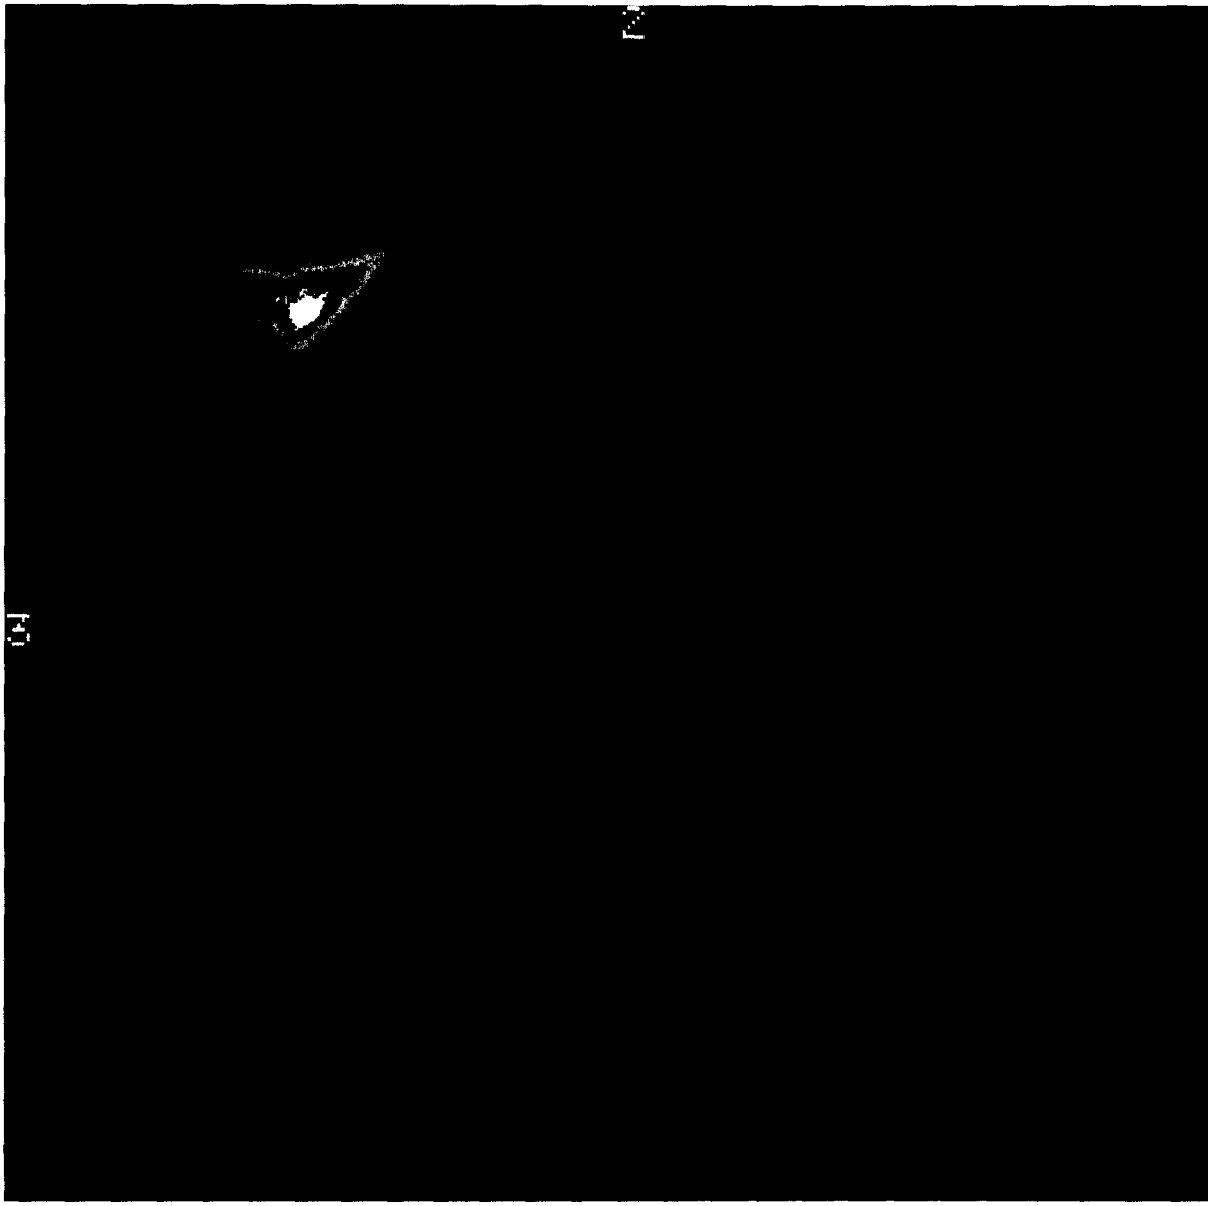

Fig. 4. Detailed PC23 score plot of the Riyadh scene. Class-masking of the upper lefthand mode (masked in olive), cf. upper right panel in Fig. 3.

Fig. 3 shows how a split-screen option is used for displaying four PC-plots, viz. PC12, PC13, PC14 as well as PC23 (default in MIA). The MIA implementation allows either this standard four-panel setup, or any specific combination of two component scores in a bivariate display (default: PC12), cf. Fig. 4.

The image analyst may delineate any 'interesting' pixel aggregation in this plot. In effect, one is delineating a tentative data class corresponding to

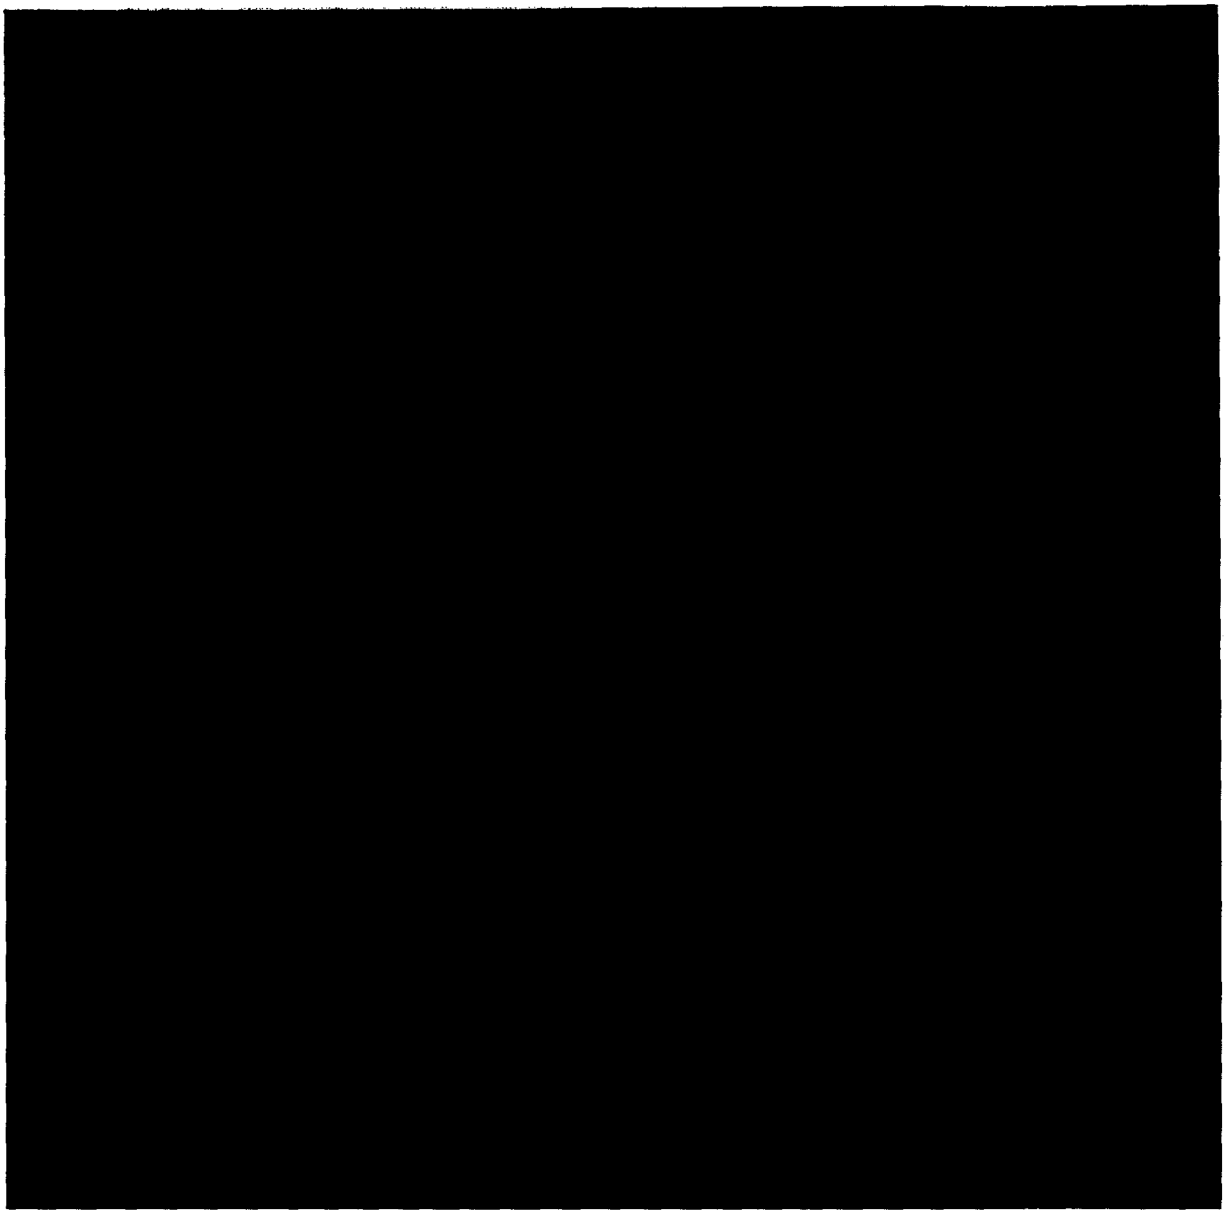

Fig. 5. Scene-space image of the class masked in Fig. 4. See Fig. 1 for a reference overview of the Riyadh scene.

pixels with similar spectral fingerprints. Fig. 4 shows one such pixel class. This class definition takes place in what we have termed the score space. This step constitutes the salient backbone of the 'reversed' mode image analysis in MIA. The complementary scene space information is subsequently presented to the user by a scene space display, in which all the pixels corresponding to this tentative class are flagged, as shown in Fig. 5.

These two 'projections' (score space/scene space) are complementary throughout any MIA image analysis. No score space class can be considered without its complementary scene space layout. (There does not always appear to be a

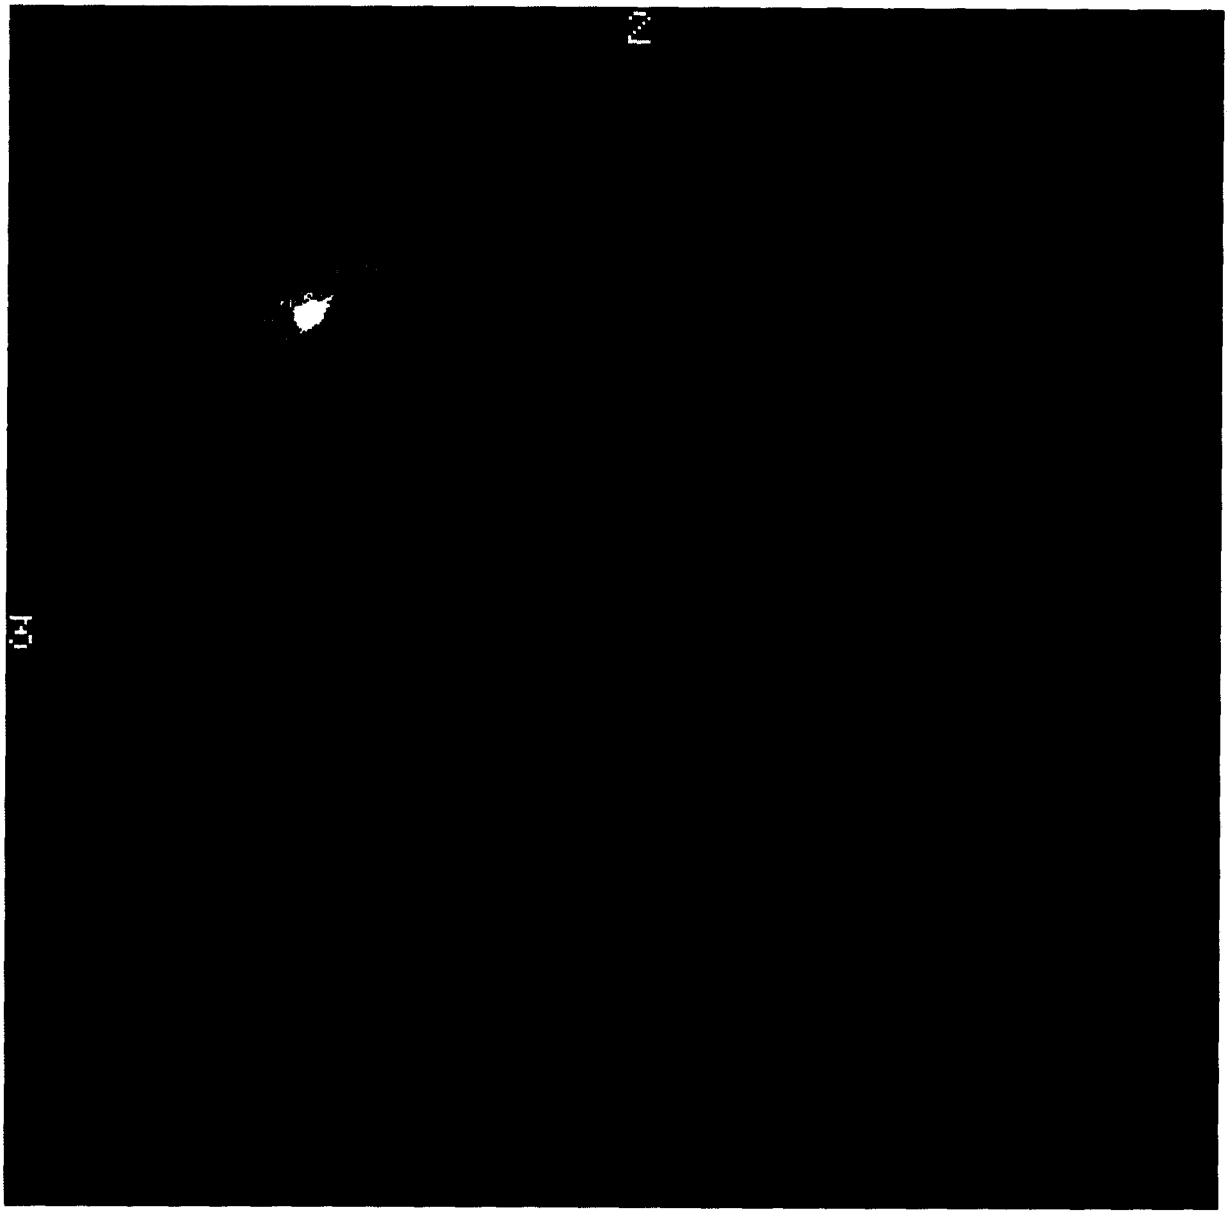

Fig. 6. Class-mask for a low-density area in the same PC23 plot as in Figs. 3 and 4.

similar obligation in the traditional IP approach, in which one may sometimes observe how a training class selected in the scene image is used as input for a 'classifier' algorithm etc. without proper evaluation in the converse feature/score space.) The MIA approach is intrinsically based upon this duality of representations of both feature/score space as well as scene space: the feature space is treated below in more detail.

MIA presupposes a series of iterations between these two spaces before 'meaningful' — i.e. problem-dependent — class definitions are to hand: it is precisely this interactive evaluation of data structures in the complementary sets of score and

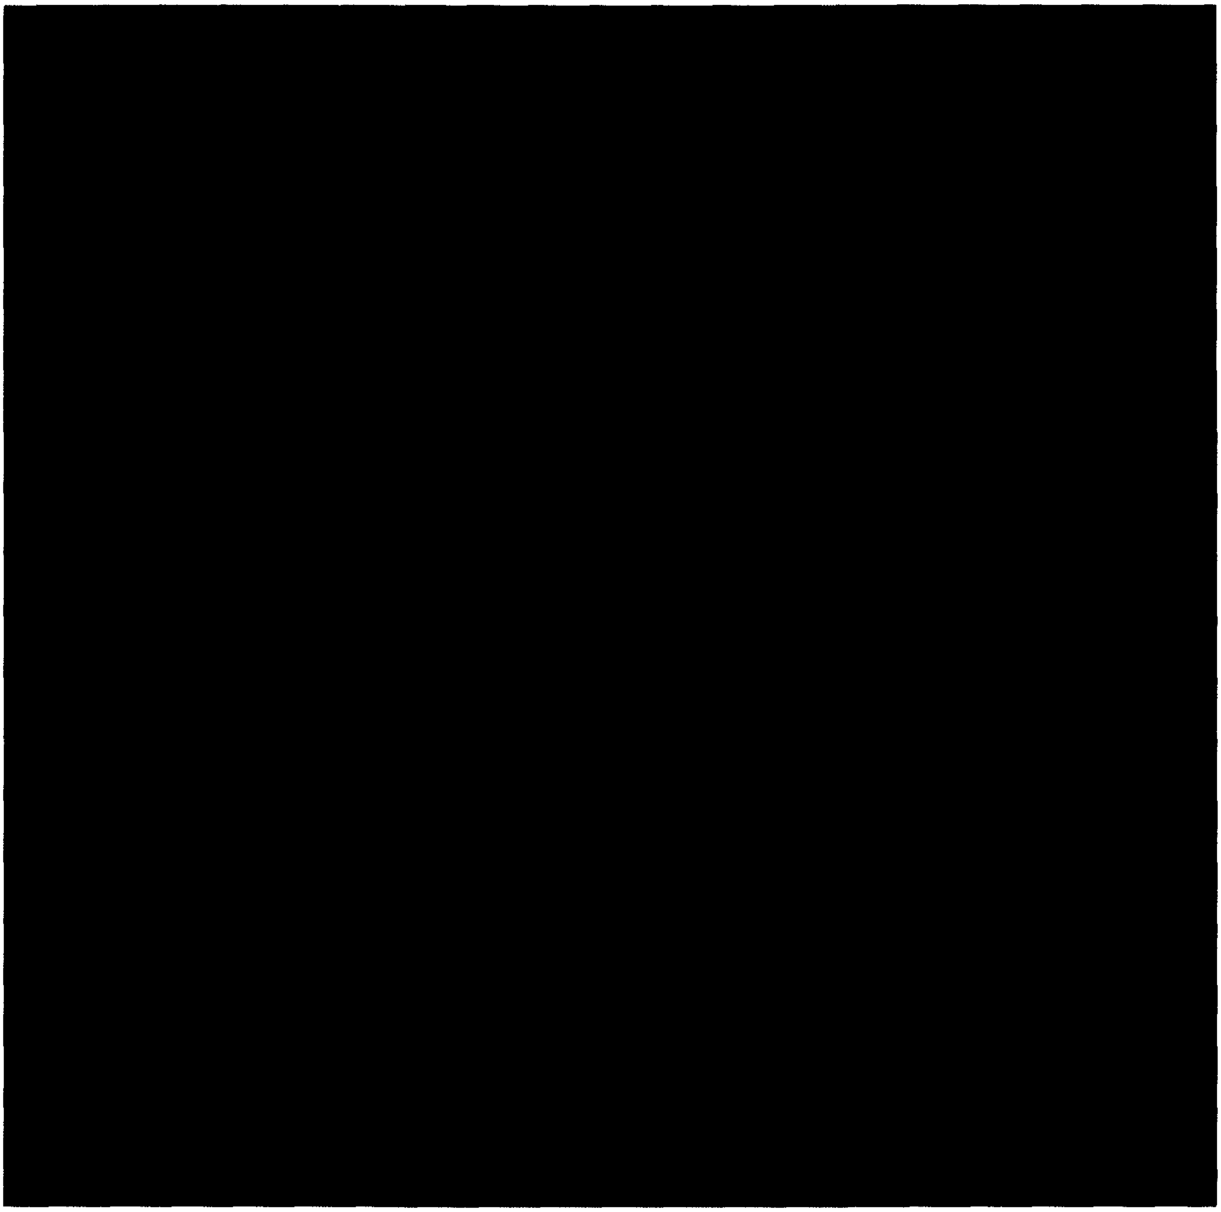

Fig. 7. Scene-space image of the class masked in Fig. 6. See Fig. 1 for a reference overview of the Riyadh scene, in which this class layout has been transferred to the overlay (white).

scene spaces that allows the image analyst to use domain-specific knowledge in interpreting the 'meaning' of each pixel class. The user assigns meaning to the class by his/her own particular domain knowledge. The MIA approach vehemently opposes any attempt at preconceived, algorithm-embedded substitutes for this crucial human analyst–graphic display interaction. This is perhaps the strongest aspect of how we intend MIA to be used in exploratory image analysis.

This class 'meaning' is often obvious in the scene space when one or more of the dominant

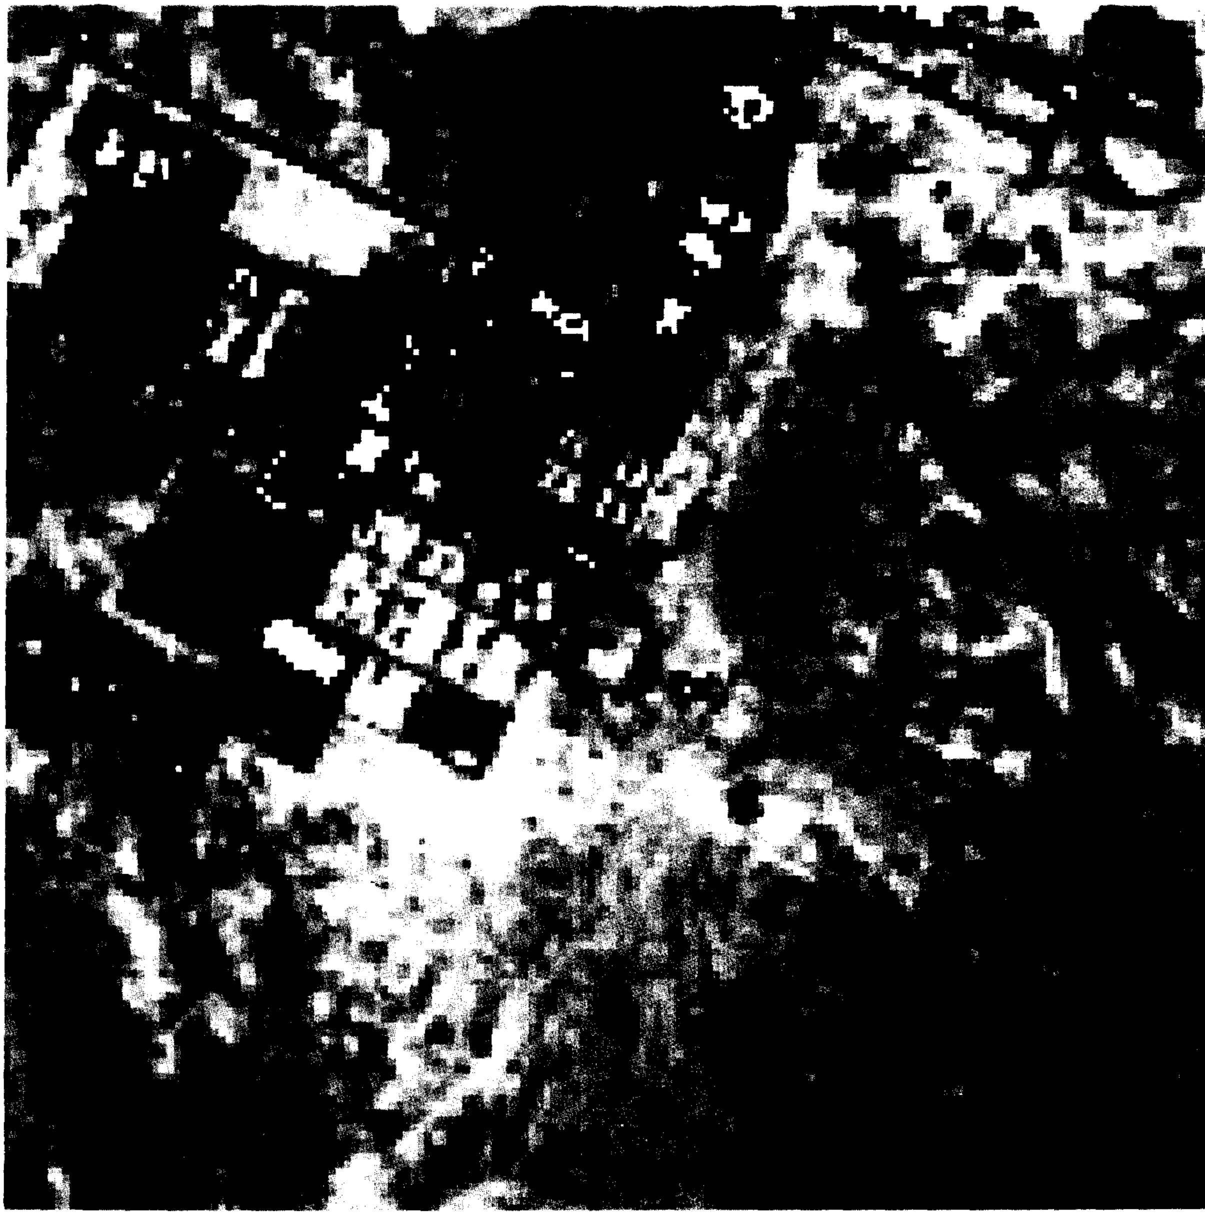

Fig. 8. Enlarged lower left part of the scene in Fig. 7, with overlay of yet another illustrative outlier pixel class. This class apparently corresponds to the interiors of certain buildings, compounds etc. (The interpretation is not the issue here, whereas the possibilities of interpretation are.) A training sample delineation in this scene image — to follow the traditional IP approach — with similar modelling coverage would be extremely difficult to obtain; see also text.

modes of the types presented in Figs. 2 and 3 have been designated. This use of MIA allows the user to ensure that all pixels are included in the relevant class. When delineating a class in the score plot, the user can be absolutely certain that all pixels belonging to the chosen class are included, no matter how they are distributed in the scene space. It is easy to see how the reverse IP ap-

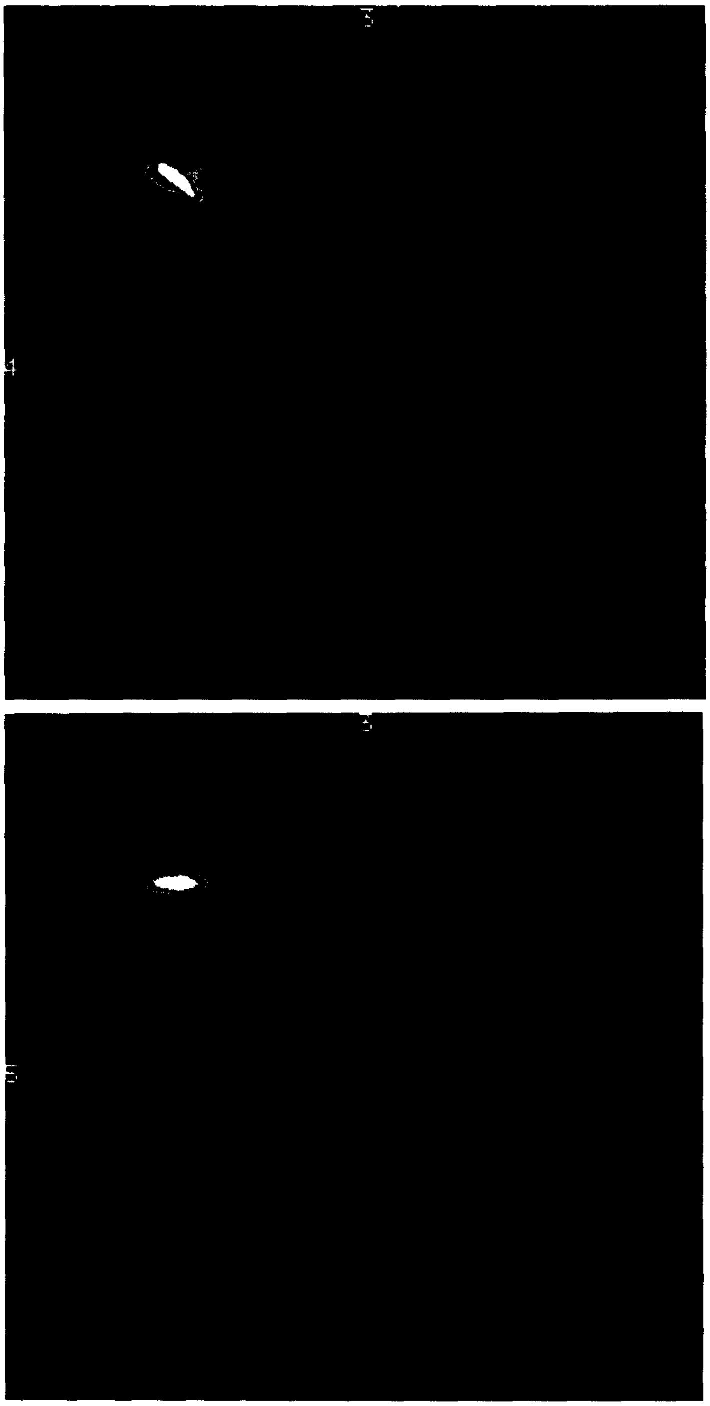

Fig. 9. Composite display of matching sets of PC34 and PC35 score plots, corresponding to the same original scene (Riyadh). Note how the shift from principal component No. 4 to 5 gives way to an enormous decrease in data structure. Figs. 3 and 9 go a long way to kill the myth that remote sensing imagery is characterised by 'quasi-multi-normal' distributions, except e.g. this fortuitous PC35 plot.

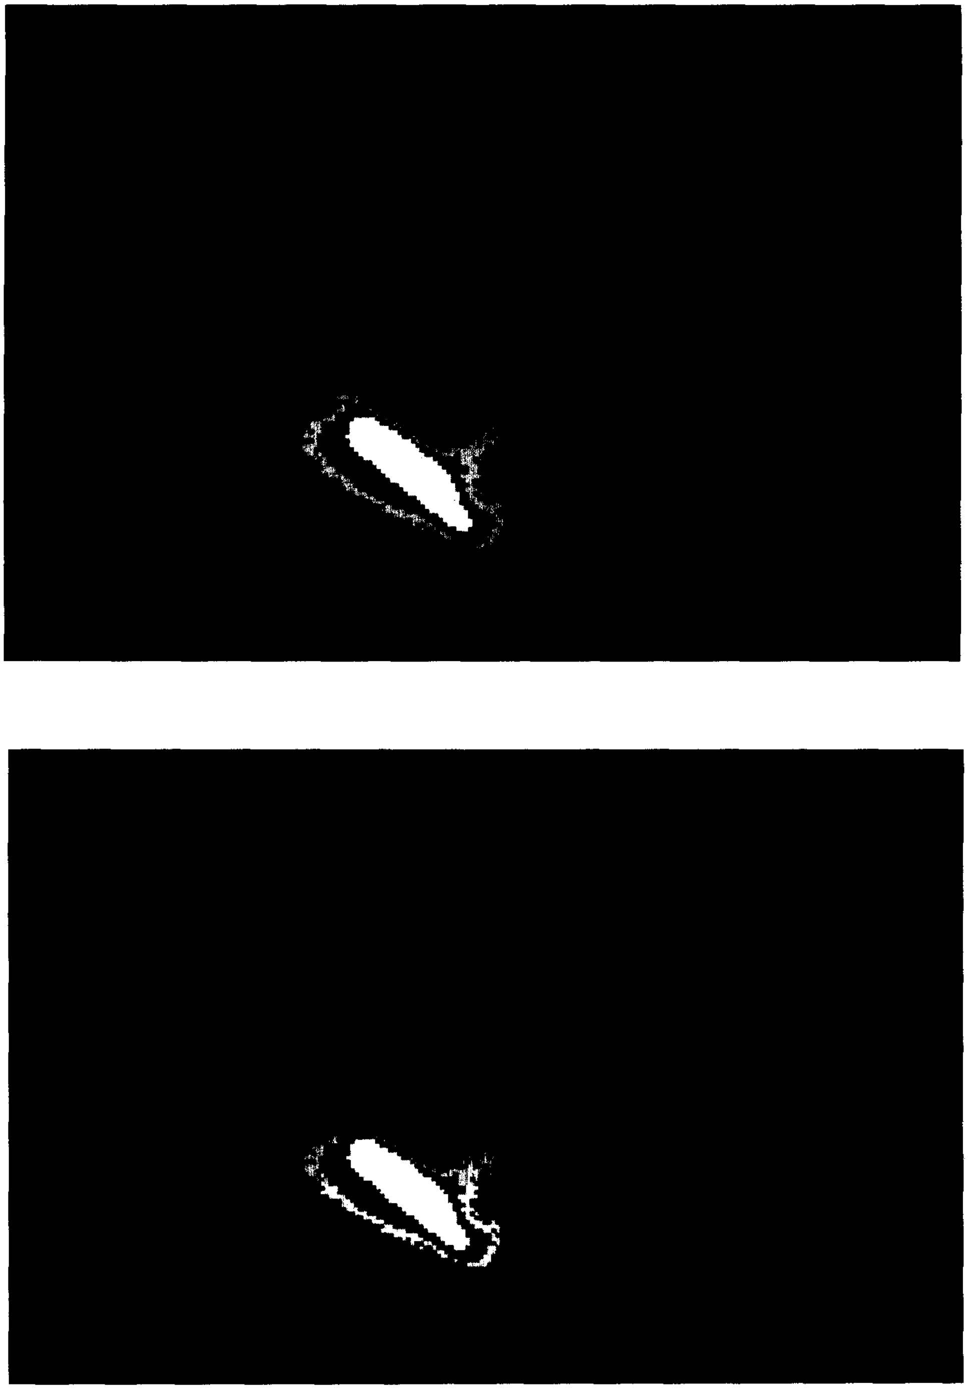

Fig. 10. Class delineation of a low-density area with 'trend' in the PC34 plot in Fig. 9. This class has been delineated because of the trend indication in the center part; cf. with and without class mask (olive).

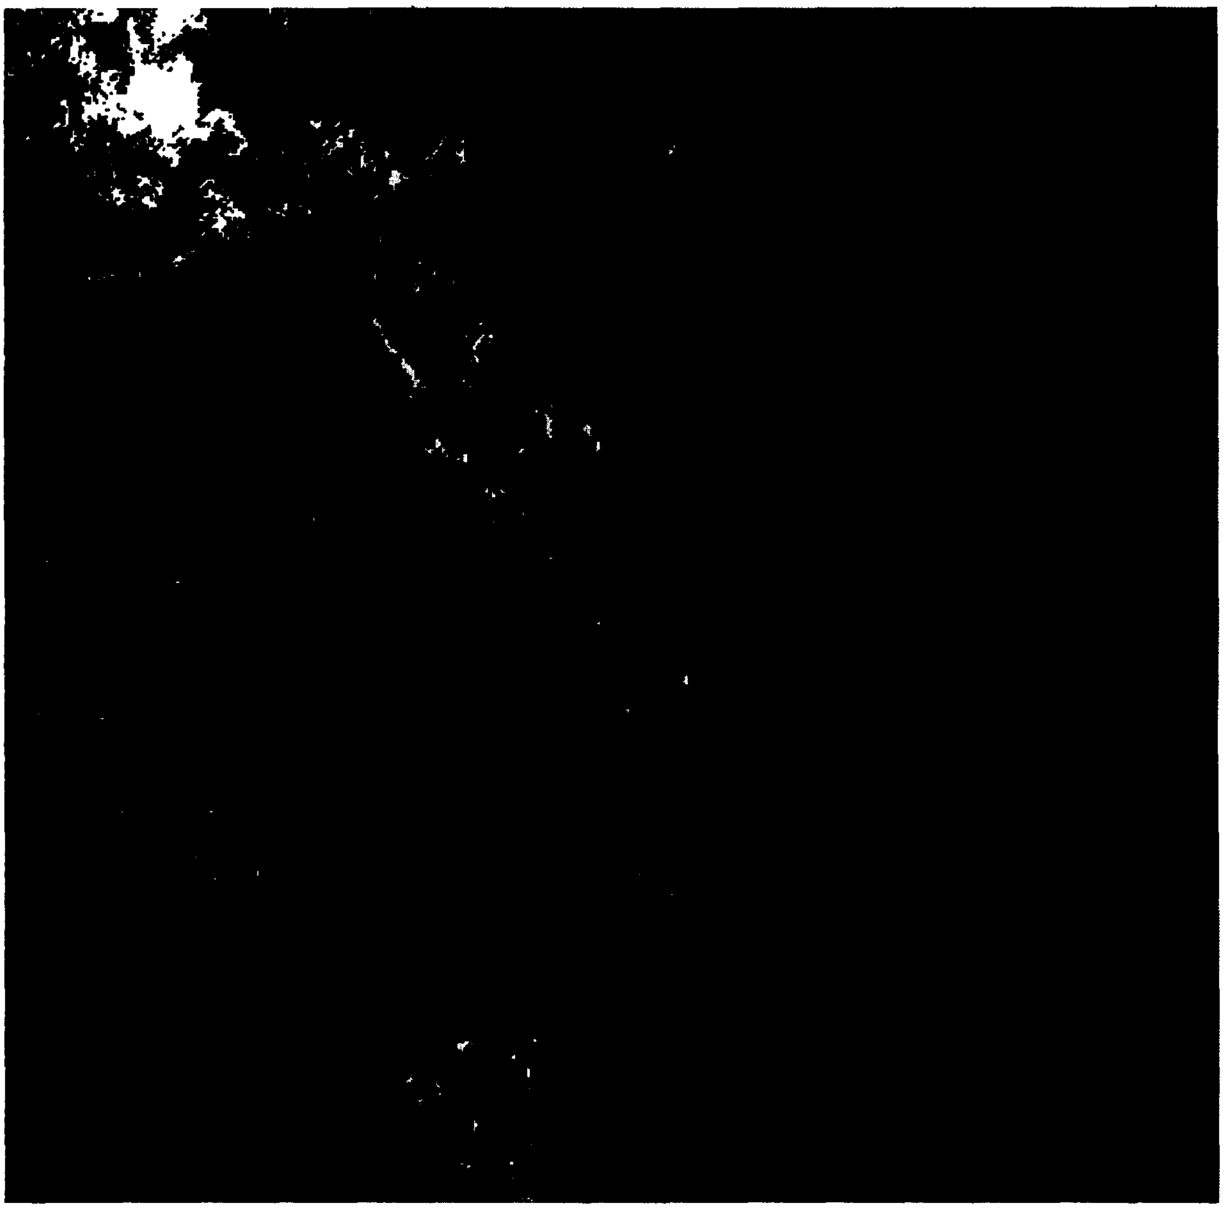

Fig. 11. Scene-space image of the class masked in Fig. 10.

proach can never be similarly certain that a particular scene space class delineation will be absolutely representative; indeed this is almost universally uncertain, see Figs. 1, 8 and 11.

Figs. 6 and 7 illustrate the delineation of a class with a very low density of pixels (cf. Fig. 3). Fig. 7 gives the corresponding scene space layout for these pixels — compare this with the scene space layout in the original scene, Fig. 1.

A similar, low-density class was also the basis for the zoomed scene space display in Fig. 8.

Here one observes how it is possible to achieve a very detailed class definition with the MIA score space class concept.

Fig. 12. MIA module: 'Overview' display of the standard split-screen layout of four different aspects of the data structure in the Riyadh scene. The statistical feature space is to the left, with the corresponding scene space information to the right. PC-plot in the upper left panel (PC12 in the present example), with the corresponding loading plot in the lower left panel. The two panels to the right give a compressed version of the pertinent PC-score images. This layout of the decomposed image structure allows the user to concentrate on the problem-dependent interpretation.

Fig. 9 illustrates the subtlety of the various PC-score plot options. Whereas the PC34 plot gives ample opportunity to evaluate the data structure, the equivalent PC35 shows a surprisingly homogenous distribution of scores. A strict 'exploration strategy' for browsing through the higher-order PC dimensions is often very much a necessity here.

In Figs. 10 and 11 yet another 'interesting' data class has been delineated. Fig. 11 gives the corresponding scene space layout; cf. also Fig. 1. In these two figures (1 and 11) it is again particularly clear how the MIA approach allows a fundamentally superior class definition in comparison to that which can be obtained by the traditional scene space training sample delineations. With MIA, one obtains a guaranteed complete class containing all pixels with similar spectral signatures regardless of whatever diverse spatial layout these pixels may display in the scene space.

Loadings from principal component analysis contain information about the importance of the original variables (channels) for each of the principal component dimensions. A very useful way of studying loadings is by plotting them against each other in 2-D (or 3-D) plots, which are only projected graphic visualizations of the correlation matrix [8]. In these plots, the importance of loadings and the degree of (dis)similarity between variables can be evaluated. Loading plots results from any combination of two or three principal components. We term this the statistical feature space proper; usually, however, this term is used to connote both the score and loading space simultaneously, i.e. in traditional IP parlance.

It has been found particularly useful to adhere to a systematic evaluation of the rapidly bewildering possible permutations of pairs (triplets) of such loadings. We recommend the use of a systematic serial approach such as PC12, PC13, PC14, PC23, PC24 etc. which will tell the analyst all about the data structure of the original variables in a very few component dimensions. The intervariable correlations are revealed in the loading relationships. We have designed a simple graphic loading plot assessment option in the MIA module as well (Fig. 12). Obviously this option increases in usefulness when truly multi-channel imagery is analyzed; the present 7-channel LANDSAT scene only serves as a modest low-dimensional illustration. Technical imagery may certainly be expected to display a (much) more multidimensional nature in this context.

It may perhaps not be immodest to view the present approach towards multivariate image analysis as a signpost of a (new) methodology that is seeking suitably (highly) complex problems, certainly both within as well as outside chemometrics. Whatever the dimensionality and the complexity of the problem formulation, MIA is ready, in theory. Of course, as the dimensionality of the imagery goes up, one may encounter specific hardware-constrained calculation time problems, but these are simply technicalities that can either be easily remedied by working on a more powerful computer, or by patience carrying you through on your modest PC. In both situations, the user meets the exact same MIA interface; the image analytic methodology is invariant.

An image may thus be decomposed into the scene space (image space) and the statistical space; this latter comprises both the score space and the loading space. The statistical space is decomposed completely along the lines of bilinear decomposition [8]. Fig. 12 shows how it is possible to employ both a score plot and the matching loading plot, as well as the two corresponding score images in the scene space. This is the most comprehensive assemblage of the decomposed information in an image. Each such four-panel 'overview' will allow the user to integrate the information gleaned from each of the partial decompositions, viz. the scores/loading decomposition vs. the PC-transformed scene imagery. We contend that this type of display will serve well for problem-specific interpretation and the like.

STRATEGY OF MIA IN IMAGE DECOMPOSITION

MIA can be used in identical fashion either to delineate a dominating mode in any particular score plot, or to delineate outlying classes, cf. above. These latter will of course be the ones that are not prominent in the image itself. Pixels belonging to such classes will be optimally describable in the score space, precisely because of their sparsity and their irregular distribution in scene space. There is a complete gradation between these two situations. Experience from both data analysis in general and remote sensing IP show that this type of image analysis is very strongly problem dependent, i.e. the image data analyst is obliged to specify some form of objective for the image de-

Fig. 13. Residual image, resulting from subtraction of the first four components from the seven-channel LANDSAT Riyadh scene. Structural information (apparently both statistical and spatial) is still present in the residuals. Olive and brown: 'within-model' pixels. Black, blue and white: farther away 'outside-model' pixels.

composition; just 'playing around' with the powerful IP tools is not advisable.

The least interesting option of MIA is the above major mode modelling, Fig. 2. More subtle class delineations, e.g. those illustrated in Figs. 4, 6 and 10, will do much more justice as illustration of the powerful possibilities for problem-dependent image analysis in MIA.

Used in a first pass, these IA facilities serve to find the objective major classes present, which are then immediately considered as objects for interpretation. Used in this fashion, MIA works in an interactive, exploratory mode. The results of this introductory analytical phase will be the segmentation of an image into several coherent, local data classes.

In particular, any 'interesting' (understandable, interpretable) local PC-model may serve to reclassify an entire image. By reclassifying the whole pixel array in this fashion, the image can be segmented according to this local model only. By successive reclassifications based upon a series of all such local models found — and duly interpreted by the image analyst — the total image will gradually be segmented accordingly. It is important to note that the final image classification reflects this gradual build-up of the interpretation of the total image structure, NOT a shotgun, once-only pass through a particular classification algorithm etc. The image analyst only has to watch her/his step in not allowing any overlap between different class delineations in the score space (at least not before acquiring some experience). Simple Boolean image plane arithmetics may serve as a useful guardian in this context. (In the present ERDAS implementation one has ready access to suitable GIS facilities [6].) Used in this fashion, MIA works in a supervised, pattern recognition mode.

For each individual local model, one may compute the corresponding residual images, which comprise a colour-sliced version of the residual distance for each pixel with relation to the local model. Fig. 13 presents one such image, in which one may assess the residual model distances for all pixels. The pixels present within the local model are distinguished by a neutral colour, while pixels with gradually greater residual distances are represented by a grading 'off-scale' colour scale. The particulars of this colour scale are relatively uninteresting; it is the spatial structure of the pixel layout in the scene space that will attest to the remaining data structure after the local model has been subtracted.

This type of image is very useful during ongoing analysis of any complex image. The user may conveniently try out any apparently suitable class-complexity for a given local model of interest; the residual image will carry important information as to the validity of this choice.

By diligent use of (combinations of) several of the above MIA options, the user will have ample freedom to choose her/his own particular approach for image decomposition, relevant to the problem definition at hand. It is emphasized that the MIA approach is a flexible tool for the iterative creation of one's own image analytical strategy, not another ready-made algorithm, or fixed methodology, for image processing.

DISCUSSION

We have not presented just another specific study involving dedicated image processing cum domain-specific interpretation. In the remote sensing literature, for example, examples abound which use a variety of standard multivariate statistical techniques (e.g. cluster analysis, discriminant analysis, training sample distribution fitting etc.), on a variety of specific interpretative problem formulations, but which (very nearly) all stay in the scene-space training sample tradition. MIA should by now be clearly recognisable as the directly opposite approach. For this reason we have deliberately left out many potential references, also because they mostly relate to other scientific communities, much more than to the readers of this journal (see below for a few, very important exceptions [7–15]). We are confident, however, that the principles behind the general MIA approach will have been appreciated even with the use of the remote sensing exemplar.

An interesting point raised by an astute referee concerns the fact that the principal component score plots relate to more or less 'abstract factors' (in factor analytic parlance), with emphasis upon these being 'not easily' interpretable. The point here is exactly that in-depth interpretation of the principal components is not an absolute prerequisite for the informed use thereof! If you can assign meaning to the MIA loading plot, this will most certainly assure a deeper insight into the data structures. If you, for some reason, cannot do so, the score and loading plots are still quite legitimate objective representations of these very same data structures. Increased understanding and interpretability of the ultimately derived imagery is what constitutes the proof of the pudding.

For the record: we are not preaching a mindless pragmatism in the above remarks. Without going into the depths of the factor analysis vs. principal components debate, we assure our readers that the chemometric experience behind the MIA image approach is indeed soundly based. We refer to refs. 1–4 and 6–8 in which may be found all the necessary referenced substantiation.

CONCLUSION

We use the statistical correlation aspects of multi-channel spectral information in multivariate imagery to take over from data analytic experience some simple strategies for exploratory data analysis and classification, also for the image regimen. Principal component analysis allows data reduction, where the most relevant information is condensed in a few principal component (or score) images. Scatter plots in this score space allow optimal delineation of problem-relevant classes, which forms the backbone of the 'reversed' MIA approach. Overlays, colour slicing and many other techniques adapted from traditional image processing, enhance the visual content of this information. A variety of derived imagery results from construction of local PC-models based upon the MIA concept. Loading plots allow the user to investigate the importance of the variables for the model constructed (this will become increasingly important for tomorrow's multi-channel imagery). Residual imagery is also shown to be of great importance. We show that both unsupervised and supervised multivariate image analysis is feasible.

These features are illustrated by a remote sensing example, but the general principles behind MIA are much more potent; indeed we perceive by far the broadest application field as being outside remote sensing, e.g. within chemometric analysis in optical and electron microscopy and other imaging techniques in the laboratory [1,2].

MIA data analysis is almost completely visually oriented, a factor that greatly contributes to its user-friendliness. We perceive MIA only as a first attempt towards a more human perception based general decomposition technique that we can now get only a few glimpses of [1–4]. This decomposition may be directed towards typical imagery [1,2] or towards higher-order data arrays [3,4]. The salient aspects will be the availability of the kind of derived imagery illustrated in the present work. For example, we directly aim at the possibilities of 'slicing up' higher-order imagery, with the present 2-D multivariate image decomposition 'units' for higher-order data arrays of the 'image type' [3], while the opposite type of arrays, with more variable ways than object ways, need more formalised approaches [4].

RELATED WORK

The gamut of contemporary image processing comprises a plethora of literature, most of which tends to fall outside the scope of this journal because of its subject matter (remote sensing, optics, tecnical IP etc.). With the risk of seeming almost offensive we originally presented only the first eight entries in the literature section below, for no other reason than that MIA was developed completely on this basis alone during the last 2–3 years. Indeed one reviewer took offence, however, and demanded a computer literature search. We also give an additional 7 references, carefully screened from the list of 95 references that showed up. Well over 95%, or more, of what's catalogued under the heading multivariate (...) image (...) analysis still remains within the traditional IP regime, and is thus not relevant for comparison with the present approach. On the other hand, all due respect goes to refs. 9–15, which appeared in many a different setting than chemometrics and remote sensing, as focused on in this paper.

The works [9–13] deal with various aspects of more or less similar approaches to that described here, though generally much imbedded in the particulars of their specific subject matters spanning: electron microscopy/delineation of complex molecular structures [9,10]; technical image processing/unmixing of 'component patterns' [11]; radiology/magnetic resonance imagery (MRI) [12]; LANDSAT MSS interpretations [13]. None of these works comprise a matching general system cum computer implementation, as does MIA. Ref. 14 looks at computer analysis and biomedical interpretation of microscopic imagery with a philosophy of image processing and the need for dedicated image analysis and interpretation very much along the lines laid out above. Ref. 15 addresses the issue of whether to use raw, i.e. unstandardised, or standardised (correlation-based) principal components, a point discussed by us in ref. 7.

ACKNOWLEDGEMENTS

We thank Nick Faust, Bruce Rado, Andrea Gernazian and Stan Quinn of ERDAS, for valuable discussions and help in Atlanta; Svante Wold, Chemometric Research Group, Umeå University, Sweden and Christer Albano, Peat Characterization project of Umeå, for interest and encouragement. PG would like to thank the latter organisation for financial and logistic support.

Two referees gave extremely different assessments of the original manuscript; both were very helpful. We thank both for illuminating the weak spots.

REFERENCES

1 P. Geladi, K. Esbensen and S. Wold, Image analysis and chemical information in images, *Analytica Chimica Acta*, 191 (1986) 473–480.
2 P. Geladi and K. Esbensen, Can image analysis provide information useful in chemistry?, *Journal of Chemometrics*, 3 (1989) 419–429.
3 K. Esbensen, S. Wold and P. Geladi, Relationships between higher-order data array configuration and problem formulations in multivariate data analysis, *Journal of Chemometrics*, 3 (1988) 33–48.
4 P. Geladi, Analysis of multi-way (multi-mode) data, *Chemometrics and Intelligent Laboratory Systems*, 7 (1989) 11–30.
5 The Beatles, *White Double Album*, 1967, EMI records. The track: 'Number Nine' does an admirable job of illustrating the tedious working part of image processing — John & Paul never knew!
6 K. Esbensen, M. Robb, G.-H. Strand and P. Geladi, Integrated image processing/GIS software/hardware for the PC. The ERDAS hw/sw system (software review), *Chemometrics and Intelligent Laboratory Systems*, 4 (1989) 95–98.
7 P. Geladi, H. Isaksson, L. Lindqvist, S. Wold and K. Esbensen, Principal component analysis on multivariate images, *Chemometrics and Intelligent Laboratory Systems*, 5 (1989) 209–220.
8 S. Wold, K. Esbensen and P. Geladi, Principal components analysis — a tutorial, *Chemometrics and Intelligent Laboratory Systems*, 2 (1987) 37–52.
9 G. Harauz, M. Stoeffler-Meilicke and M. van Heel, Characteristic views of prokaryotic 50S ribosomal subunits, *Journal of Molecular Evolution*, 26 (1987) 247–357.
10 J. Frank, The role of multivariate image analysis in solving the architecture of the Limulus Polyphemus hemocyanin molecule, *Ultramicroscopy*, 13 (1984) 153–164.
11 S. Kawata, K. Sasaki and S. Minami, Component analysis of spatial and spectral patterns in multispectral images. I. Basis, *Journal of the Optical Society of America, series A*, 4 (1987) 2101–2106.
12 H. Grahn, F. Delaglio, M.A. Delsuc and G.C. Levy, Multivariate data analysis for pattern recognition in two-dimensional NMR, *Journal of Magnetic Resonance*, 77 (1988) 294–307.
13 S. Folving, Studies of tidal flat environments with LANDSAT MSS data. ESA/EARSeL Symposium on Europe from Space. ESA SP-258 (Dec. 1986) 289–293, *NTIS Report*, N87-28157/XAB.
14 P.H. Bartels and G.L. Wied, Computer analysis and biomedical interpretation of microscopic images: current problems and future directions, *Proceedings IEEE*, 65 (Feb. 1977) 252–261.
15 A.R. Harrison and T.S. Richards, Principal components analysis of LANDSAT MSS data, *Journal of the British Interplanetary Society*, 39 (1986) 561–566.