

Automated Analysis of SEM X-Ray Spectral Images: A Powerful New Microanalysis Tool

Paul G. Kotula,* Michael R. Keenan, and Joseph R. Michael

*Materials Characterization Department, Sandia National Laboratories, P.O. Box 5800, MS 0886,
Albuquerque, NM 87185-0886, USA*

Abstract: Spectral imaging in the scanning electron microscope (SEM) equipped with an energy-dispersive X-ray (EDX) analyzer has the potential to be a powerful tool for chemical phase identification, but the large data sets have, in the past, proved too large to efficiently analyze. In the present work, we describe the application of a new automated, unbiased, multivariate statistical analysis technique to very large X-ray spectral image data sets. The method, based in part on principal components analysis, returns physically accurate (all positive) component spectra and images in a few minutes on a standard personal computer. The efficacy of the technique for microanalysis is illustrated by the analysis of complex multi-phase materials, particulates, a diffusion couple, and a single-pixel-detection problem.

Key words: spectrum image, spectral image, statistical analysis, principal components analysis, X-ray microanalysis, information extraction, multivariate statistical analysis

INTRODUCTION

It is a conundrum that faces microanalysts every day: How do you perform a complete and comprehensive survey of the chemistry of a microstructure in a reasonable time? Traditional solutions in the scanning electron microscope (SEM) would include taking an image and then collecting X-ray spectra from a series of points. This method relies heavily on the operator's ability to identify chemically distinct regions. In addition, for a complex microstructure, this could take a considerable amount of time to do correctly. On the positive side, if care is taken, these "point spectra" can be quantified. An alternative method that has been used widely is X-ray mapping (Goldstein et al., 1992), where a window about a range of X-ray energies is integrated and

displayed as an image. This has the advantage that qualitative elemental distributions from an area of a microstructure can be visualized. The disadvantages are that maps are generally not quantitative, can be susceptible to artifacts (Newbury, 1997; Newbury and Bright, 1999), rely on foreknowledge of the elements to map, and cannot discern elemental correlations (e.g., Al versus Al_2O_3). More recently, with increases in computer CPU speed, memory, and storage space, X-ray spectrum-imaging systems have become available. A spectral image is a two-dimensional array of points in the microstructure with a complete X-ray spectrum from each point (Legge and Hammond, 1979; Jeanguillaume and Colliex, 1989; Mott et al., 1995; Anderson, 1998). Spectral images potentially overcome the shortcomings of point analyses (e.g., how to pick the points) by being an imaging technique that allows the analyst to qualitatively analyze and perhaps quantify the spectra. It should be noted that for the spectral image to have as much spatially re-




Figure 1. Schematic of the spectral image data cube.

solved chemical information as possible, the pixel size should be somewhat finer than the X-ray generation volume (Lyman, 1986; Goldstein et al., 1992). In addition, spectral images have potential to overcome the problems with mapping (e.g., choosing the elements to map) because there is a complete spectrum from each pixel. The problems with spectrum imaging up until recently were that there were no tools for extracting the chemically relevant information from the massive amount of data (>65 million individual data points for a 256×256 -pixel \times 1024-channel spectral image). The tools that have been available for analyzing spectral images amount to ones that map after the fact and then allow summing of spectra that have been thresholded from maps (Mott et al., 1995; Mott and Friel, 1999). Additionally, it can take a considerable amount of time to collect spectral images with individual spectra containing sufficient counts to perform even a qualitative, let alone quantitative, analysis of a given spectrum. It is for these reasons that new analysis methods are needed to objectively analyze spectral images (Anderson, 1999, 2000; Kotula et al., 1999; Kotula and Keenan, 2000).

One inherent problem with spectrum imaging is that the data sets cannot readily be visualized in their entirety. This point can be appreciated by looking at the three-dimensional representation of the spectrum image data set in Figure 1. Each element of the cube has associated with it three dimensions: x -position, y -position, and X-ray energy channel. Further complicating the complete analysis is the fact that much of the spectrum image is redundant. This would include regions of the spectra where no peaks are found as well as pixels from the same chemical phase. Redundancy is, in fact, absolutely necessary for the analysis method described below to work. The single greatest prob-

lem is that for a robust unbiased analysis, no assumptions can be made as to what may or may not be present. In other words, you cannot blindly ignore parts of the spectra just in case there really is a peak there.

The key requirements of a robust system for the automated analysis of spectral image data sets are: no assumptions about the absence or presence of any constituent; the ability to handle noisy data; the ability to handle significant spectral overlap; spectral image processing times of the same order or much less than the acquisition time for the data; and quantitative agreement between raw and reconstructed data. The system developed in the present work meets all of the above requirements.

MATERIALS AND METHODS

Multivariate Statistical Analysis

Multivariate statistical analysis (MSA) describes a general set of techniques that are utilized for analyzing data sets such as opinion polls, new-drug trials, and series of spectra (Harman, 1976; Malinowski, 1991; Geladi and Grahn, 1996). For all of these applications of MSA there are from several to thousands of variables in the data sets, and the goal is to explore the variation in multiple dependent variables simultaneously. For X-ray spectral image data, the variables are the energy channels and will typically number from 1000 to 4000. The end product of MSA is a transformation from the raw data, with all its inherent complexity, to a simple solution that captures the important features of the original data in a more compact way. This is equivalent to performing a dimensional reduction on a data set where the original spectral image could contain 65,536 spectra (i.e., 256×256 pixels) when in reality there are only a few spectral shapes that describe the relevant information in the spectral image. This has the advantage that redundancy can be exploited by combining all the like spectra. In fact, redundancy is used to advantage in multivariate methods, so that a collection of similar noisy spectra can, with suitable algorithms, be combined to form a superspectrum with the composite counting statistics of the individuals.

The goal of MSA in the present context is to extract the relevant information from the raw X-ray spectrum image. A spectral image will typically contain from several thousand to tens of thousands of X-ray spectra, each with 1000 or more energy channels. The desired result from MSA of a

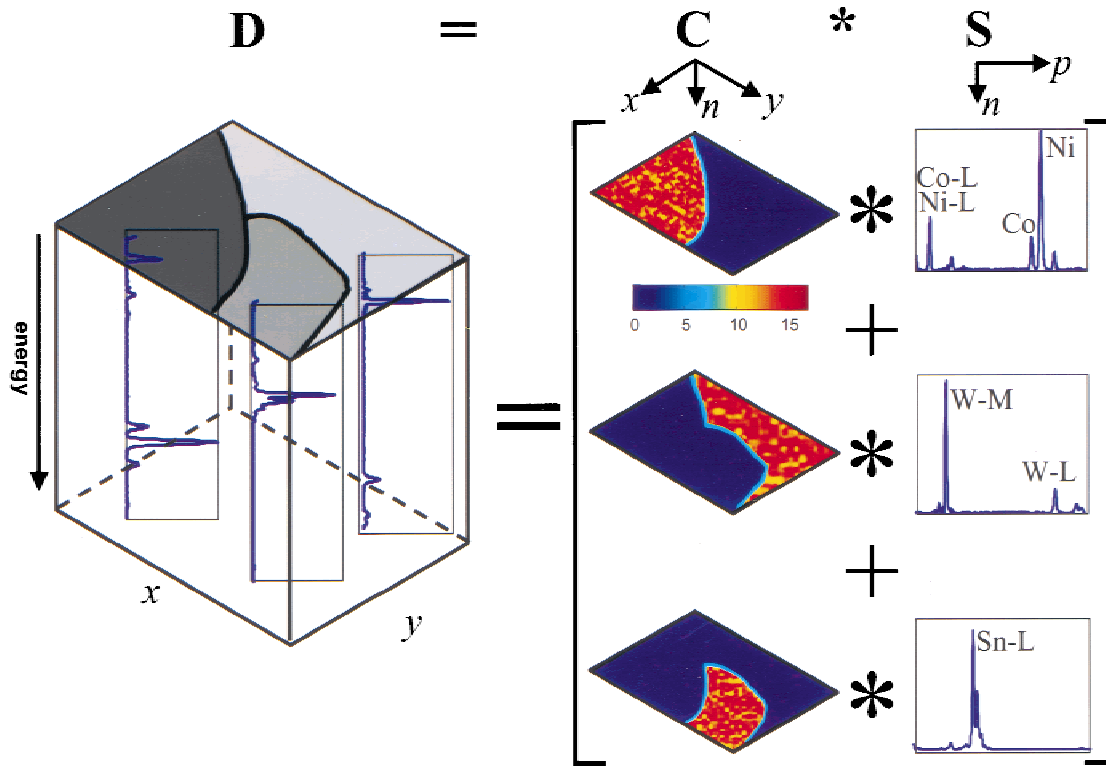


Figure 2. Relationship between the data cube and pure-component images and spectra for a hypothetical case where n , the number of components automatically determined by the software, is 3, x and y are real-space pixel coordinates, and p is the number of energy channels per spectrum.

spectral image such as this is a spectrum from each distinct chemical component present and a set of images describing where and to what extent in the microstructure the chemical components are. Any one of a number of MSA methods can be applied to a problem such as this but one key assumption was made in this work to simplify the necessary computations. We assume that any raw spectrum can be reconstructed by a linear combination of the product of the MSA-determined component spectra and their respective amplitudes or weighting (the pixel intensity in the component image). This can be written as an equation where s_i is the i th raw spectrum from the original spectrum image:

$$s_i = C_{(i,1)}S_1 + C_{(i,2)}S_2 + C_{(i,3)}S_3 + \dots + C_{(i,n)}S_n + E_i. \quad (1)$$

The $C_{(i,n)}$ terms are the amplitudes or weighting of the component spectra S_n where n describes the number of nonnoise components found. The C terms can be used to construct component images that then describe the spatial distribution of the signal described by the component spectra. The last term, E_i , describes the residuals due to Poisson noise. Equivalently, we could describe the raw spectral im-

age data matrix, \mathbf{D} , as the product of two new matrices, \mathbf{C} and \mathbf{S} :

$$\mathbf{D} = \mathbf{C} \cdot \mathbf{S}^T. \quad (2)$$

\mathbf{C} is a matrix of concentrations or component images of dimension $(x \cdot y) \cdot n$, and \mathbf{S}^T (where \mathbf{T} denotes the matrix transpose) is a matrix of spectral shapes of dimension $p \cdot n$, where p is the number of channels per spectrum. The relationship between the original data cube and the MSA solution is shown schematically in Figure 2 for a hypothetical microstructure with three chemical components.

Typically, the first step in MSA calculations is to perform orthogonal factorization of the data into the product $\mathbf{C} \cdot \mathbf{S}^T$ by, for example, a singular value decomposition (SVD; Malinowski, 1991). SVD is simply one method for computing a principal components analysis (PCA). Trebbia and Bonnet (1990) have described the details of PCA performed on the raw spectral image data and the importance of suitable normalization. From a geometric point of view, PCA begins by finding the single direction in the p -dimensional space that best describes the location of the

data within the space. The vector describing that is the first principal component. Once found, a second direction orthogonal to the first is determined that best accounts for the variation in the data that is orthogonal to the first direction. This is the second principal component. The process continues with each new principal component maximally accounting for the variation in the data orthogonal to all preceding components. The first few principal components will typically contain the chemical information of interest, while the rest will describe noise or error. In practice, PCA will return an answer of the same dimension, p , as the original spectral image that can then be manually reduced by inspection (or automatically as in this work). The result of orthogonal factorization (i.e., PCA) is the imposition of orthogonality on the estimated principal components. The drawback of PCA, by itself, is that the answer it gives is abstract and not readily interpretable or even physical as principal component spectra and images will have negative intensities. Real spectra are not orthogonal as PCA assumes they are. There will therefore be mixing of signals from chemically distinct chemical constituents in a single principal component. To further complicate interpretation of PCA results, information about different chemical constituents will be convolved not only in the principal component spectra but in the images as well. Given the complexity of the principal components, the utility of spending a large amount of time analyzing PCA results for the typical analyst is questionable.

It is for these reasons that a new method for converting from *abstract* principal components to *physically meaningful* pure components has been developed, based in part upon multivariate curve resolution techniques (Andrew and Hancewicz, 1998; Tauler and deJuan, 2002). The first application of the present method has been presented elsewhere (Kotula and Keenan, 2000). The terminology of pure components implies that the components are no longer mixed as the principal components were. This is achieved by relaxing the orthogonality constraint imposed by PCA and applying constrained alternating least squares algorithms to equation (2) above. Although orthogonality has been relaxed, linearity has not and this will result in nonintuitive results under some circumstances as will be described. This typically manifests itself as extra components that are linear approximations to nonlinear effects such as absorption and can be extremely useful for understanding geometrical effects in microanalysis. This novel approach to the analysis of spectral image data sets will be demonstrated for a metal-ceramic braze joint with seven chemical phases, a Cu-Ni

(continuous solid solution) diffusion couple, particle analysis, and detection of a single pixel of Pt in a Si matrix. For several of these examples, a quantitative comparison between raw and processed data will be made to demonstrate the effectiveness of the current spectral image data-analysis approach.

RESULTS AND DISCUSSION

Analysis of a Metal-Ceramic Braze

Braze joints between metals and ceramics are typically rich in complex but technologically important chemical interactions. It is for this reason that the first example of automated X-ray spectral image analysis is of a braze between copper and alumina in which seven distinct chemical components were detected with the automated spectral image analysis technique described above. The spectral image was acquired with a JEOL 840 SEM with W-filament, operated at 20 kV with a 10 mm², 138-eV-resolution Si (Li) X-ray detector. The spectral image was collected with the NORAN Vantage Digital Imaging with Spectral Imaging (DSI) with a resolution of 128 × 128 pixels (16,384 spectra) at 1024 channels per spectrum, a pulse processor shaping time of 9 μs, and 50% dead time. To build the spectral image, 36 frames were collected at 1 ms/pixel/frame for a total acquisition time of 10 min or 36 ms/pixel. This resulted in an average of 300 total counts per spectrum. Figure 3 shows the results of PCA of this spectral image data set. The principal components are abstract and difficult to interpret. Sn, for example, is present to some degree in every principal component. Additionally, all of the principal component images and spectra are mixed and contain both positive and negative intensities. It is for these reasons that the present algorithms were developed to automatically convert abstract principal components into physically meaningful and more readily interpreted pure components. The automated spectral image analysis in this case took approximately 1 min on a dual 500-MHz Pentium III PC with 1 Gbyte RAM. This included both the PCA and new algorithms. The matrix size was 134 Mbyte and the maximum amount of memory utilized by the analysis software was 170 Mbyte of RAM (i.e., 36 Mbyte memory overhead). Figure 4 is a RGB composite of all the chemical components found. Associated with each of the pure component images used to construct the RGB image is a pure component spectrum

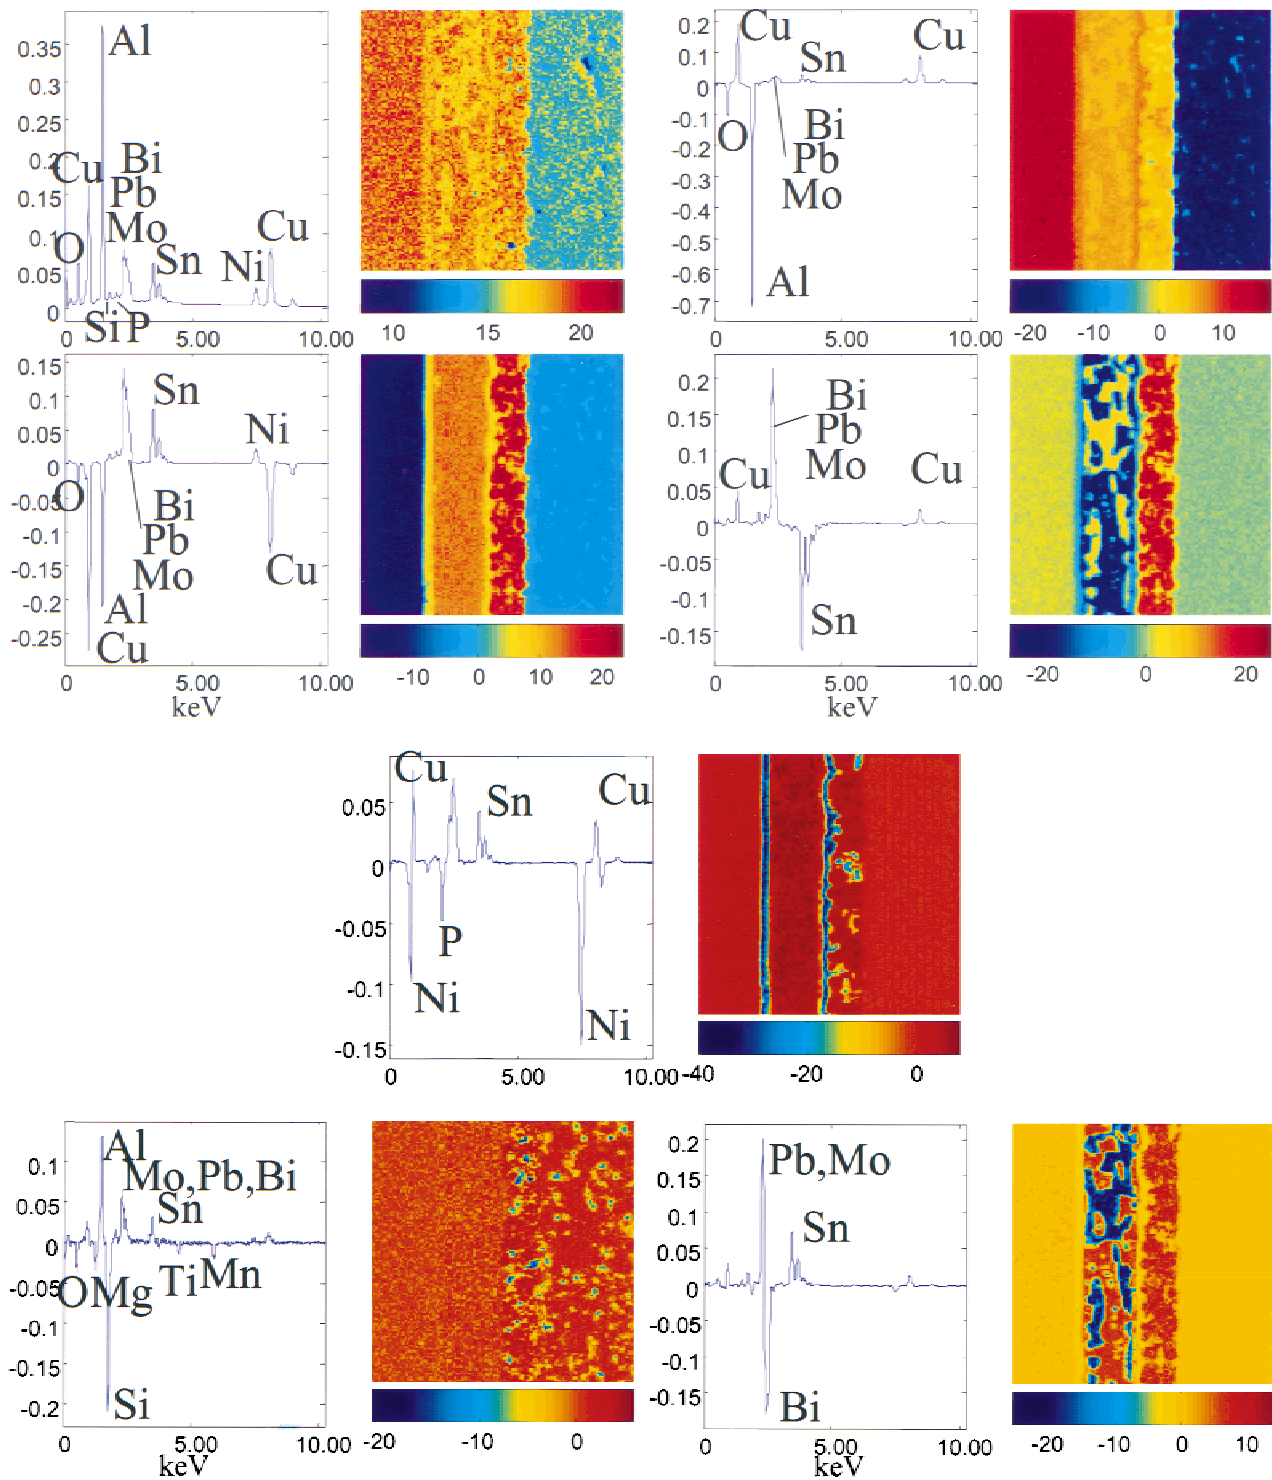


Figure 3. Results of a principal components analysis of the spectral image from a braze illustrating the difficulty of interpreting the abstract principal components.




Figure 4. RGB composite of the results of an automated analysis of a spectral image from a braze in a thermoelectric cooler. Cu is red, Ni-P is green, Al_2O_3 is blue, Bi is magenta, Sn is yellow, Mo/Pb is cyan, and silicate glass is white.

from which the qualitative chemical composition of the component was determined. The pure component spectra and respective images from the analysis are shown in Figure 5. The pure component or concentration images in Figure 5 (and elsewhere in this article) are based on a color scale with dark-blue being zero absolute concentration and dark-red being maximum absolute concentration or amplitude of the respective pure-component spectrum. The actual numbers corresponding to a given color are shown below each image. The pure-component spectral intensities are normalized counts. Therefore, to get back to actual counts for a given pixel, simply multiply the number corresponding to the color of a pixel by the normalized counts in the pure-component spectrum. To form the RGB image in Figure 4, the real-number-intensity pure-component images were converted into 8-bit TIFF format by scaling the minimum and maximum values to 0 and 255 respectively.

Ultimately a direct comparison between raw and reconstructed data is required to judge the effectiveness of the automated analysis. Figure 6 compares one raw pixel from each of the Sn and Pb-Mo phases with the respective reconstructed spectra from the same pixels. To reconstruct the spectrum for a given pixel of a distinct chemical component we simply multiply, using Sn as an example, the

vectors $\mathbf{C}_{\text{Sn}} \cdot \mathbf{S}_{\text{Sn}}^T$, which returns a matrix the same size as the raw spectral image. We can then directly compare the same raw versus reconstructed pixel.

The braze-analysis example shown above demonstrates the power of the present automated X-ray spectral image analysis technique. The solution, as shown in Figures 4 and 5, was produced with no operator intervention in approximately 1 min. While Figure 4 is useful for displaying the interpreted solution, it is only part of the answer to the microanalysis problem at hand. Figure 5 illustrates the two-orders-of-magnitude reduction in the amount of data needed to describe the relevant information contained in the raw spectral image. The concentration matrix \mathbf{C} in this case has dimensions of 128×128 pixels \times 7 pure components and displays the spatial variations (i.e., as images) of the respective pure-component spectra. The respective spectral-shape matrix \mathbf{S} has dimensions of 1024 energy channels \times 7 pure components and consists of the pure-component spectra. Therefore, the original 16.8 million data points (16,384 spectra \times 1024 energy channels) has been reduced to 122 thousand data points in 1 min with no loss of nonnoise information. This dimensional reduction allows the analyst to concentrate on interpreting the results without spending much time on either acquisition or analysis and with less fear of missing something important within the limits of sampling (i.e., pixel density) and counting statistics. The comprehensive solution in Figures 4 and 5 also illustrates another significant advantage to using MSA over more manual analysis methods—the ability to perform analysis on noisy data. For traditional X-ray mapping, acquisition times are typically quite long, as the goal is to visualize differences in X-ray image intensity. In addition to the long acquisition times and potential artifacts involved with X-ray mapping, it is also a time-intensive analysis method, as correlations between regions in different maps must be made manually. MSA methods, on the other hand, make full use of the redundancy inherent to spectral image data, and have, in this work, been fully automated and made computationally efficient. The spectral image of the braze above was collected in 10 min, resulting in an average of 300 counts in each spectrum, and was comprehensively analyzed in 1 min. The typical raw spectrum has a full scale of 10 to 20 counts, making manual analysis of such noisy data impractical. In part because of the redundancy embodied by a large number of noisy spectra, such data pose no great difficulty for MSA techniques.

Direct comparison between raw and reconstructed pixels for the case of Sn, in Figure 6A, shows that the entire

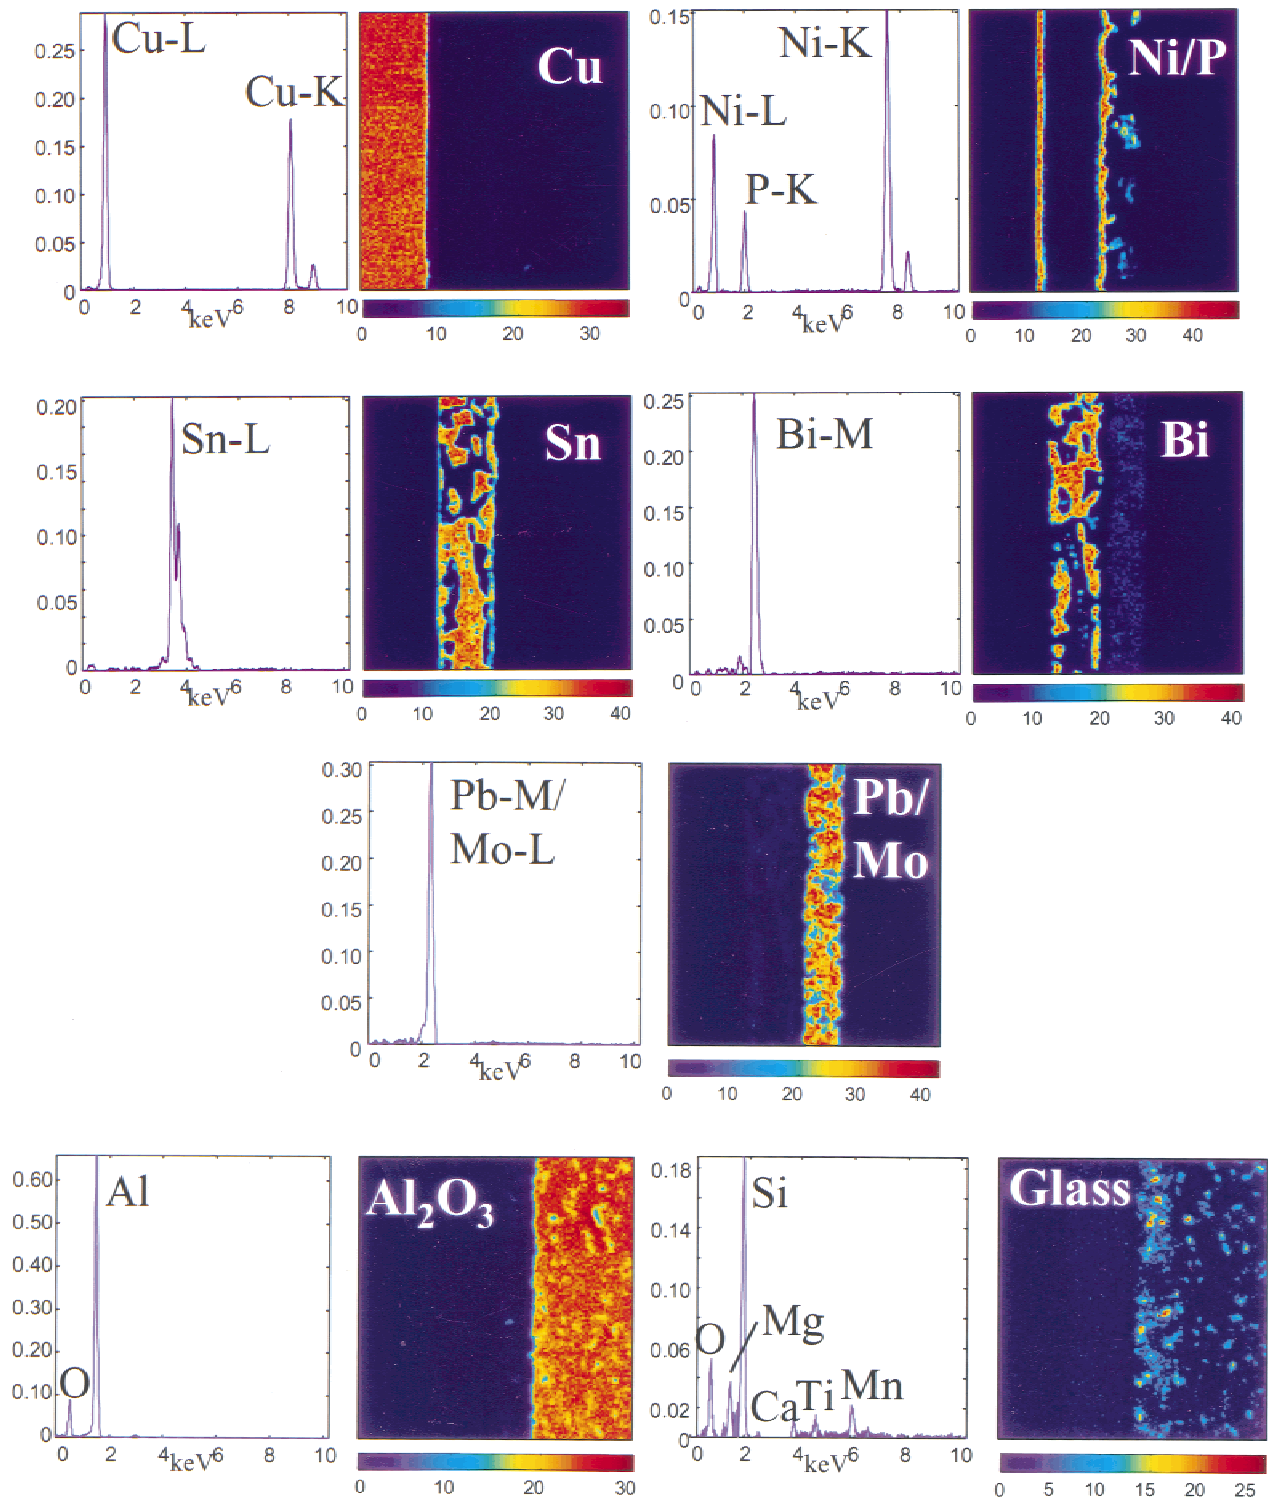


Figure 5. Pure components from the automated analysis of a spectral image from a braze in a thermoelectric cooler. The vertical units on the spectra are normalized counts. The pure-component images were used to construct the RGB image in Figure 4.

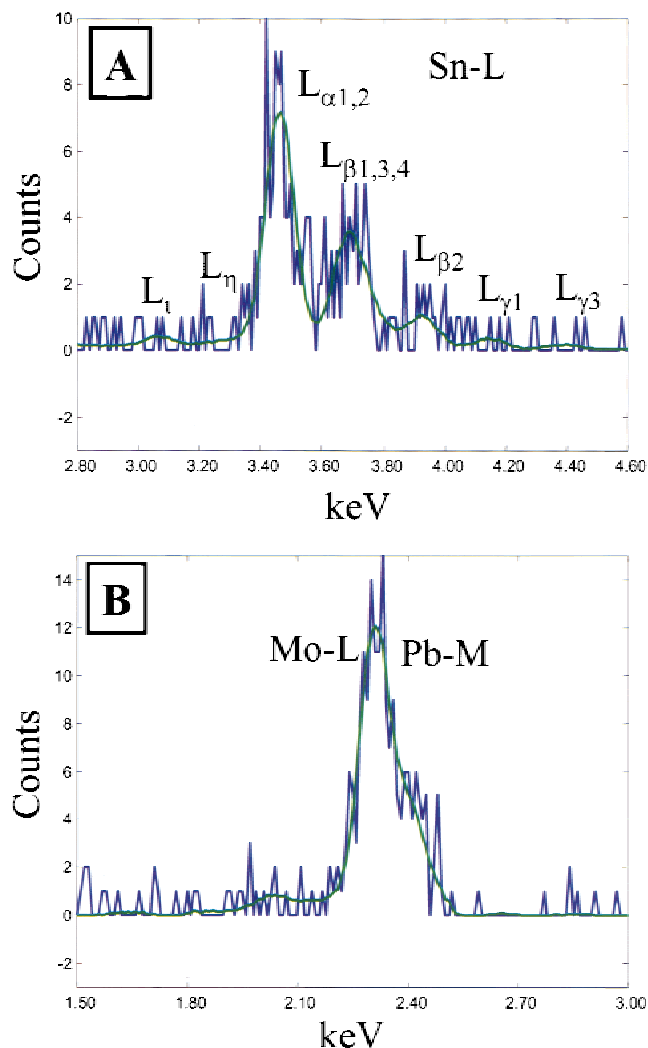


Figure 6. **A:** Comparison of raw spectrum from one pixel in the Sn phase (blue) and the spectrum from the same pixel reconstructed (green) from the braze spectral image data set. **B:** Comparison of raw spectrum from one pixel in the Pb/Mo phase (blue) and the spectrum from the same pixel reconstructed (green).

family of peaks for a given element is preserved. Even though the software makes no assumptions about the shapes and relative intensities of X-ray lines for a given element, the information is maintained through the analysis. The primary difference between the raw and reconstructed spectra for Sn is the relative level of noise. The reconstructed spectrum is a statistical aggregation of all the Sn in the spectral image. The residual, from the subtraction of the raw and processed spectra (not shown), is quite small except where there are a large number of counts, as under the large peaks. This is a direct outcome of Poisson statistics, where the noise level of a measurement is equal to the square root

of the number of counts. There will be more noise, in an absolute sense, under a large peak as there are more counts there. The same phenomenon is observed for the linear-least-square (LLSQ) fitting of any spectroscopic data whose acquisition is governed by counting discrete events. In fact, the way in which the automated analysis software works can be thought of as LLSQ fitting of internally determined “reference” spectra to all the raw pixels. The reference spectra however bear statistical aggregation and filtering, and therefore have a larger number of counts than the individual raw spectra.

Close examination of the pure-component images (concentration maps) from Figure 5 shows that the boundaries between chemical components are not abrupt, but span several pixels. For example, the Sn pure-component spectrum derives from regions of the spectral image where Sn is by itself and also from mixed regions where it is not spatially resolved, for example, at chemical phase boundaries. At a boundary between two compounds where signals from both are mixed, the software quantitates the proportion of the spectral shapes in each. The mixed pixel is then composed of a linear combination of the spectral signals from the adjacent phases as far as the software is concerned. This is equivalent to subpixel spectral demixing. Perhaps the best example of this in Figure 5 is the glass component. There are no pure glass pixels (spectra) in the raw data because the signal from the glass is always mixed with that of the alumina and Pb/Mo components. In spite of this, the automated analysis software identifies the glass as a chemically distinct component of the microstructure. There are analysis methods that look for the pixels that most represent the chemical phases in a raw data set and then categorize the raw spectra from this basis of “pure” (or purest) spectra. Such methods fail when there are not pure (unmixed) spectra for every chemical component. It is important to note that the MSA method employed here does not categorize pixels as belonging to a class. Classification methods such as pattern recognition will typically describe “fuzzy” boundaries as separate “boundary” components or as one or other of the components adjacent to the boundary. If, for example, Sn were adjacent to Pb, then all along the boundary between the two either a second component corresponding to Sn-Pb would be identified or the pixels would be classified as either Sn or Pb.

Examination of the Pb/Mo component spectrum in Figure 6B illustrates again the lack of assumptions about the generation and detection of X rays in the automated analysis software. In this case, Pb ($M_{\alpha 1,2}$ 2.342 kV) and Mo ($L_{\alpha 1}$

at 2.293 kV and L_{β_1} at 2.394 kV) are present in the same component and appear to be one set of peaks. It is by closer inspection of the pure-component spectrum that it is clear that two elements are present although not spectrally separated. In fact it is one of the great strengths of MSA techniques that these correlations are maintained, with data that is noisy. The typical Sn spectrum has about 10 counts full-scale while the typical Pb-Mo spectrum has about 15. It is clear from examination of Figure 5 that the individual elemental spectral correlations (i.e., families of X-ray lines) as well as correlations between different elements (e.g., as in the Pb-Mo and Al-O cases) are maintained. There are, however, several circumstances when this may not be the case and some of these will be illustrated in the following examples.

While the braze example is particularly straightforward to interpret, this is not always the case, as will be illustrated in the following examples. One should note that the present automated analysis software does just that; analyze spectral image data sets in an objective way. It does not however interpret the solutions. There are a number of circumstances when the solution may seem nonintuitive. First, it is important to note that the software makes no assumptions about the nature of X-ray generation and detection. The software does not know that for a given element there may be a family of X-ray lines present. This can be illustrated quite dramatically by scrambling all the pixels in the spectral image. Analysis of the scrambled spectral image will produce the same pure component spectra as that of the original spectral image. The pure-component images will, however, carry the same scrambling as the original spectral image. Precisely the same argument can be made for the energy channels. Therefore, physical proximity of pixels or channels is meaningless for the analysis itself, but still essential for the interpretation of the results.

Analysis of a Cu-Ni Diffusion Couple

Diffusion processes occur in virtually every class of materials. The frequent observation of diffusion phenomena in materials is the primary reason for analyzing a spectral image from a Cu-Ni diffusion couple. Cu-Ni is a simple binary system with complete solid solubility across the entire compositional range and is well understood. Discs of high-purity Cu and Ni were carefully polished flat and smooth, cleaned, and assembled in a clamp. The entire assembly was then heated under inert gas for 24 h at 900°C. A cross section was then prepared by conventional metallog-

raphy. It should be noted that this specimen is not perfectly smooth and contains Kirkendall voids resulting from the interdiffusion process. Spectral image data sets were acquired from the diffusion zone at accelerating voltages of 7 and 20 kV in a JEOL 5900LV using a NORAN Vantage-Digital Imaging with Spectral Imaging. At 7 kV, only the L -lines for Cu and Ni are excited, so this case should pose the greatest challenge for the automated spectral image analysis due to spectral overlap. The L_{α} peaks for Cu and Ni differ in energy by only 79 eV. The data sets were all acquired with a pulse processor shaping time of 9 μ s and a dwell time per pixel per frame of 40 μ s with 2000 frames to build the spectral image. The total acquisition time was 22 min or 80 ms total dwell time per pixel. For the 7-kV data set, the dead time was 25% with 1.2M counts in the spectral image and an average of 70 counts per spectrum. The 20-kV spectral image was acquired with 25% dead time and had 5.5M counts total with an average of 337 counts per spectrum. The data analysis times for both spectral image data sets were between 1 and 3 min. Figure 7 is the result of the analysis of the 7-kV spectral image data set. In this case, the automated analysis returned two pure components to describe the diffusion couple—pure Ni and pure Cu. The results of the analysis of the spectral image data set collected at 20 kV are shown in Figure 8. In this case, the software has identified three components.

In the case of a simple diffusion couple, the results of the automated analysis may at first seem nonintuitive. In fact, however, the result is a natural outcome of the linearity assumption present in the MSA calculations and will have implications for a whole range of related microanalysis problems. Without *a priori* knowledge of the nature of the problem, the automated analysis software has provided a solution that can readily be interpreted. The two components identified in the analysis, pure Cu and pure Ni, were separated at 7 kV even though there is considerable spectral overlap between the L -lines of the two (Fig. 7). The results of the analysis of the 20-kV spectral image, shown in Figure 8, yielded three components, two of which are similar to those from the 7-kV analysis, pure Ni and pure Cu. The third can readily be explained due to nonlinearities induced by absorption and fluorescence of X rays.

In virtually every X-ray spectral image acquired, there will be adjacent chemical constituents that may or may not be mixed in reality but will appear mixed due to the size of the X-ray generation volume. In the braze example, in which there was no apparent chemical interaction between adjacent components, there is still mixing of signals, and to

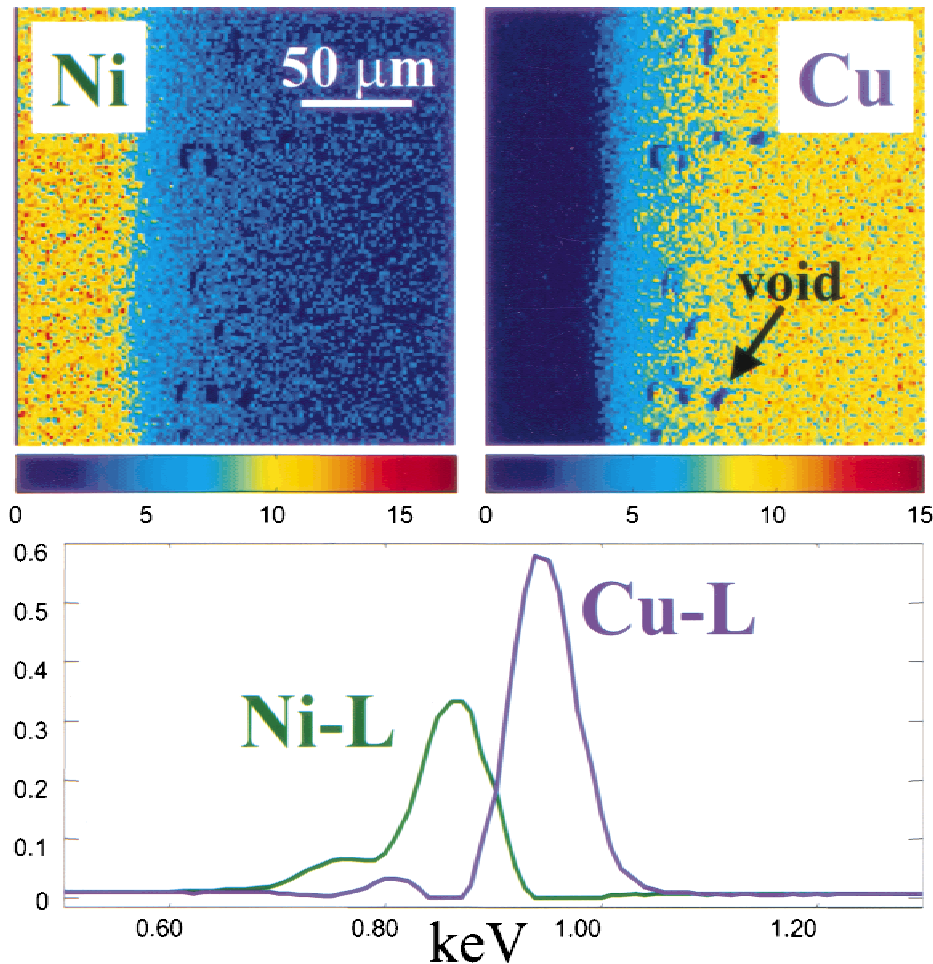


Figure 7. Results of the automated analysis of a spectral image acquired at 7 kV from the Cu-Ni diffusion couple. The vertical units on the spectrum are normalized counts.

a first-order approximation, the mixed signal can be said to be a linear combination of the signals from the adjacent components. If the two components happened to be chemically reactive, then interdiffusion might occur under the right processing conditions and there would now be a chemical gradient present in the materials. The automated analysis can make no distinction between these two cases, so for a binary diffusion problem, two pure components will describe the spectral image data set. It is important to note that analysis methods that use thresholding or pattern recognition would turn the continuous diffusion profile into a step function: either Cu or Ni but not both. The problem is, in fact, more complicated as, in reality, the signals from mixed elements are not a linear combination of the pure signals from the elements. This can readily be explained due to absorption and fluorescence of X rays. If we take a pure Ni spectrum and a pure Cu spectrum and mathematically sum the two in equal parts, it is not equivalent to acquiring a spectrum from a 50–50 mixture of Cu-Ni. However, if a

series of spectral image data sets are acquired for different times, several different analysis regimes will occur. For extremely short acquisition times resulting in less than ~ 50 total counts per spectrum, the automated analysis solution will typically contain one component which is an average of all the pixels with no distinction of the fine details. In this case, the individual spectra are so noisy that no difference between pixels can be determined. For short to intermediate acquisition times (~ 50 – 500 total counts per spectrum), the solution will begin to take on more detail, but the linearity approximation will hold. This is, in fact, the most appropriate and convenient regime to work in. The solutions are typically intuitive and the acquisition times might range from 5 to 30 min. For long acquisition times, nonlinearity effects will become apparent as extra pure components that approximate the nonlinearity. This is essentially a signal-to-noise (S/N) issue. As S/N increases, more spectral features will be distinguishable from noise. Ideally, the acquisition should be tailored to answer the question desired. Identify-

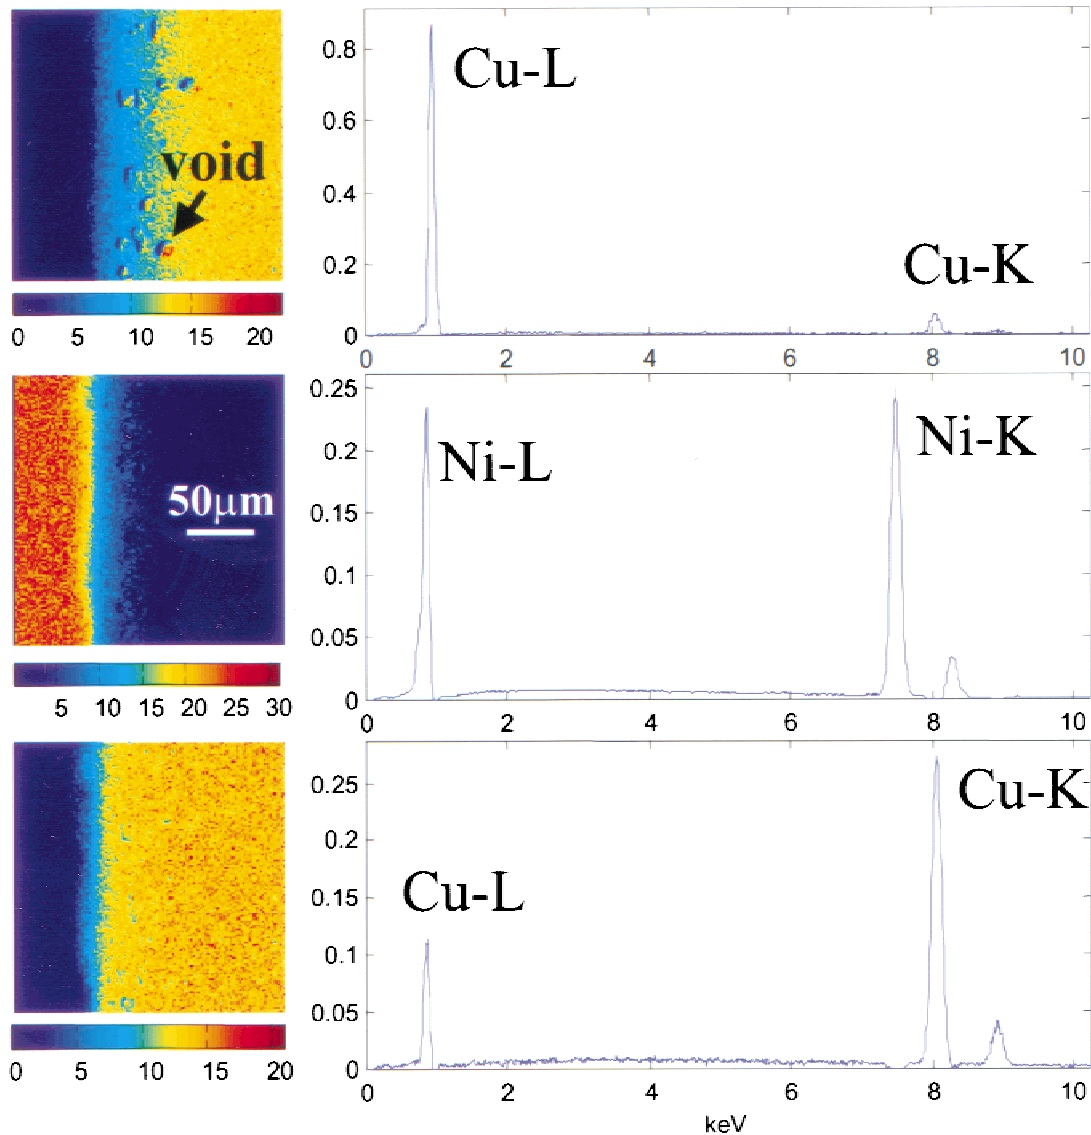


Figure 8. Results of the automated analysis of a spectral image acquired at 20 kV from the Cu-Ni diffusion couple. Note the extra component resulting from geometrical effects. The vertical units on the spectra are normalized counts.

ing several major components in a microstructure from spectral image data is easier, from an acquisition standpoint, than finding minor ones. An analyst looking for a minor component (the so-called needle in the haystack) and acquiring high S/N data should be prepared to have the MSA techniques return not only information about minor chemical components but also ones due to nonlinearity and varying X-ray generation/detection geometries.

Bulk microanalytical characterization may not always be performed on flat, polished specimens, and in the case of analysis of rough surfaces, the collected X-ray spectra may change depending upon geometrical conditions. This can readily be seen by examination of the results of the analyses

of the 7- and 20-kV spectral image data sets (Figs. 7 and 8). It was noted above that this diffusion couple is not perfectly flat and contains Kirkendall voids. The voids add topography to the surface that changes the X-ray collection geometry. Where there is an opened void, *L* X rays will be detected only from the part of the void directly facing the detector. The *L* X rays generated on the other side of the void (well below the surface) will be absorbed and not detected. For the present analysis, the detector looks towards the lower right-hand quadrant of the images (towards 4 or 5 o'clock). For the 7-kV spectral image analysis, where only the *L*-lines are excited, there will be an X-ray "shadow" facing the detector, and this can be seen as indi-




Figure 9. Pure-component spectra from a specimen with Ni on the surface and buried below carbon. The dashed spectrum shows almost no L X-ray intensity (Ni buried, for example) while the solid spectrum shows a much-reduced K X-ray intensity. The schematic shows the electron beam in two positions corresponding to the solid and dashed spectra, and the resultant X-ray generation volumes.

cated by the arrow in the Cu pure-component image in Figure 7. For the 20-kV spectral image analysis, three components were found as seen in Figure 8. In addition to pure Cu and pure Ni, there is a component showing mostly Cu-L. This component shows up most strongly in the lower right-hand side of the voids that face the detector. The voids are also farther into the Cu as the Cu diffuses faster than Ni. The arrow in the “void” component points to the same void as indicated by the arrow in Figure 7. The reason that for 20-kV there is an additional component from the analysis is that the K X rays are not similarly absorbed as the L X rays are. As stated above, there is no *a priori* knowledge in the analysis. So there is no expectation that the Cu K and L intensities should be related. In much the same way as the software recognizes the independence of the variations of L X rays from Cu and Ni, it also recognizes the same independence in this case for the Cu- K and Cu- L X rays. It is an important point that absorption can have much the same effect as variable composition, resulting in unexpected components from the analysis. This phenomenon is not limited to rough samples, as illustrated in Figure 9, which is a schematic of a sample with a smooth surface that showed much the same effect as the voids in the diffusion couple. The sample contained both Ni and C in a complex geometry. The Ni in some of the regions analyzed was at the surface and in others was below a layer of C thick enough to absorb most of the Ni- L X rays. This resulted in two “Ni” components that vary independently of each other, one showing almost no L X rays coming from the regions of buried Ni and the other looking more typical of a bulk Ni spectrum with both K - and L -lines in the proportion expected for analysis at 15 kV.

Particle Analysis

The analysis of particles is a topic of critical importance for the semiconductor industry as well as atmospheric science. For both groups, a rapid and unbiased analysis of particulates is desired (Anderson and Small, 1998). In the present example, a particulate specimen was fabricated from known powders of FeCo and alumina that were dispersed on a carbon support. A spectral image data set was acquired with a JEOL 5900LV (W-filament) operating at 15 kV using the NORAN Vantage-Digital Imaging with Spectral Imaging. The spectral image was acquired with a pulse processor shaping time of $9 \mu\text{s}$ and a dwell time per pixel per frame of $40 \mu\text{s}$ with 158 frames to build the spectral image. The pixel sampling density (i.e., the spacing between pixels in both x and y directions) was $\sim 4 \mu\text{m}$. The total acquisition time was 2 min or 6 ms total dwell time per pixel. The dead time was 30% with 1.8M counts in the spectral image and an average of 108 counts per spectrum. The computation time for the automated analysis was 2 min and six pure components were found. The computer memory overhead for the calculation was the same as for the braze example above. In addition to the carbon support, three types of particles were found: alumina, FeCo, and, unexpectedly, Ca-S-Si-O. The FeCo particles were described by two components, one containing Fe and Co K -lines and one containing Fe and Co L -lines. One additional component was found for the regions of the carbon support shadowed by the particles. This component consisted of a weak C peak but no O as was found elsewhere mixed with C on the unshadowed support. Figure 10 is a RGB composite of the pure-component images shown in Figure 11. For the purpose of forming the




Figure 10. Example of automated particle analysis. In addition to the carbon support (red), three types of particles were found: FeCo (blue), alumina (green), and Ca-S-Si-O (cyan). The black regions are carbon support shadowed from the X-ray detector.

RGB image in Figure 10, the two FeCo component images were summed. As before, the real-numbered pure-component images were rescaled to 8-bit TIFF images in order to form the RGB image. The pure-component images and corresponding pure-component spectra are shown in Figure 11.

The problem of identifying unknown particles can be particularly challenging and time consuming if done manually. In the present automated particle analysis, over 50 particles were characterized with the spectral image acquired in 2 min and analyzed in another 2 min. In addition to finding the expected FeCo and alumina particles, a third type of particle was unexpectedly found to contain Ca-S-Si-O. Only one particle of this type was found and it was probably from atmospheric contamination, as no special precautions were taken in the preparation of the sample. The Ca-S-Si-O particle was found in a noisy data set even though it covers only ~ 40 pixels out of 16,384 (128×128). Finding this particle is essentially equivalent to finding the proverbial needle in the haystack. In fact, close examination of both Figures 10 and 11 reveal that particles one pixel or smaller in size were identified as either alumina or FeCo.

Geometrical factors will become important in the analysis of particles as they typically have rough, irregular surfaces. It is the irregular geometry of the present sample that

has resulted in the two extra pure components. Of the six pure components, our interpretation indicated that two were seemingly redundant. The first was the carbon support shadowed from the detector. The X-ray detector faces the lower right-hand quadrant of the images and the particles sit out on the surface, thus shadowing some of the support. Comparison of Figure 11A and B shows that the primary difference between shadowed and unshadowed support is the intensities of carbon and oxygen. The relative intensities of the carbon signals differ by a factor of 5 due to the absorption of these soft X rays by the particles. In addition to absorption of C and O K X rays, the background (not shown) also exhibits the effect of absorption by the particles. Spectrally then, there is a significant difference between the shadowed and unshadowed support which is identified as significant by the software. The analyst can, however, ignore the extra C component or add it back manually into the primary C component. The second seemingly redundant component was from the FeCo. In this case, two components were identified consisting of either *K* or *L* X rays only. The reason for this is entirely analogous to the example shown in Figure 9 and discussed in the Cu-Ni diffusion-couple section above. In short, for geometrical reasons, the *K*-to-*L* X-ray ratios vary depending upon whether the spectra come from the side of the particle facing the detector or not. Pixels towards the detector have a lower *K*-to-*L* intensity ratio than ones away from the detector. As the software makes no *a priori* assumptions about relative intensities of families of X-ray lines, it sees two signals that vary independently from one another. As was the case with the Cu and Ni components in the Cu-Ni diffusion couple, the FeCo *K* component and FeCo *L* component share a number of the same pixels. This implies that there are regions of intensity (in the pure-component images) common to both *K* and *L* components. To reconstruct the actual spectrum from a given location, it would be necessary to take a linear combination of the two. In reality, if we wanted to quantify the composition of the FeCo, we would use the *K* component and disregard the *L* component.

The spectral image in the present example undersampled the microstructure. Data sets with 256×256 pixel spectral images have been acquired for comparable pixel dwell times as above and analyzed on the same PC. In this case, the computer memory required to perform the calculation was 700 Mbytes, still comfortably within the 1 Gbyte of RAM present on the computer used for this work and readily available on new computers. This opens up the possibility, for the purposes of particle analysis, of fully

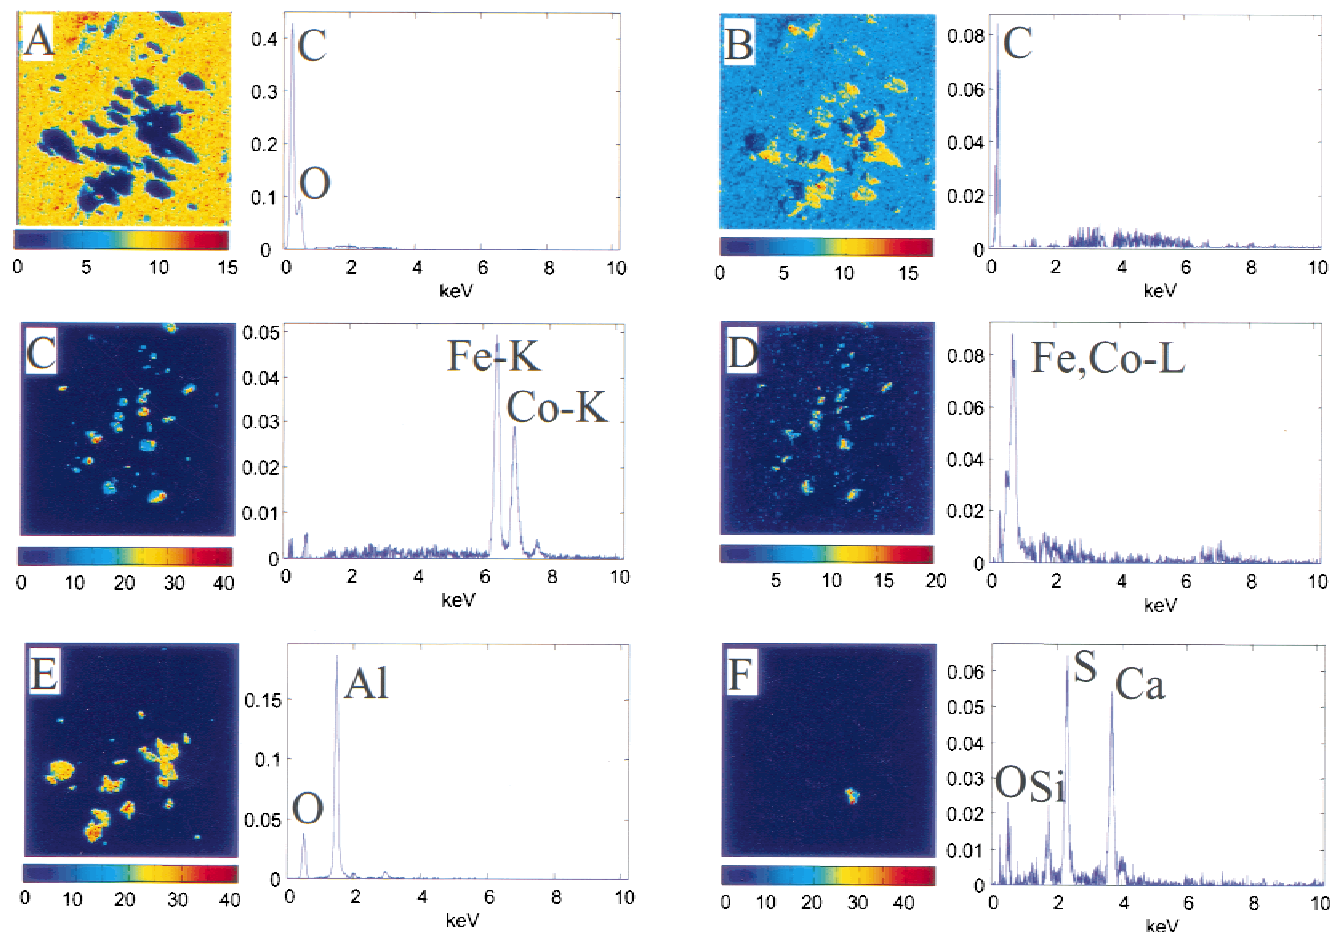


Figure 11. Pure components (images and spectra) for the particle-analysis spectral image. **A:** Carbon support. **B:** Carbon shadowed from detector by alumina particles. **C:** FeCo K-component. **D:** FeCo *L* component, towards detector. **E:** Alumina particles. **F:** Ca-S-Si-O particle. The vertical units on the spectra are normalized counts.

sampling large areas. At $500\times$ magnification on the present system, the field of view for a spectral image is $\sim 200\ \mu\text{m}$. At 256×256 pixel resolution, this allows comprehensive sampling of a surface at submicron resolution. The factor of 4-times-longer ($4\times$ total pixels) time necessary to match the acquisition shown above results only in a 10-min acquisition and approximately 5 min for the automated analysis.

Single-Pixel Detection

The question of detection limits and sampling is often posed in microanalysis (Goldstein et al., 1992) and is of particular relevance to multivariate statistical analysis methods (Anderson and Small, 1998). In an effort to understand the limitations of noisy spectral image data and the present automated analysis technique, a specimen was fabricated to approximate a single-pixel detection problem. The specific

question would be at what point could one pixel of Pt be detected in Si. The specimen was a single crystal of silicon with a “pixel” of Pt, $1 \times 1\ \mu\text{m}$ in lateral dimension $\times 2\ \mu\text{m}$ deep. The specimen was fabricated in a FEI DB-235 Focused Ion Beam (FIB)/SEM. The hole was drilled into the surface of the silicon and then filled with Pt using the Ga ion-beam-assisted Pt deposition source in the FIB. The entire surface to be analyzed was then scanned several times with the ion beam to clean it of Pt incidentally deposited around the edge of the hole. Nine spectral image data sets were acquired with a range of per pixel total dwell times from 1 to 500 ms. The data were acquired on a JEOL 5900LV operated at 20 kV as described in the previous examples. The dwell time per pixel per frame was $1000\ \mu\text{s}$ with between 1 and 500 frames to build each spectral image and a pulse processor shaping time of $50\ \mu\text{s}$. The dead time for all acquisitions was 30%, which resulted in between ~ 3 and 1400 total

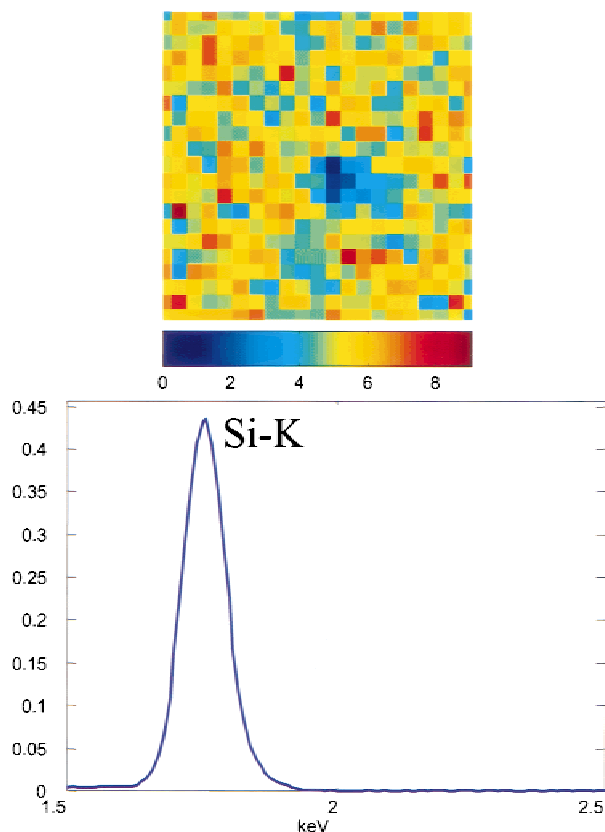


Figure 12. Pure components (image and spectrum) for the analysis of the single-pixel-detection spectral image with 10 ms total dwell time per pixel. The silicon pure component image is zoomed to 20×20 pixels about the location of the Pt. The vertical units on the spectrum are normalized counts.

counts per 1024 channel (10 eV/channel) spectrum, respectively. All spectral image data sets were acquired with a resolution of 128×128 pixels with a pixel sampling density of $\sim 0.8 \mu\text{m}$. This sampling density was slightly finer than the size of the target feature (the Pt “pixel”) in order to fully sample it. It was impractical to position the center of a given pixel over the center of the Pt. This and the fact the interaction volume is relatively large for Si at 20 kV means that the true number of pixels showing Pt could be nine or more. Calculation times for all data sets were less than 3 min each. Figures 12 and 13 show the results for spectral image data sets acquired with per pixel total dwell times of 10 and 20 ms. The pure component images have all been zoomed to the location of the Pt. The analysis software automatically found only one component, namely Si, for the 10-ms spectral image, while it found two, Si and Pt, for the 20-ms spectral image.

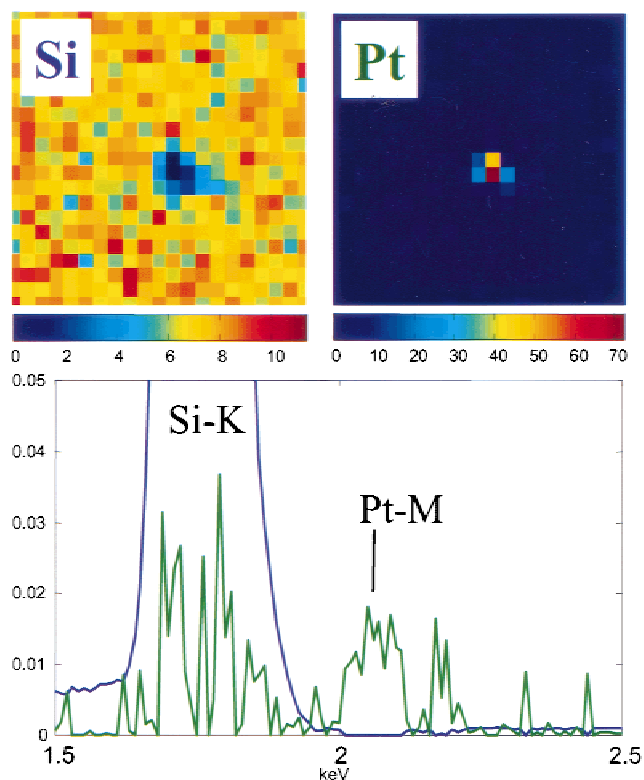


Figure 13. Pure components (images and spectra) for the analysis of the single-pixel-detection spectral image with 20 ms total dwell time per pixel. The pure-component images are zoomed to 20×20 pixels about the Pt. The vertical units on the spectra are normalized counts. Full scale for the silicon pure-component spectrum was 0.62.

One motivation for trying to resolve the one-pixel-detection problem is equivalent to asking what is the real advantage of using MSA techniques to analyze large data sets. Alternatively we might wish to ask how long we should acquire data to be sure that small, localized contaminants are not present. MSA techniques, as implemented in this work, clearly provide a robust and objective way to analyze large spectral image data sets. In the present example, a particularly difficult problem was posed to the software, namely to find one pixel of something other than Si in Si. In fact, for experimental reasons, it was not possible to achieve exactly one pixel but rather one to two total pixels. Given the experimental limitations, there was a clear crossover from one to two pure components, corresponding to the detection of Pt, between per pixel total dwell times of 10 and 20 ms, as shown in Figures 12 and 13, respectively. This corresponded to 28 and 56 total counts per pixel, or ~ 3.5

and 7 counts full-scale in the Si peak, respectively. For the Pt-*M* peak, this corresponded to ~ 1.5 counts full-scale for the 20 ms/pixel acquisition. To characterize the relative signal and noise levels in each of the spectral image data sets, the eigenvalues for the respective data sets were computed. The resultant vector of eigenvalues has a length of p (i.e., the number of variables or energy channels) and contains the absolute contribution of each of the components (or eigenvectors) to the solution. The majority of eigenvalues, plotted in descending order, fall on a straight line on a semi-log plot and simply describe the Poisson statistics of counting photons and therefore the noise baseline. Any component that contains nonnoise information will be above this line and the relative S/N level can be calculated by the ratio of the nonnoise versus noise eigenvalues. If this is done for the Pt pure component in Figure 13, a value of 1.09 is obtained (a value for S/N of 1 corresponds to only noise). A similar calculation was done for the first apparent noise component for the 10-ms spectral image analysis and a value of 1.01 was obtained. These S/N numbers take into account the collective statistics of all the similar pixels (i.e., spectra from the original spectral image) as well as all the channels that vary similarly in those pixels. For the 20-ms spectral image, at least four pixels contribute some part to the Pt pure component, but these signals could be folded back into a smaller number of Pt-only pixels.

The silicon pure-component image for the 10-ms spectral image is qualitatively the same as for the 20-ms spectral image, shown in Figure 13. That is, the number of dark-blue pixels in the Si pure-component image is similar. From the Pt pure component image in Figure 13, it appears that the Pt signal originated from between one and two pixels while the number of dark-blue Si pixels was significantly greater. This is probably due to Pt absorbing some of the Si-*K* X rays, as the pattern of pixels looks similar to those seen for the particle analysis example. Overall, the number of total counts per pixel (in a 128×128 pixel spectral image) to detect approximately one pixel of something different is about 56. A more practical experimental value (i.e., that provides better S/N) for the number of counts per pixel to acquire to find one different pixel would be ~ 300 . In fact, test calculations on fabricated data sets placed the crossover from nondetection to detection of one different pixel between 30 and 300 counts per spectrum. The present experimental example of Pt in Si was chosen to avoid a great degree of spectral overlap. If the example of one pixel of Rb in Si had been chosen, the answer would be much

different, as the two differ in relevant peak energy (Rb-*L* and Si-*K*) by ~ 45 eV. With no other practical peaks to help distinguish the two below 10 kV, it is expected that this would take significantly more counts per pixel in the spectral image to solve. It is expected that the case of one pixel of W in Si (~ 35 eV difference between W-*M* and Si-*K*) would be straightforward to solve, although less so than the Pt-Si case, as the W-*L* lines would help differentiate the two.

CONCLUSIONS

Automated analysis of large spectral image data sets has been demonstrated with the application of MSA techniques developed in this work. Such analysis techniques clearly provide a robust and objective way to analyze large spectral image data sets in a computationally efficient way. Calculations on 128×128 pixel \times 1024 channel spectral image data sets take less than 3 min in general. No *a priori* assumptions about X-ray generation or detection are made by the software, which results in the greatest analytical sensitivity to the unexpected. As a result of this, however, under certain circumstances, the analysis may return nonintuitive, but ultimately understandable results. These include, for example, geometrical effects resulting in preferential absorption of one family of X-ray lines such as was the case with the voids in the Cu-Ni diffusion couple and with the analysis of particles. In this case, there will not be one total spectral shape that is characteristic of a given chemical phase or type of particle. Absent geometrical effects, two pure components will be necessary to describe the mixing of signals from two adjacent chemical constituents as will occur at phase boundaries (due to the large X-ray generation volume) and in diffusion couples. At both sharp and diffuse boundaries, the resultant mixed signals will be described as a linear combination of the adjacent pure components. Additionally nonlinearity effects resulting from absorption and fluorescence— $(A_{\text{spectrum}} + B_{\text{spectrum}}) \neq (A + B)_{\text{spectrum}}$ —will also become apparent under the right circumstances (e.g., Cu-Ni diffusion couple at 20 kV). Although not explicitly mentioned in the above examples, artifacts from the acquisition such as correlated electronic noise or spectrometer drift will also cause additional components in the automated analysis, but these will not, in general, correlate with microstructural features and can easily be identified/ignored. In spite of the potential for some nonintuitive results, the automated analysis performs a powerful

service for the analyst, quickly analyzing large, complex data sets, but does not interpret the results. It is clear from the examples shown in this work that the total number of counts per spectrum needed to be able to perform the automated MSA is quite small. Fewer than 300 total counts per spectrum can potentially find spectrally distinct components of the order of one pixel.

ACKNOWLEDGMENTS

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy (DOE) under contract DE-AC04-94AL85000. The authors thank David Rohde and Jim Harrington of Thermo NORAN for provision of the braze data set.

REFERENCES

- ANDERSON, I.M. (1998). Multivariate statistical analysis of low-voltage EDS spectrum images. *Microsc Microanal* **4** (Suppl. 2), 272–273.
- ANDERSON, I.M. (1999). Quantitative chemical mapping with low-voltage EDS spectrum imaging and multivariate statistical analysis (MSA). *Microsc Microanal* **5** (Suppl. 2), 318–319.
- ANDERSON, I.M. (2000). Spectrum imaging: Microanalysis for a new millennium. *Microsc Microanal* **6** (Suppl. 2), 1048–1049.
- ANDERSON, I.M. & SMALL, J.A. (1998). Multivariate statistical analysis of particle X-ray spectra. *Microsc Microanal* **4** (Suppl. 2), 202–203.
- ANDREW, J.J. & HANCEWICZ, T.M. (1998). Rapid analysis of Raman image data using two-way multivariate curve resolution. *Appl Spectrosc* **52**, 797.
- GELADI, P. & GRAHN, H. (1996). *Multivariate Image Analysis*. Chichester, England: John Wiley and Sons.
- GOLDSTEIN, J.I., NEWBURY, D.E., ECHLIN, P., JOY, D.C., ROMIG, A.D., JR., LYMAN, C.E., FIORI, C. & LIFSHIN, E. (1992). *Scanning Electron Microscopy and X-Ray Microanalysis*, 2nd ed. New York: Plenum Press.
- HARMAN, H.H. (1976). *Modern Factor Analysis*, 3rd ed. rev. Chicago: The University of Chicago Press.
- JEANGUILLAUME, C. & COLLIEX, C. (1989). Spectrum image: The next step in EELS digital acquisition and processing. *Ultramicroscopy* **28**, 252–257.
- KOTULA, P.G., ANDERSON, I.M. & STEPHENS, J.J. (1999). Statistical analysis of SEM-EDX spectrum images of a metal-ceramic braze. *Microsc Microanal* **5** (Suppl. 2), 806–807.
- KOTULA, P.G. & KEENAN, M.R. (2000). Information extraction: Statistical analysis to get the most from spectrum images. *Microsc Microanal* **6** (Suppl. 2), 1052–1053.
- LEGGE, G.J.F. & HAMMOND, I. (1979). Total quantitative recording of elemental maps and spectra with a scanning microprobe. *J Microsc* **117**, 201–210.
- LYMAN, C.E. (1986). Digital X-ray imaging of small particles. *Ultramicroscopy* **20**, 119–124.
- MALINOWSKI, E.R. (1991). *Factor Analysis in Chemistry*, 2nd ed. New York: John Wiley and Sons.
- MOTT, R.B. & FRIEL, J.J. (1999). Saving the photons: Mapping X-rays by position-tagged spectrometry. *J Microsc-Oxford* **193**, 2–14.
- MOTT, R.B., WALDMAN, C.G., BATCHELER, R. & FRIEL, J.J. (1995). Position-tagged spectrometry: A new approach for EDS spectrum imaging. In *Proc. Microscopy and Microanalysis, 1995*, Bailey, G.W., Ellisman, M.H., Hennigar, R.A. & Zaluzec, N.J. (Eds.), pp. 592–593. New York: Jones and Begell Publishing.
- NEWBURY, D.E. (1997). Chemical compositional mapping by microbeam analysis at the micrometer scale and finer. *Microelectronics J* **28**, 489–508.
- NEWBURY, D.E. & BRIGHT, D.S. (1999). Logarithmic 3-band color encoding: Robust method for display and comparison of compositional maps in electron probe X-ray microanalysis. *Microsc Microanal* **5**, 333–343.
- TAULER, R. & DE JUAN, A. (2002). Multivariate curve resolution home page. <http://www.ub.es/gesq/mcr/mcr.htm>, last updated May 2002.
- TREBBIA, P. & BONNET, N. (1990). EELS elemental mapping with unconventional methods I. Theoretical basis: Image analysis with multivariate statistics and entropy concepts. *Ultramicroscopy* **34**, 165–178.