Chemometrics and Intelligent Laboratory Systems, 14 (1992) 357–374
Elsevier Science Publishers B.V., Amsterdam

Strategies for multivariate image regression

Kim H. Esbensen *

Center for Industrial Research (SI), P.O. Box 124 Blindern, N-0314, Oslo, and E&P Research Center, Norsk Hydro, Bergen, P.O. Box 4313 Nygårdstangen, N-5028 Bergen (Norway)

Paul L. Geladi

Research Group for Chemometrics, University of Umeå, S-90187 Umeå, and Center for Peat Research, P.O. Box 5097, S-90005 Umeå (Sweden)

Hans F. Grahn

Astra Arcus AB, S-151 85 Södertälje (Sweden)

(Received 2 September 1991; accepted 29 November 1991)

Abstract

Esbensen, K.H., Geladi, P.L. and Grahn, H.F., 1992. Strategies for multivariate image regression. *Chemometrics and Intelligent Laboratory Systems*, 14: 357–374.

We present multivariate image regression (MIR) as a set of typically problem-dependent strategies for image decomposition guided by the nature of the Y variable and/or training data set delineation in the (X, Y) image domains. Regression techniques common in chemometrics may be applied also to the image regimen (in this paper we treat mainly two-dimensional images). We present applications of both IMPCR and IMPLS-DISCRIM in an effort to delineate the various possibilities for image regression. IMPCR builds directly on our earlier bilinear multivariate image analysis projection approach, while IMPLS-DISCRIM is trained on scene space binary classification masking with subsequent off-screen partial least squares analysis; the results are back-projected as images in the original scene space. Regression may either be carried out for modelling purposes and/or for subsequent prediction purposes. In the image domain this duality is accompanied by several optional training data set delineations in the scene space and/or in the spectral domain. We try to cover as complete a survey as possible of typical, representative regression problem types. We illustrate some of these MIR strategies with an MR-imaging example as well as a simple didactic MIR calibration from analytical chemistry.

INTRODUCTION

We have earlier presented a concept for multivariate image analysis (MIA) which is based on bilinear projection [principal component analysis (PCA)] of multivariate images in an otherwise mostly univariate image processing environment [1,2]. While the principal component transform is a well-known tool of most image processing systems, it is mostly used as a visualization aid: the PC-transformed imagery may be presented as principal component transformed images, either as an alternative to the unprocessed images or as a transformed basis for, e.g., classification training. The sequence: first training class delineation in the image space, followed by classifier calcula-

tions/segmentation etc. characterizes what we have termed the 'traditional image analytical approach' [1]. This approach is not optimal; we have argued that the MIA approach comprises a more flexible, problem-dependent approach [1,3–7].

In the traditionalist approach it is customary, for example, to plot the individual wavelength (or radiometric, or channel) intensities against each other, or even the calculated PC scores, but apparently most often as a post-classification/segmentation graphic illustration of the already established classes and their relationships. The MIA approach is the very opposite: the starting point of most image analytical efforts is in the same PC score cross plots, only exploratory data analysis and/or training class delineation is carried out in this data space, irrespective of (or rather: temporarily disjoint from) the (con)textural image relationships in the scene space [1–3]. This approach, among other advantages, leads to superior sequential image segmentation, complete with pixel outlier detection and quantifications, optimal class definitions and a built-in guarantee against ill-defined training data sets in classification [1–5,7].

There are — at least — two main types of multivariate image analysis, viz. exploratory data analysis (EDA), which deals with just one block of variables: **X**, and regression modelling dealing with two (or more) blocks of variables: **X** → **Y**; we have earlier outlined the MIA approach both for EDA as well as classification/discrimination in ref. 1, while ref. 3 treats in full detail the relationships between problem definition and proper data analytical methodology, including the regression case. In the following we presume that a proper regression problem is formulated, i.e. that a representative **X**- and **Y**-block training data set has been set up for modelling, usually with some subsequent prediction purpose: **X**(new) → **Y**(predict). While it is practically possible to use the regression modelling features alone, ostensibly for EDA, this is in general considered bad form in our point of view. If one does not know which variable(s) to designate as **Y**, perhaps one is better off not using a regression model; there are plenty of one-block alternatives in this situation.

REGRESSION ON IMAGE DATA

In ref. 4 we presented image regression focussing on image principal component regression (IMPCR) and especially on the algebraic details. In the present paper we focus on the application potential of this methodology that can be put to use within chemistry — as well as outside, since regression is so widely used in the sciences — and augment the PCR method with PLS1-DISCRIM, i.e. a one **Y** variable partial least squares (PLS) application in which this **Y** variable is of a simple binary (0, 1) type. This type of **Y** variable has been termed a 'dummy variable' in the classification parlance, but this is a most unfortunate choice of terminology, as we shall see from the power of the IMPLS1-DISCRIM example below. For in-depth methodological aspects of IMPCR, the reader is referred to ref. 4; we shall present the various strategies for multivariate image regression (MIR) below, directly following the outlines and terminology presented in ref. 4.

REGRESSION CASES

Image regression comprises basically a regression relationship between a stack of **X** images (an **X** image with a **P**-dimensional feature vector associated with each pixel) and one or more congruent **Y** images (there is either one **Y** variable or, in the case of several **Y** images, there is a complementary Q-dimensional vector for the same set of pixels). In ordinary regression modelling there is 'only' the problem of selecting the proper, i.e. the most problem-relevant, training data set on which to set up the regression model. There is also the inherent ambiguity of either a nominal, ordinal, or ratio scale for the **Y** variable. In the image regression situation this is augmented by various cases as to the delineation of the **X** space (and **Y** space) training data set basis, for example a training area in the **X** space (traditional image space training class delineation), or in the complementary statistical score space (MIA training class delineation). Moreover, one will often have access to more **X** data than **Y** data in the scene

space, but the opposite case may also be encountered. A first attempt at systematization is the following:

1. **Y** contains only one variable
2. **Y** contains several variables
3. **Y** variable(s) are categorical (nominal, binary scale)
4. **Y** variable(s) are continuous (ratio scale)
5. **Y** data and **X** data cover the same scene span
6 **Y** data and **X** data do not cover identical scene spans
7. **X** and **Y** data represent two normally disjoint spectral regions.

It is readily acknowledged that real-world data probably most often exhibit several of these characteristics at the same time, e.g. either (1) or (2) with either (3) or (4) and either (5) or (6) — or more complex combinations.

In ref. 4 we treated two representative cases, both with one continuous **Y** variable on a ratio scale, differing only with respect to criteria (5) and (6) above, one of which was based on a significantly underrepresentative training data set. The size of a training area is usually limited by the **Y**-data coverage. There is always a limit as to how far you can go in terms of prediction extrapolations, based on small training bases.

While it would appear that the use of the regression modelling formalism for probing alternative **X**–**Y** designations (i.e. exploratory **X**–**Y** block definitions, cf. above) has not been too well reflected in the literature, there is a much more legitimate EDA possibility for application in the image regression case. If, for example, interesting **Y** data are available, albeit only for a limited scene area, it may still be well worthwhile to put up a MIR model in order to investigate the resulting **X** imagery alone (i.e. no subsequent **Y** prediction). It is important to realize that one has made good use of the available **Y** information in this **Y**-guided regression modelling. Usually, however, the primary reason for regression modelling is for prediction purposes — also in the image regimen. For instance, predicting those areas in the **Y** image that were not present for modelling purposes (criterion (6) above), dependent upon all other features of the **X** and the **Y** data being comparable for this prediction.

Observe that while such an undertaking might well be a type of fill-in in the scene space (a prediction interpolation in the case where the spatial disposition of the **Y** coverage effectively spans the entire scene space), this may at the same time happen to be a prediction extrapolation in terms of the statistical span of the training data set (prediction samples for which the **X** data are outside reasonable comparison with the training data). This situation relates to the complementary nature of the scene space and the statistical space:

The regression modelling in the MIR approach follows directly that of our earlier MIA approach, in that the individual pixels are temporarily subjected to the regression modelling, irrespective of their (con)textural relationships. This is tantamount to temporarily taking the regression modelling into the statistical realm in which the modelling takes place exactly as in the non-image regression situation. This means that the 'usual' potential problems as to outliers, heterogeneity in the internal training data sample dispositions, i.e. sub-groupings etc., still apply. It is thus easy to appreciate how samples (in reality pixels) may well be spanning the scene space nicely, and may still exhibit outlying characteristics in this familiar statistical sense. It is one of our strongest arguments in favour of the MIA approach, that such relationships are bound to show up in the MIA score cross-plot approach to the image analysis. It is mandatory to assess the pertinent data structures in both these complementary spaces, the scene space (the image space) as well as the statistical (score cross-plot space). This is true for both MIA [1,2,5,6,7] as well as for MIR [3,4,8].

MULTIVARIATE CALIBRATION

In the discussion on regression modelling reference should be made to the recent textbook on the slightly broader topic: *Multivariate Calibration* by Martens and Næs [9]. Although this treatise covers non-image cases of calibration (train-

ing modelling explicitly for subsequent prediction purposes), its dual thrust for the user as well as the statistically accomplished reader really relates pretty much all the necessary background in order to bring the regression/multivariate calibration concepts safely into the image regimen. There is here a wealth of well-documented material relating to nearly all possible pittfalls and ill-informed use of the powerful tool of regression modelling in the predictive multivariate calibration setting as well as a well-credited list of other important work in this area.

SOME REPRESENTATIVE IMAGE REGRESSION APPLICATION FIELDS

In refs. 5–8 we have outlined typical and representative application fields within chemistry, remote sensing, and MR tomography. Some typical examples of image regression may be characterized.

Multi-wavelength optical inspection and spectroscopy

The P-dimensional wavelength vector carries (chemical) signals, to be transformed into (chemical) information, e.g. absorption intensities to be calibrated as concentrations etc. The same situation applies also — in principle — to, e.g., macroscopy (anything that is so large that it can be observed without a microscope). See Table 1 for examples.

Nuclear magnetic resonance (NMR) imaging, MRI

Traditional radiological approaches typically use only very few pulse repetitions and echo times as variables, often totalling some three, or at best four, variables (the latter representing a chemical contrast-induced extra variable) — a very small multi-variable set indeed. It is, however, entirely possible to augment this set appreciably with only a minor effort, by varying the experimental MR parameters spin-relaxation time (TR) and spin-echo time (TE) [2,8,10], routinely obtaining up to nine or more variables, and thus being able to tap this potentially much more differentiating higher-order realm. The MRI example below, as well as the paper by Grahn and Sääf [11], goes into these aspects in more detail:

> By varying TE (spin-echo time) and TR (spin-relaxation time) in the MR pulse sequence, one can weigh the imagery towards the main physical parameters spin–lattice relaxation time (T_1 weighted imagery) and the spin–spin relaxation time (T_2 weighted imagery). Typically, three different weightings are performed in an MR experiment, undoubtedly related to a common well-spread impression of MR scientists, that these images contain most of the essential characteristics in biological tissue a.o. However, as seen in many of our studies, higher-order component images (4th or 5th components) preferentially on the basis of twice this number of variables, from both MIA as well as MIR imagery, do indeed often contain tissue information and thus also describe significant structure/texture. Use of the full multivariate image analytical approach allows for a designed MR experiment where optimal tissue characteristics can be extracted [2,8,10].

See the remarks below on contextural variables (dealing with remotely sensed data) which probably may come to play an increasingly important role in furthering the MIA/MIR approach in NMR imaging as well; cf. also Grahn and Sääf [11].

TABLE 1

Variable types from typical image-producing methods

Variable type	Used in
Visible/UV electromagnetic radiation	Light microscopy Fluorescence Remote sensing
Infrared electromagnetic radiation	Microscopy Macroscopy Remote sensing
X-ray electromagnetic radiation	X-ray tomography X-ray microscopy Electron microscopy Micro-PIXE
Radio wavelengths	Radar remote sensing
Electron energy	Electron microscopy
m/e	Secondary-ion mass spectrometry
Pulse sequence/time intervals ± chemical contrasts	NMR imaging (MRI)
Ultrasound frequency	Ultrasound imaging

Remote sensing

(Readers not interested in the remote sensing regimen need only skim through the outline below for a guided tour of the traditional way of doing image processing.)

Typically, the remote sensing setting pays tribute to a strong paradigm where classification training almost invariably takes place in the scene space, influenced to a great extent by the nature of the remote imaging situation. This, almost perforce, must have a significantly smaller span in the Y-data coverage than in the **X** space. The **Y** variable in this setting typically is some kind of 'ground truth' observed/collected/measured on the ground or by a low-flying, airborn platform by, e.g. geologists, geophysicists, agricultural experts, foresters, environmental chemists, regarding for example the occurrence/prevalence of specific mineralizations, soil types, plant or forest coverage and plant diseases. The potentially useful **Y** calibration parameters are almost legion — but their practical implementation and wide usage has seen unforeseen delays up to this time. The remote sensing 'boom' has been eagerly awaited all through the eighties, but has so far failed to materialize. Now, however, there are several indicators that tend to show that this situation might well come around at last in the last decade of the millennium. In this setting, we believe that the MIA/MIR methodologies are destined to become very useful. We have used remote sensing examples extensively in our earlier work [1,4,7] and will not reiterate them in the present work, with the exception of one particular matter.

There is one aspect of augmented MIA/MIR application in this area which will be of great interest in other application areas as well: that of contextural variables. In refs. 12 and 13 the application of a whole suit of ten potentially useful contextural parameters are tested on complex SAR (synthethis aperture radar) data from an active ice field in the Davis Strait. The particular combination of the best of these contexture characterizing parameters and MIA proved to be a powerful tool for unsupervised pattern cognition, i.e. the situation in which the image analyst (glaciologist in this particular case) did not have access to ground truth (for obvious reasons: even getting to the ice-packed Davis Strait may be fraught with difficulty!); cf. also MRI, above.

In ref. 14 we presented a powerful demonstration of the capabilities of MIA when working with very many variables (channels) without any modifications to the software or the firmware necessary for a set of up to 256 variables. This feature, which has been signalled as one of the hallmarks of the MIA approach ever since its conceptual inception (vide refs. 1, 3, 6 and 7 and earlier references therein), can be taken over to MIR directly. There are plans — and some of these airborne platforms have already been test flown — for truly 'hyper-spectral' ('very many' channel) capabilities [15] of up to 1054 channels, even though the call for massive parallel supercomputer power in this context appears to us almost ludicrous in its limited, brute force approach. MIA/MIR would appear to constitute a viable everyman's alternative with much more modest computer requirements [1,3,7,14].

Disjoint variable blocks

Consider the situation in which, for example, we have identical scene space coverage for two, normally disjoint radiometric wavelength regions, be this either for a full or a fractional scene image. As a case in point we shall use an astronomical example in which we have access to both the 'conventional' visible region and the newly developed infrared (and NIR) regions for the same segment of the sky. The astronomers now have simultaneous access to both these regions, which traditionally is interpreted in a more or less disjoint fashion. The (N)IR region of the electromagnetic spectrum is considered to reveal hidden, embedded thermal sources, i.e. stars in the very process of being born, a.o. [16].

Considered from a chemometric point of view, this situation is simply an identical, congruent image, for which we have two disjoint wavelength regions, which could be designated as the **X** and a **Y** block, respectively. By rights, the visible region can probably be considered the most familiar (at

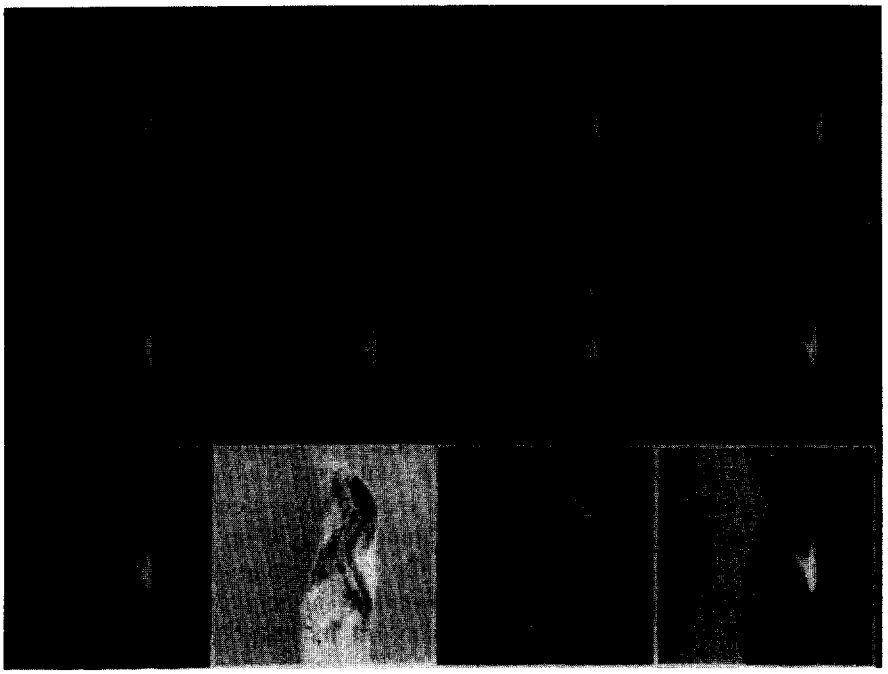

Fig. 1. Rat study. Selected set of eight MR parameters (top two rows), i.e. eight MIR X variables. Third row: the first four PCA images (left to right).

least for the astronomers) and perhaps one might favour an X block designation for this reason, with the new (N)IR region as the complementary Y block. One can then easily perform a MIR, etc. However, the snag here lies in what may be termed the most adequate problem formulation [3]. Consider in a little more detail exactly what goes on in the PLS analysis. The X information is gauged by the Y information; the Y information guides the decomposition of the X matrix data structure. In this context it would appear more relevant to put the 'common, well-known' visible region into the Y block, thereby allowing this familiar aspect to act as the guiding hand for the decomposition of the new — and potentially interesting — X block, i.e. the exact opposite of the first designation. In ref. 3 we have discussed in detail the essential need for this type of proper problem formulation.

By allowing PLS to decompose the (NIR)-X data in this fashion, we are now in a position to allow the subsequent back-projection into the image (scene) realm to let the X score images show us the (N)IR image in a new fashion, viz. the (N)IR image that correlates optimally with the known visible image (in the Y space) in the PLS regression sense. Displaying this type of PLS-decomposed X image allows the astronomer to take cognizance of both these two, previously disjoint wavelength regions simultaneously, not by some simple unreflected image processing addition/subtraction or the like, but in a very specific, narrow problem-dependent intercorrelated fashion. But there is more to be gained from this judicious use of problem-dependent PLS block designation. Perhaps the greatest potential possibility for calculating truly new derived imagery lies in inspection of the residual X image from the above analysis (the residual image is covered in detail in refs. 1, 3, 4 and 7).

In the residual X image we will find every item of radiometric information that does not correlate with the known visible information present in the Y image — in its pertinent contextural setting in the X scene space. This residual image may be termed the problem-dependent decomposed im-

age, exactly because of this specific, problem-dependent PLS block designation. It would appear that the residual imagery will constitute a valuable augmentation to the raw, unprocessed (N)IR images. For example, the visible-decomposed (N)IR image may now be further processed. It may, for example, be subtracted from the raw image, in a much more intelligent fashion, a.o.

In chemistry in general, and in analytical chemistry in particular, where multivariate image analysis is required, there are ample analogous situations with the above with respect to disjoint spectral regions; see the tabulation above. The only requirement that need be complied with is the identical spanning of the scene space — and, of course, that the problem formulation is indeed contingent upon a similar set of disjoint, or otherwise not actively combined, spectral regions.

CHEMICAL (MULTIVARIATE) IMAGE ANALYSIS (CIA)

The above four typical application fields for multivariate image analysis serve to illustrate the emerging concept of chemical image analysis (CIA). CIA is primarily characterized by the interdependent spectral and image (area/volume) analyzing capabilities. It combines all the well-known analytical attributes of multivariate calibration from chemistry/chemometrics with the powerful facilities of multivariate image analysis — MIA/MIR. As in non-image chemometric applications, as well as in the image regimen, one is most often in a situation where both EDA as well as classification/segmentation/calibration and prediction are on the agenda. There is a truly bewildering application potential for CIA and we have only tried to cover the first few, most obvi-

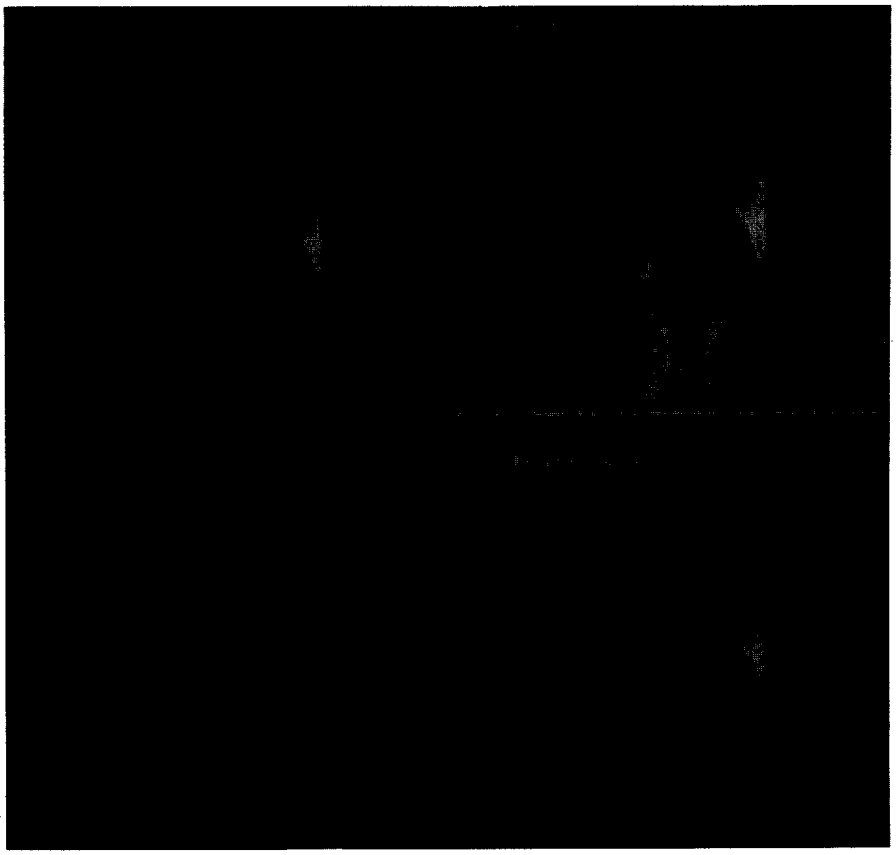

Fig. 2. First four PCA images of rat study in full detail (upper left → right; lower left → right); loading relations not shown here.

ous possibilities in this paper. Future applied work will be especially important in finding new areas in need of further methodological development; but we suspect that the present status of operative MIA [1,2,5,6,7], and soon MIR as well [3,4,8], will span a more than wide enough space for work for many years to come.

EXAMPLES

Below two illustrative expositions of the salient aspects of MIR are outlined, one using three 'dummy variables' in the **Y** block, the other using

TABLE 2

Experimental MR1 variables

Variable	Experimental MR parameters	
	TR	TE
1	500	21
2	500	30
3	500	45
4	800	21
5	800	30
6	800	45
7	1200	21
8	1200	30
9	1200	45
10	500	60
11	800	60
12	1200	60

Fig. 3. PLS1-DISCRIM prediction images for rat study. (Clockwise): PLS1-DISCRIM prediction image: muscle; PLS1-DISCRIM prediction image: brain; (ignore: 3rd PCA image); PLS1-DISCRIM prediction image: fat. Observe the satisfactory histological contrasts.

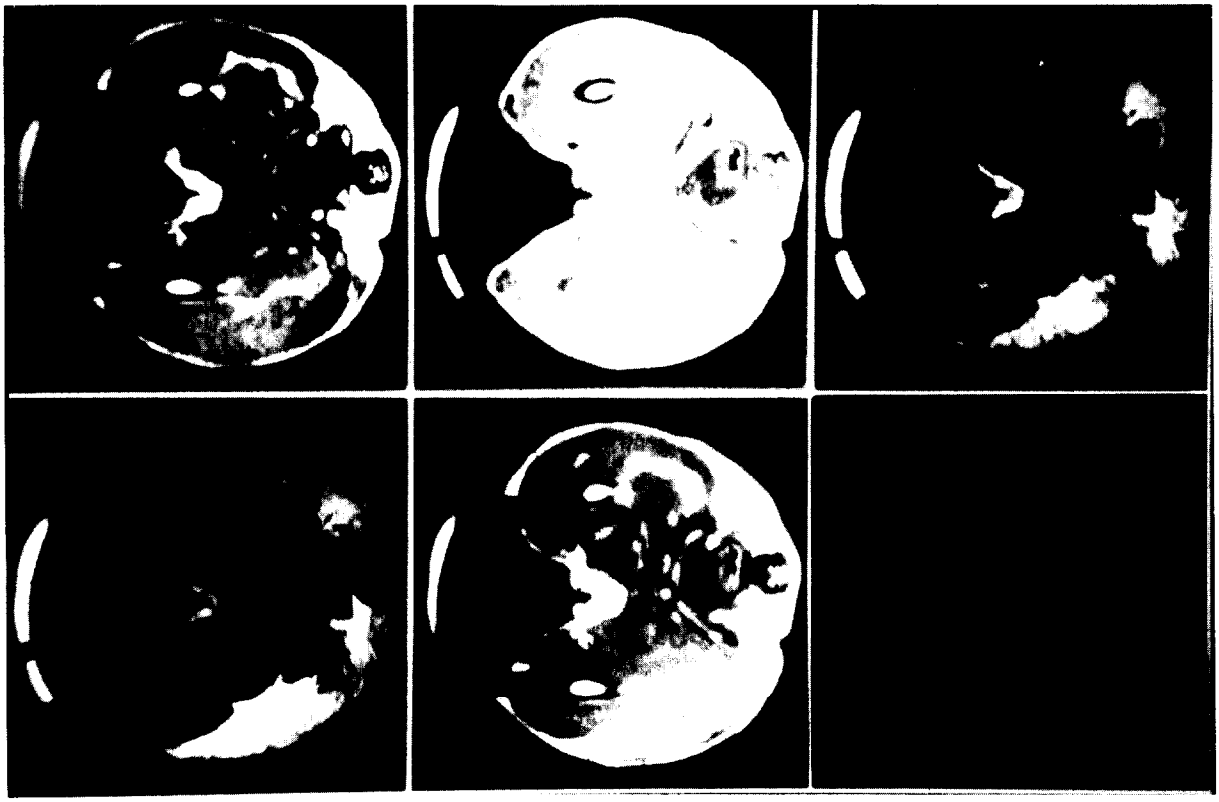

Fig. 4. Beagly study. Set of five selected MR parameters, i.e. five MIR X variables. Observe how none of these original images may serve to delineate the different features (muscle, oedema, abscess etc.) alone.

one continuous Y variable, which corresponds to concentrations of selected chemical compounds. The first study, relating to MRI, which is actually composed of two separate MIR in vivo studies but with roughly identical MR experimental parameters, is intended to illustrate the power of MIA/MIR when used in tomography. The second study is a very simple, didactic exposition of some of the more apparent possibilities in image-based analytical chemistry when the matrix characteristics are of a very difficult nature, i.e. where the physical makeup of the analyte(s) play a major role in shaping the spectral signatures besides the quantitative analyte concentration contributions.

MR imaging (MRI)

The magnetic resonance images were recorded on a General Electric chemical shift/imaging (CSI) instrument. The images had a field of view of 80 × 80 mm for the first study (termed the 'rat' study below) and 120 × 120 mm for the second (the 'beagle' study). The images were obtained with a spin echo sequence with pulse repetition times (TR) and echo times (TE) according to Table 2. The slice thickness was 2 mm, using a single echo. Digital resolution is 256 × 256 pixels. (Readers not well-versed in MRI need only concern themselves that the tabulation in Table 2 outlines twelve X variables, all eager to please.)

Rat study

In the rat study, eight of the twelve original images in Table 2 showing the whole body of the rat are given in Fig. 1, with Fig. 2 depicting the corresponding score imagery. A subset of eight X variables was chosen from the initial MIA analysis, in which a plot of the loading relationships showed which variables were most influential in

modelling the score space data structure — a standard MIA procedure [1,2,8].

The rat study employed a set of three 'dummy' Y variables — binary classification variables, in which supervised class belonging in the PLS training stage was denoted '1' (and '0' for 'not belonging to') by visual inspection of the scene space image by a trained MR analyst/pathologist/veterinarian. These three classifications were chosen so as to cover a total of 15 × 15 pixel each. The variables correspond with the rat's tongue, liver and brain, respectively. These training pixels were chosen so as to represent the pertinent training feature as coherently as possible (i.e. in the middle of areas of 'known' organs). Observe the 'traditionalist' approach here, i.e. training set delineation in the image space; in other, more complex, MIR studies, we routinely also use the MIA score space approach of course.

Fig. 3 shows the results of the ensuing prediction of the entire image. The three training areas (3 × 15 = 45 pixels) correspond to less than 0.06% of the entire image field of view (FOV). Allowing for the effective fraction of this FOV that is actually covered by the rat (roughly a third), the training areas still comprise a truly unsignificant part — certainly much less than 1% by area. Consequently we have not bothered removing these areas in the prediction run. Even though this is strictly speaking not correct as a quantitative prediction validation, this small fraction of 'doubly modelled-predicted' pixels does not constitute any serious infringement at all of the visual impression in Fig. 3 that a clear and concise,

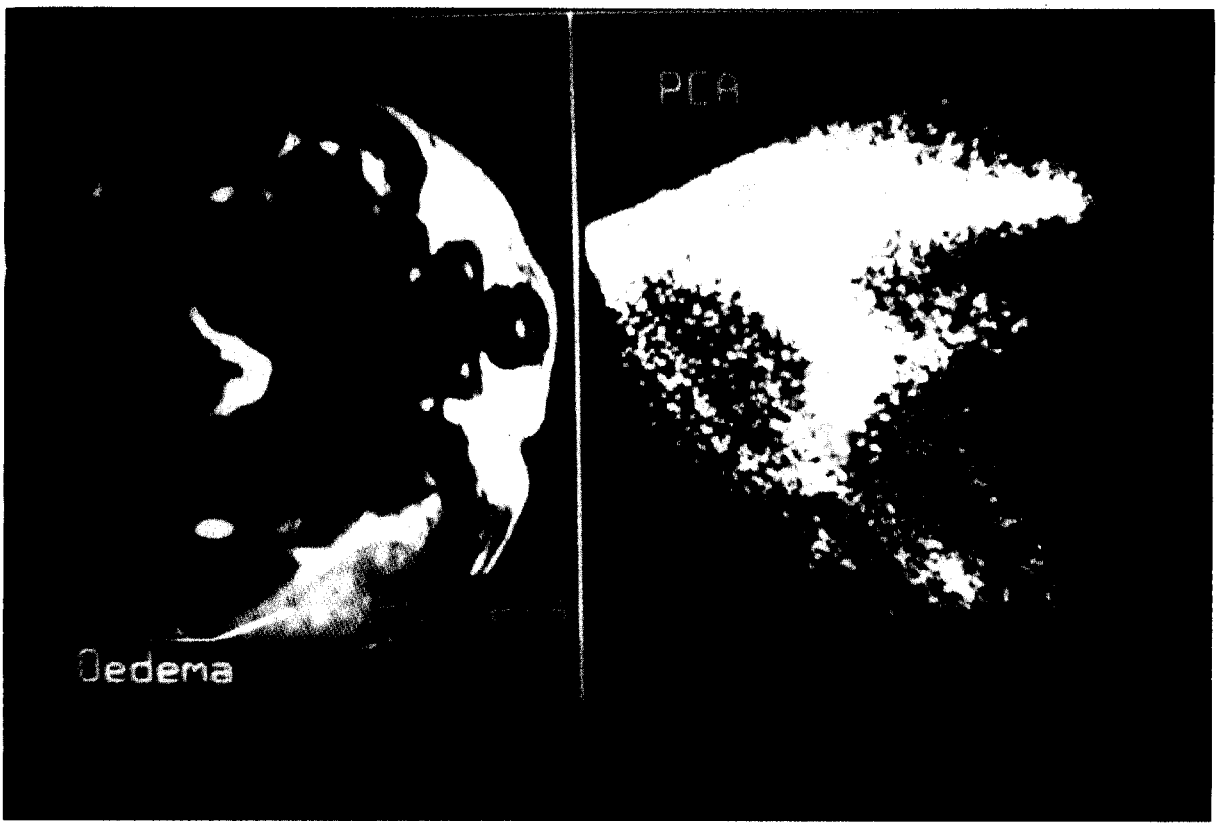

Fig. 5. Beagle study PCA score plot, components 1 and 2 (right half), with tissue designations in left half (score image). Abscess class shown in red; oedema class in yellow.

radiologically useful classification of these three tissue types has indeed been obtained.

Beagle study

The second study was performed on an anaesthetized dog — a beagle — where a suspension of live bacteria was injected in one of the dog's thighs in order to induce infection. This region developed an abcess (infected tissue), and surrounding this a region of water-filled tissue was also found (oedema). These two tissue types constituted the primary basis for a PLS training stage, with a third **Y**-discriminating variable corresponding to unaffected muscle tissue in this case.

This PLS1-DISCRIM study (identical in its methodology to the rat study above) used five congruent images, i.e. five **X** variables. Figs. 4–9 show the results from this study. A detailed appreciation of the figure captions in this sequence will reveal some of the salient aspects of the typical MIA/MIR modus operandi, from original image assessments (Fig. 4), through initial problem-dependent projections — PCA (Fig. 5) and PLS (Fig. 6), further from more refined training sample delineations (Fig. 7) to the final prediction imagery (Figs. 8 and 9).

The results compiled in Figs. 8 and 9 show the distinct classification of these three tissue types in the derived imagery. It is possible to fully resolve the abcess/oedema complex, cf. left and right images in Figs. 8 and 9; witness also the muscle prediction image, in which these two types are entirely absent (Fig. 9, left half). The MRI potential of this, the most simple binary PLS1-DISCRIM possible, would appear to be great. We infer that a continuous **Y** variable PLS1 facility should augment this appreciably in the near future.

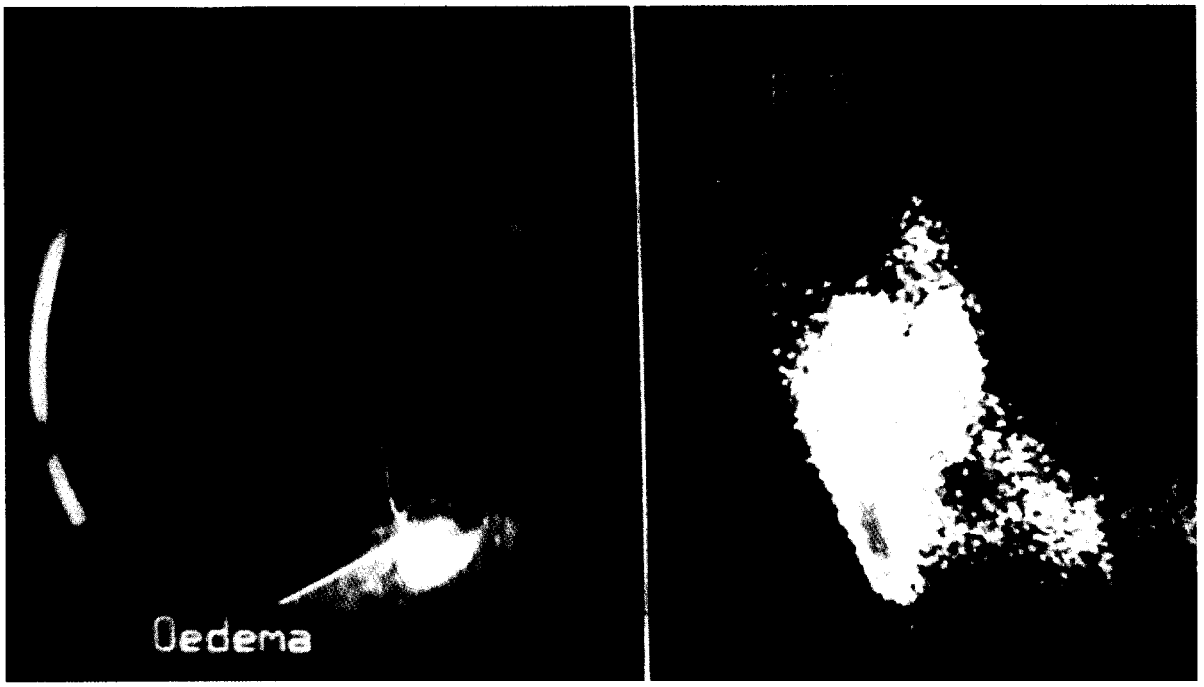

Fig. 6. Beagle study PLS training class delineation (abscess: red) in left half (PLS component image), with corresponding PLS score 1 vs. score 2 plot in right half. Observe how training selection may be initiated in either of these two domains.

The MRI calculations were performed using UNIX software written in C for use on X11 window systems.

Analytical chemistry — a didactic visible (VIS) / near infrared (NIR) example

The NIR region (800–3000 nm) has been used very succesfully in combination with multivariate techniques for data analysis, for example in the classification and calibration of organic and biological materials, wood and pulp, peat and other agricultural media. The current measurements are made in the 580–1200 nm region with the variable selection shown in Table 3; a total of eight variables was chosen for this application. The light was filtered at the source, preventing unnecessary optical path adjustments (refocussing).

The experiment below uses a NIR camera as the input medium for MIR with the objective of quantitative analysis of powder mixtures. The present example uses model powder mixtures of sugar and common salt (sodium chloride) in various mixtures, some of which ground to different average grain size. Thus the experiment relates both to the chemical analytical signal itself as well as to physical interference/matrix effects. The present example is only meant as a very simple didactic illustration of the application potential of this type of CIA; thus there is no room here for full grain size modelling. The full study will appear elsewhere.

The powder mixtures are made up of four identically flattened piles' laid out on a black velvet background cloth, a material that absorbs light well in the VIS and NIR regions and which gives minimal specular reflection. The mixture

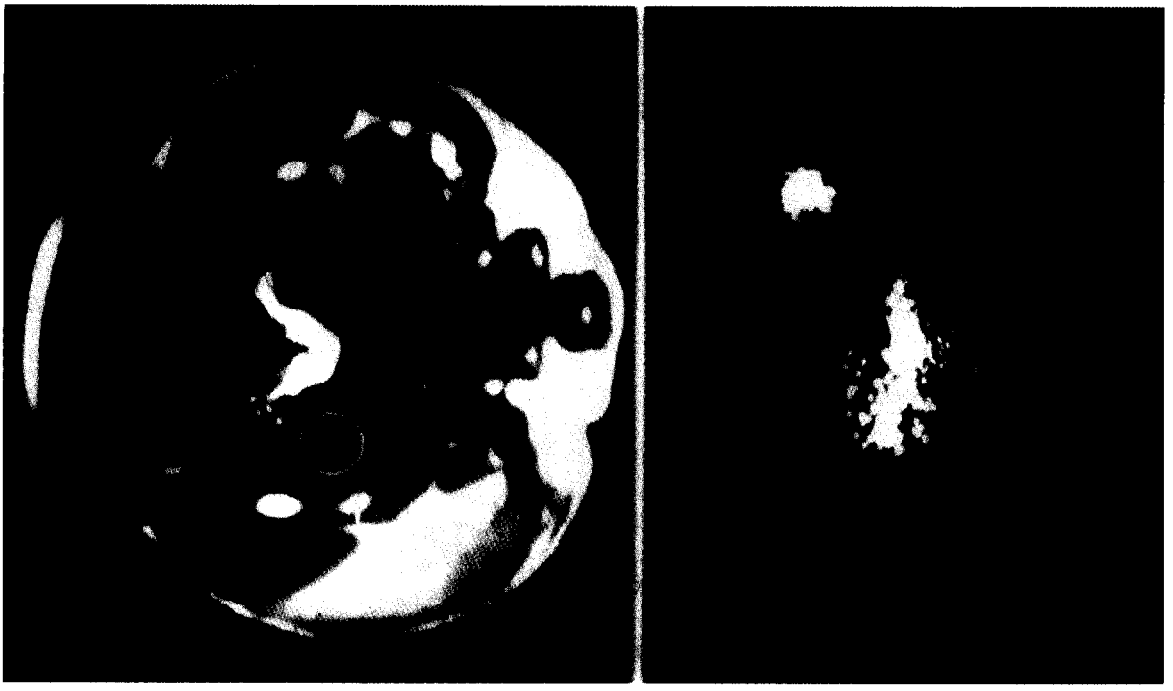

Fig. 7. Beagle study training class delineations for abscess (white), oedema (blue) and unaffected muscle (red) in righthand score space (components 1 vs. 2). Muscle (red) alone has been highlighted in left image.

compounds are sodium chloride (kitchen grade) and sugar (same source); there is also a 50/50 (w/w) mixture of identical grain size as the two pure end members. Finally, another 50/50 (w/w) mixture was ground to a significantly finer grain size. See the caption to Fig. 10 for an explanation of the standard positioning of these four mixtures in Figs. 10–15.

The training image consisted of the central 128×128 portion of the entire 256×256 image scene, see Fig. 10A. For this study we employed an IMPCR model identical to that presented in ref. 4, which was calibrated for sugar concentration only, cf. above. This model was mean-centered, but no rescaling was deemed necessary [3,7]. The MIR model used in the results presented below comprise five PCA components with the following **X** fraction sums-of-squares (SS) modelled, of which we used three and four in the predictions (Table 4).

Fig. 11 is a colour composite of the three first principal components. Observe that there would appear to be systematic information related both to sugar concentration as well as to grain size (upper left quadrant versus all other three), even though it is not calibrated for here. It would appear that a more detailed training calibration simultaneously may address both concentration as well as grain size quantifications.

Fig. 12 shows that rank = 3 and rank = 4 IMPCR models by and large compete favourably with one another; in this plot pixel clustering along the diagonal reflects good agreement between the two alternate rank models, while clustering away from the diagonal will indicate (severe) differences. We experimented with both

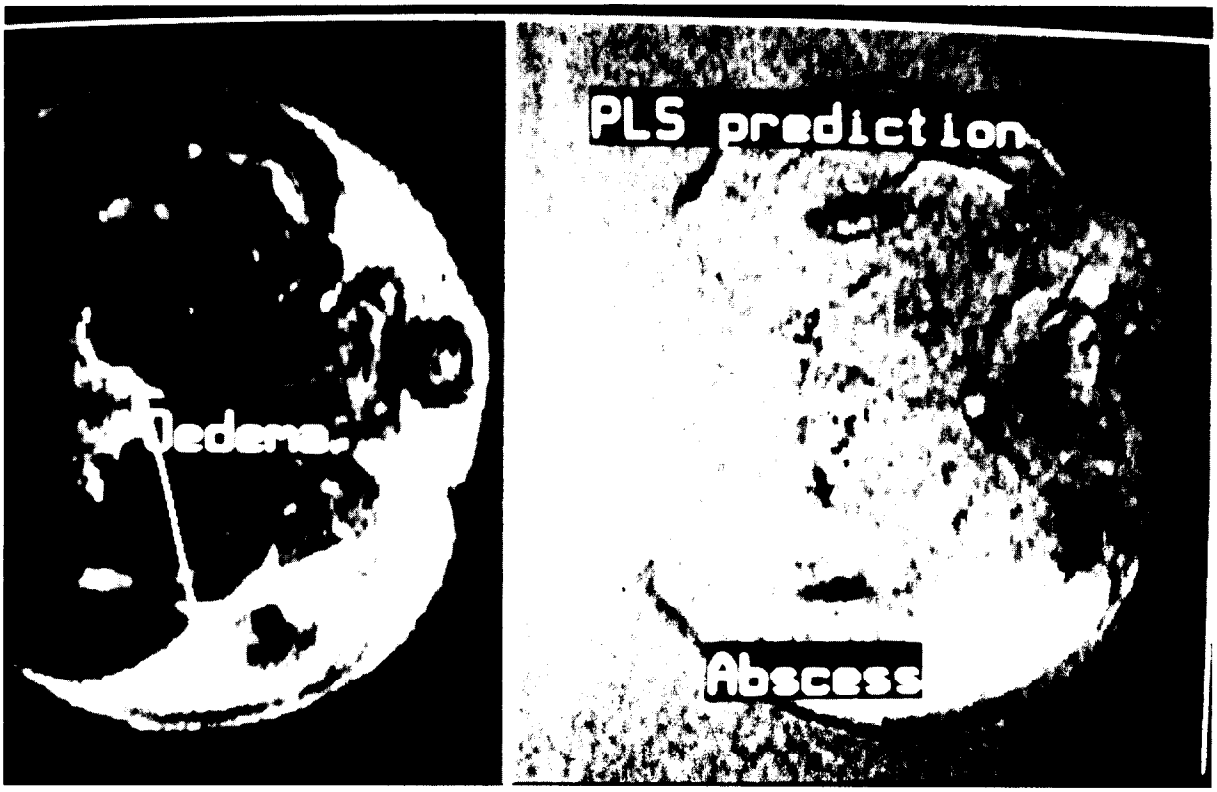

Fig. 8. Beagle study PLS1-DISCRIM prediction images for oedema and abscess. Note the excellent histological distinctions due to PLS combinations of the original images, cf. Fig. 4.

TABLE 3

VIS/NIR variable selection

Variable No.	Wavelength (nm)
1	1210
2	1010
3	900
4	800
5	740
6	680
7	630
8	580

TABLE 4

X variable modelled by PCA

PCA component	X fraction SS modelled (%)
1	97.8
2	1.1
3	0.8
4	0.1
5	0.1

models (rank = 3, 4) in order to see if any overfitting effects could be produced — with somewhat ambiguous results, however, compare Figs. 14 and 15.

Fig. 13 gives an overview of the entire suite of rank = 1, 2, 3, 4 models, with the rank = 4 model in full detail in Fig. 14. Observed in the proper progression (upper left → right; lower left → right) one appreciates that scene space fill-in predictions are apparently quite satisfactory (inner rectangle), with rapid deterioration in the marginal, less than optimal areas. By comparing the upper and lower left quadrangles in Fig. 14 it will become clear that the CIA signal information necessary to deal with the 'embedded' grain size differences is apparently present and thus available for incorporation in more complex models.

Fig. 9. Beagle study PLS1-DISCRIM prediction image for muscle (left half). The abscess prediction image was also shown in Fig. 8 (here in a slightly different contrast-stretched version, which performs better than Fig. 8).

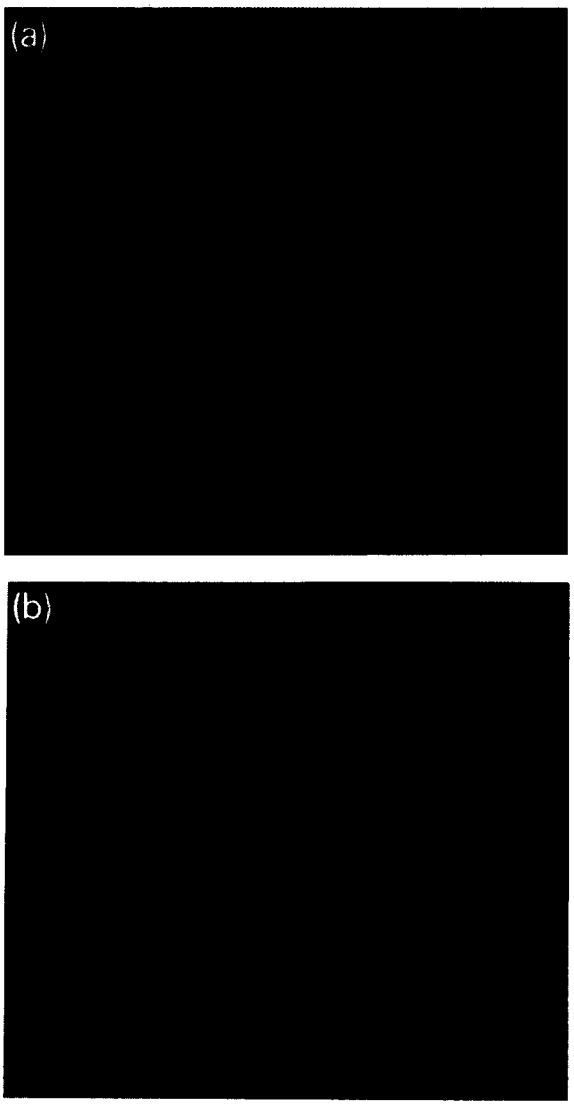

Fig. 10. Mixture CIA setup in false color composite (in brackets: lookup table colors): upper right: 0% sugar (i.e. 100% salt) (black); upper left: 50/50 sugar/salt (fine ground) (orange); lower left: 50/50 sugar/salt (orange); lower right: 100% sugar (green). The inner (128×128) square in (a) is the training domain. The particular **Y** variable setup in (b) uses only sugar concentrations as the PCR **Y** variable, i.e. no grain size information; see text. Note some uneven light conditions in the marginal areas.

Finally, Fig. 15 shows an alternative type of graphic prediction evaluation that we presented in ref. 4; white areas denote the best match(es) between predicted and known 'true' values. There

Fig. 11. A composite of the first three principal components for the training data set. PC1 = red; PC2 = green; PC3 = blue. Note that grain size-related information does enter into these components – cf. upper left vs. all other quadrangles.

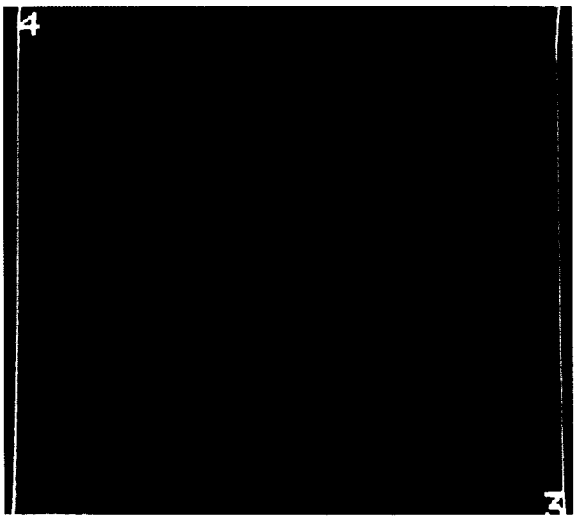

Fig. 12. Scatter plot of two IMPCR models, one with rank = 3 (horizontal) and the other with rank = 4 (vertical). Close inspection reveals a tendency for three (four) pixel groupings, though the first-order pattern is of a quite satisfactory (1, 1) overall correlation between these two models. The four groups correspond to the four **Y** classification classes. See text for details.

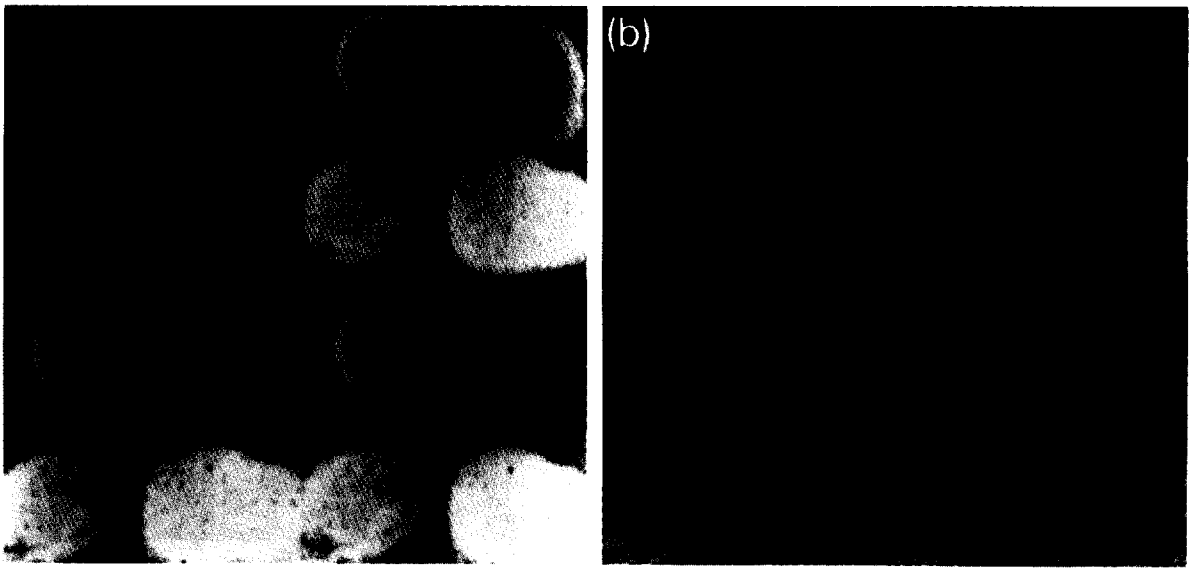

Fig. 13. IMPCR prediction images for all rank = 1, 2, 3, 4 models (upper left, right; lower left, right, respectively). Lookup table colors: 0% sugar (black); 50% sugar (orange); 100% sugar (green). The development of more satisfactory models can be graphically appreciated.

are still plenty of subtle differences between all four mixtures after the present first-order IMPCR modelling that makes for an optimistic view on the possibilities for more refined MIR CIA.

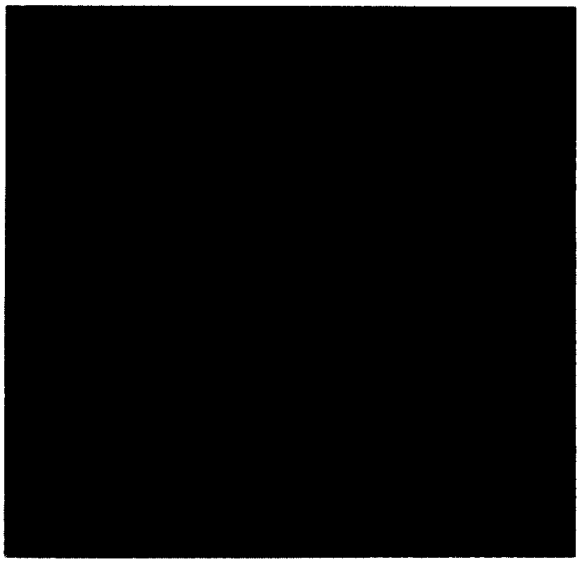

Fig. 14. IMPCR (rank = 4) Y predictions of sugar concentrations. Predictions are satisfactory inside training area delineation; note rapid deteriorations in the marginal areas with imperfect light conditions and/or different grain size.

This experiment was digitized using Kontron IBAS hardware controlled by IBAS 2.0 software on a 386/387 PC host. The MIA software was standard ERDAS modules, augmented using the ERDAS toolkit, also running on a 386/387 PC host. Pictures were photographed directly from the image monitor on 200 ISO colour negative film.

IMAGE REGRESSION — A NEW REALM

The above types of MIR CIA are only the first forays into a whole new realm of image-based analytical chemistry (area/volume-based), the development of which has only just begun. The question whether a continuous Y-variable IMPLS1 version can do a better job than IMPCR1 is particularly interesting. In the image regimen there is nearly always a potentially very big training set present; there may — or may not — be sufficiently strong data structures present in the image domain in particular cases, that the well-known PLS/PCR ambiguities [9] are reduced. We will not venture to pass judgement on this issue before a considerable base of MIR experi-

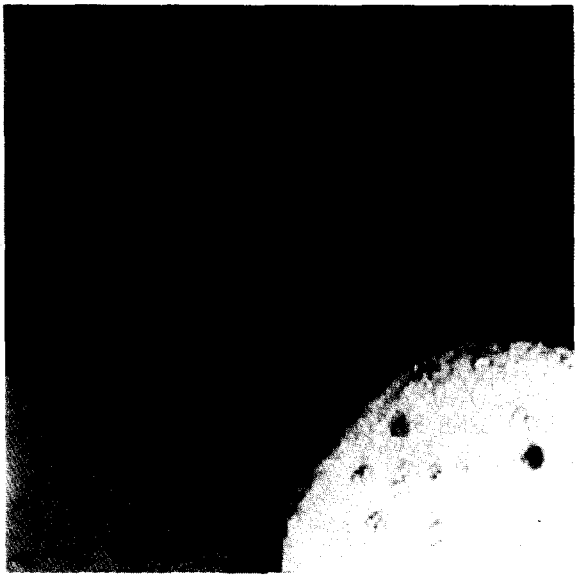

Fig. 15. Prediction evaluations using red/green/blue (RGB) color compositing. 'True' values in red channel; rank = 3 IMPCR predictions in green + blue = cyan channel. White areas (or weak red/cyan): good/satisfactory match between true and predicted values. Red areas: true values higher than predicted. Blue areas: prediction values higher than true. Black: background.

ence has been built up, preferably on many differing types of problems/data sets. An IMPLS1 version encompassing a full, continuous **Y** variable is currently under development; whether a full-spectrum IMPLS2 (more than one **Y** variable) version is also needed, has not been decided upon yet and will likewise be dependent upon forthcoming practical experiences.

FUTURE DEVELOPMENTS: HIGHER-ORDER IMAGE DECOMPOSITION

We have illustrated the emerging concept of chemical image analysis (CIA). CIA combines the analytical attributes of multivariate calibration from chemistry/chemometrics with the facilities for multivariate image analysis — MIA/MIR. There is a very large potential for CIA application ahead of us. Future applied work will also be especially important for finding still new areas, in need of even further methodologic development. As was laid out in ref. 3, we perceive the ultimate development objective as an extention of the present MIA → MIR to encompass more complex data arrays (higher than two-dimensional) in the direction → MAD (multi-way array decomposition), i.e. MIA/MIR in the CIA guise applied on the pertinent multi-way data array organizations in their proper problem-dependent contexts. It is important to realize that rank ⩾ 3 problems will have several alternative unfolding possibilities, all directly related to the original (regression) problem formulation [3,4].

ACKNOWLEDGEMENTS

We gratefully acknowledge help from Professor Nikolaus M. Szeverenyi, SUNY Health Science Center, Department of Radiology, Syracuse, New York in collecting the MR images used here. IMPCR calculations and other MIR work was done by us. Helén Bergner, Centre of Peat Research, Umeå, assisted in sample preparation. K.H.E. gratefully acknowledges support from NORSK HYDRO, Exploration and Production Research Centre, Bergen, during a 6-month resident statistician scholarship.

REFERENCES

1 K.H. Esbensen and P. Geladi, Strategy of multivariate image analysis (MIA), *Chemometrics and Intelligent Laboratory Systems*, 7 (1989) 67–86.
2 H. Grahn, H. Szeverenyi, M. Roggenbuck, F. Delaglio and P. Geladi, Data analysis of multivariate magnetic resonance images: a principal component analysis approach, *Chemometrics and Intelligent Laboratory Systems*, 5 (1989) 311–322.
3 K.H. Esbensen, S. Wold and P. Geladi, Relationships between higher-order data array configurations and problem formulations in multivariate data analysis, *Journal of Chemometrics*, 3 (1988) 33–48.
4 P. Geladi and K.H. Esbensen, Regression on multivariate images: principal component regression for modelling and prediction purposes, *Journal of Chemometrics*, 5 (1991) 97–111.
5 P. Geladi and K.H. Esbensen, Can image analysis provide information useful in chemistry?, *Journal of Chemometrics*, 3 (1989) 419–425.

6 P. Geladi, K.H. Esbensen and S.Wold, Soft modelling on multivariate images, *Technical Report SAND/23/1987*, Norwegian Computing Center, Oslo, 1987.
7 P. Geladi, H. Isaksson, L. Lindqvist, S. Wold and K.H. Esbensen, Principal component analysis on multivariate images, *Chemometrics and Intelligent Laboratory Systems*, 5 (1988) 209–220.
8 H. Grahn, N. Szeverenyi, M. Roggenbuck and P. Geladi, Tissue discrimination in magnetic resonance imaging: a predictive multivariate approach, *Chemometrics and Intelligent Laboratory Systems*, 7 (1989) 87–93.
9 H. Martens and T. Næs, *Multivariate Calibration*, Wiley, New York, 1989.
10 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.
11 H. Grahn and J. Sääf, Multivariate image regression and analysis — useful techniques for the evaluation of clinical magnetic resonance images, *Chemometrics and Intelligent Laboratory Systems*, 14 (1992) 391–396.
12 E. Holbæk-Hansen, A.H. Schistad and K.H. Esbensen, Semi-supervised classification of sea ice based on SAR images, *Proceedings IGARSS' 90 (10th International Geoscience Remote Sensing Symposium), University of Maryland, USA, May 20–24, 1990*, pp. 2245–2248.
13 E. Holbæk-Hansen and A.H. Schistad, Delvis styrt klassifikasjon av SAR-bilder av hav-is, *Norsk Regnesentral Report 835*, Norwegian Computing Center, Oslo, 1989.
14 N. Faust, W. Sharp, P. Geladi and K.H. Esbensen, Application of multivariate image analysis (MIA) to analysis of hyperspectral image data for mineral exploration, *Proceedings of the Eight Thematic Conference on Geologic Remote Sensing, Denver, Colorado, USA, April 29–May 2 1991*, pp. 215–228.
15 J.R. Schott, Remote sensing of the earth: a synoptic view, *Physics Today*, 42(9) (1989) 72–79.
16 F.C. Gillett and J.R. Houck, The decade of infrared astronomy, *Physics Today*, 44(4) (1991) 32–37.