Eurographics/ IEEE-VGTC Symposium on Visualization (2007)
Ken Museth, Torsten Möller, and Anders Ynnerman (Editors)

Parametric Visualization of High Resolution Correlated Multi-spectral Features Using PCA

A. Broersen[1], R. van Liere[1] and R. M. A Heeren[2]

[1]Centrum voor Wiskunde & Informatica, CWI, Amsterdam, The Netherlands
[2]FOM Institute for Atomic & Molecular Physics, Amsterdam, The Netherlands

Abstract

An imaging mass spectrometer is an analytical instrument that can determine the spatial distribution of chemical compounds on complex surfaces. The output of the device is a multi-spectral datacube; a three-dimensional (3D) dataset in which the xy-dimension represents the surface position and the z-dimension represents the mass spectral distribution. Analysts try to discover correlations in spectral profiles and spatial distributions inside a datacube. New technological developments allow mass spectrometric imaging on a higher spatial and spectral resolution. In this paper we present a parametric visualization technique which allows an analyst to examine spectral and spatially correlated patterns on the highest possible resolution. Principal component analysis (PCA) is used to decompose the datacube into several discriminating components. We represent these extracted features as abstract geometric shapes and use three parameters to allow for data exploration. The first parameter thresholds the spectral contribution at which an extracted component is visualized. The level of detail the shapes is controlled by a second parameter and a third parameter determines at which density-level the extracted feature is represented. This new visualization technique enables an analyst to select the most relevant spectral correlations and investigate their specific spatial distribution. With this method, less noise is included in the visualization of extracted features and by introducing various levels of detail the full spectral resolution can be utilized.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling I.4.7 [Image Processing and Computer Vision]: Feature Measurement I.4.10 [Image Processing and Computer Vision]: Image Representation

1. Introduction

An imaging mass spectrometer can determine the spatial distributions of chemical compounds directly on biological surfaces. Measuring a mass spectrum at a multitude of different locations with a spatial resolution around one micrometer results in large, three-dimensional (3D) multi-spectral datacubes. The xy-dimension of the datacube represents the locations at which the spectra were acquired. The z-dimension represents the molecular weight or mass-spectral dimension. The different surface compounds are separated by the mass-to-charge ratio of the intact molecules or their fragments. The combination of peaks associated with different molecular weights present in one location is often called a mass spectrum or *spectral profile*. The scientist uses these mass spectral profiles to analyze the surface composition. The problem they face is to determine/find spectral correlations related to different compounds and establish the spatial distributions of these correlated peaks within the millions of spectral variables [MSSK05]. Unfortunately, it is not always evident which spectral peak or combination of peaks to look at and to determine how they are distributed across the measured surface. Such a combination of correlated peaks and their spatial distribution can be defined as a *feature*. It is our intention to create exploratory visualization techniques with as few as possible data-specific denoising or complex clustering methods and still be able to visualize features within the full spectral and spatial resolution of these enormous datacubes. In the most simple case of exploration a spectral window is selected by hand using the histogram in Figure 1a, from which all intensities on one location are summed to create one image. Other methods almost always use a lim-

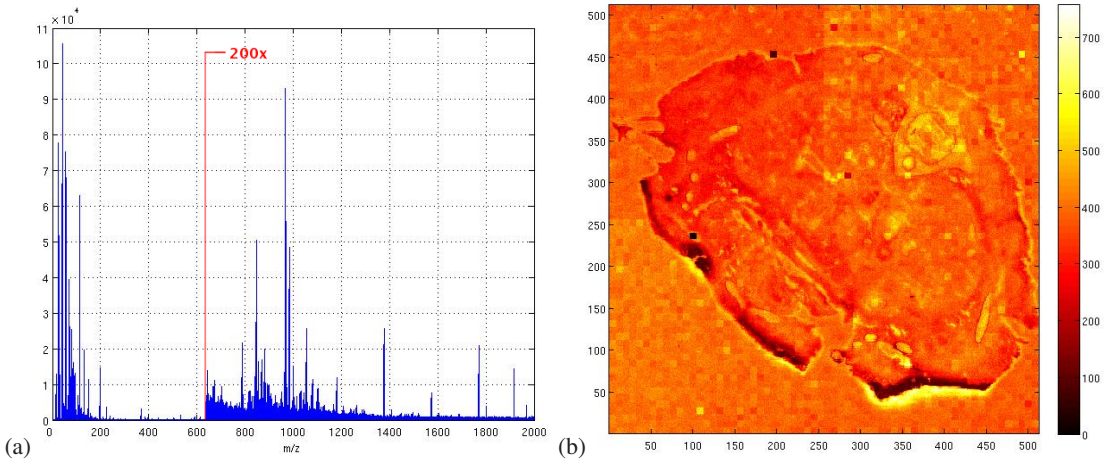

Figure 1: *(a) the sum of all* binned *spectral profiles with left part zoomed 200 times and (b) the matching sum of image planes*

ited set of spectral planes compared with the amount that imaging mass spectrometry supplies. Figure 1a shows the sum of all spectral profiles in a datacube, but this view had to be simplified by combining several neighboring spectral levels into one *bin*. This is commonly used and necessary in imaging spectrometry (see [Pac04]) to be able to create a spectral level for a peak instead of a peak distributed on multiple spectral levels. A second reason is that applying an analytical technique to an unbinned dataset is practically impossible due to time and memory restrictions. A third reason is that it increases the signal-to-noise level, but in doing so it removes some high resolution characteristics.

We present a new visualization technique that enables the user to:

- visualize spectral and spatially correlated and anti-correlated patterns
- use the highest possible resolution instead of a spectrally binned one
- extract features as 3D shapes with better defined boundaries
- parametrically explore multiple features within the same view

The feature visualization is controlled by three parameters. The first parameter α is set as a threshold on the spectral contribution of an extracted feature. In this way, only the spectrally correlated windows with the highest contributions to that extracted feature are included in the visualization. This enables a user to remove smaller peaks and noise that can clutter the visualization. A second parameter β controls the level of detail of a 3D feature. A feature can be represented as a simplified smooth 3D shape or on a high resolution that contains more details of the structure of a shape. The third parameter γ is used to set at which level of density in the data a geometric shape is created. A family of iso-surfaces can be created to explore the areas in the feature with different densities. Less noise is included when extracting iso-surfaces instead of extracting two-dimensional (2D) contours from intensity images. With these three parameters multiple correlated features can be displayed as 3D geometrical shapes that include less noise compared to a traditional 2D view.

2. Feature extraction and visualization

Various visualization techniques have been proposed to inspect datacubes. The most basic technique with a *summed intensity* image (Figure 1b), in which a side-by-side view of spectral and spatial domains can be analyzed. However, it is left to the user to identify which chemical compounds are present in the datacube and whether or not their spatial distributions are correlated. There are some complex fuzzy logic segmentation algorithms ([WSH99]) as well, but these can only be applied on a limited number of spectral windows. A few implementations also exist to visualize 3D spectral imaging data in the spectral or in spatial domain. Visualizing both implies coping with a number of difficulties. First, the 2D spatial information with added one-dimensional (1D) spectral information can not be treated in the same way as 'real' 3D volumes for instance as a result of a CT or MRI scan. To overcome this problem, most techniques first apply feature extraction using factor analysis, for instance Kenny et al. [KNM*97] or Keenan [Kee05]. Feature extraction is closely related to compression or dimension reduction techniques and target the removal of redundant data or data that mostly contains noise. Both approaches do not use the full available spectral resolution in their final visualizations. The second problem is to find the most appropriate technique for feature extraction or dimension reduction, which is specific for each spectral dataset. A third problem is the ever increasing size and resolution of the datasets. This problem makes both visualization and feature extraction more difficult even with increasing computational

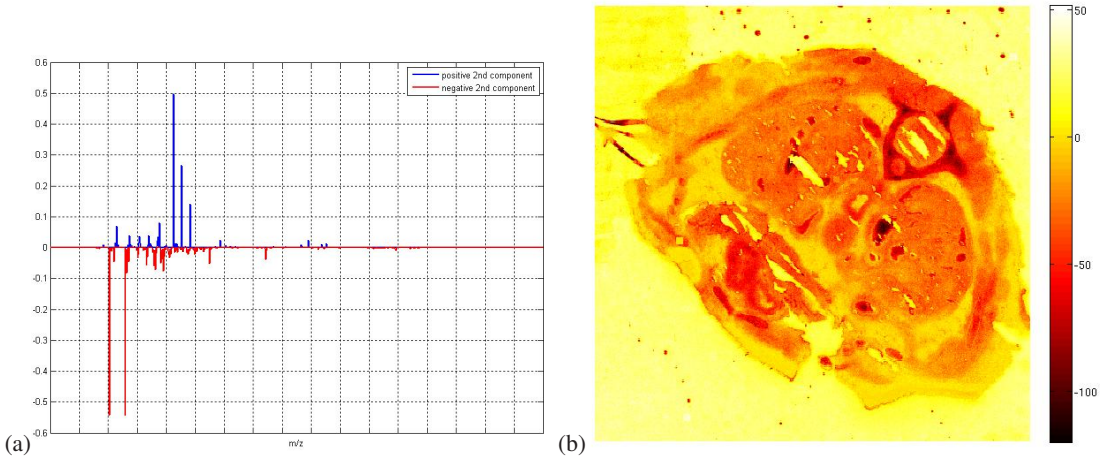

(a) (b)

Figure 2: *(a) the spectral profile of a second principal component and (b) the matching second image component*

power. Haigh et al. [HKR*97] for instance use correlation partitioning on five spectral channels after which they have to turn to dimension reduction techniques. Other visualization techniques use between $100-300$ spectral channels to visualize spectral data with flat image overlays with color weighting envelopes [JG05] or apply volume visualization techniques [PvdH01]. Both techniques work with a spectral dimension that falls within visible light and has continuous intensity values in spectral dimension with the wavelength. The spectral dimension resulting from mass spectrometry can have $\sim 2 \cdot 10^6$ intensity values and can be considered as a cloud of single 3D points. These two differences make the afore mentioned approaches for feature visualization not applicable on the datasets of mass spectrometry.

Multivariate statistical analysis tools are used in almost all attempts to explore and visualize the enormous datasets resulting from mass spectrometry. For instance, the tool AXSIA ([SOKK04]) statistically aggregates spectral profiles to identify features in the data, but the results are still shown as separate spectral profiles and summed spatial distributions. It claims to decompose the datacube more intuitively by disallowing negative spectral contributions. While most improvements to the data exploration tend to focus on denoising ([WKC03]) or specific 1D filtering techniques ([KK04]), traditional approaches like Principal Component Analysis (PCA) are fast and still one of the most successful multivariate tools ([KBF*07, Pac04]) for spectral feature selection and unsupervised exploration. One approach ([BvL05]) combines spectral and spatial results from PCA into one visualization of complete datacubes. Although correlations between spectral peaks and their spatial distribution can be studied in one view, one weakness is that there is no possibility to parametrically control feature extraction. A second shortcoming is that noise inside the spectral bins is also included in the resulting volume rendering. Spectral information is lost when combining the high-resolution spectral channels data into one bin, which is necessary before PCA can be applied. A last weakness is that it is impossible to select a spatial region of a specific spectral window inside an extracted principal component for further examination.

In our approach we focus on visual parametric exploration of the datacube. Although correlations between chemical components can be found unsupervised with PCA, much spectral information is lost when visualized in the traditional two dimensions. We use the full spectral resolution in feature visualization to reduce noise as much as possible without having to focus on advanced and computationally expensive algorithms. Not only positively correlated features are highlighted, but also their negative correlated counterparts in one parametrically simplified view.

3. Method

The features are extracted in a four-step process. First, principal component analysis is used to discriminate specific components present in the datacube according to their spectral correlation. Then, the most important spectral windows are parametrically selected to exclude smaller spectral contributions that contain more noise. In the third step, the selected windows are convolved into continuous scalar fields to be able to extract appropriate iso-surfaces from those regions where the data is the most dense. In the last step, correlations between extracted features are visualized in on their 2D locations with the additional high resolution spectral dimension. The adjustment of three parameters allows the user to interactively analyze and highlight the spatial and spectral distributions of the chemical elements and molecules on the surface of the material.

3.1. Principal Component Analysis

Different methods for decomposition or factor analysis can be used for this feature visualization. PCA ([Jol02]) has still satisfying results with respect to the speed, discrimination between extracted components and ability to identify correlations as well as anti-correlations between spectral and spatial dimensions. Keenan and Kotula [KK04] showed that in the case of mass spectrometry common preprocessing steps normalization or auto-scaling can lead to less than satisfactory results. For that reason we decided not to normalize or auto-scale in this technique.

First PCA is applied on the datacubes to extract the most important correlated spectral profiles. PCA is used to decompose and compress the thousands of spectral profiles into a few main components that capture the main characteristics of the data. The components contain those spectral peaks that are correlated. When sorted according to their eigenvalues, the first few components describe the most variance in the spectral data and therefore have the most contrast in the peak intensity. PCA is used by unfolding a non-preprocessed m by x by y datacube in such a way that a 2D m by $x \times y$ matrix X is constructed. The standard PCA model is used to compute a sorted list of principle components in an orthonormal matrix P (see Equation 1) using eigenvector decomposition.

$$Y = P \cdot X^T \tag{1}$$

The first principle components in P describe those spatial loadings that can take account for the most spectral information in the datacube. Each component is then used as a new base to project the original datacube. This results in a matrix Y with spectral score vectors, which can be interpreted as *spectral* components. Each peak in the resulting spectral component represents the contribution of a specific ion. An example of an extracted component is shown in Figure 2a. All positive peaks are colored blue and all negative peaks in red. The positive and negative part of this component are anti-correlated. The transposed datacube can be multiplied with the spectral component matrix (see Equation 2) to obtain the spatial distributions of these spectral correlations.

$$I = X^T \cdot Y \tag{2}$$

Each row in the resulting matrix I contains an unfolded image component containing the spatial contributions of each profile in Y.

All of the positive and negative values in a profile in Y contribute to a component, even when they are close to zero. The most important contributions in a component have the highest -positive or negative- contribution to the component. A threshold parameter α is defined to reduce the number of spectral bins that are used in the feature visualization. Equation 3 and 4 show that only those peaks that are above the threshold of α will remain in the part of the profile with the positive contributions (Y_α^+) or the part with the negative contributions (Y_α^-).

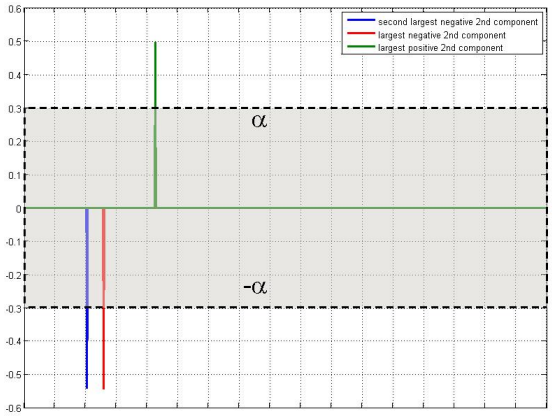

Figure 3: *selection of spectral peaks outside the greyed area with* $\alpha = 0.3$ *from the spectral profile of the second principal component from Figure 2a*

$$Y_\alpha^+(y) = \begin{cases} y, & if\ y \geq \alpha \\ 0, & otherwise \end{cases} \tag{3}$$

$$Y_\alpha^-(y) = \begin{cases} -y, & if\ y \leq -\alpha \\ 0, & otherwise \end{cases} \tag{4}$$

A good initial choice for α is often the highest possible value, so that only those peaks with highest positive or negative contribution remain for further processing. In this way, small or less important contributions that could contain more noise remain hidden at first. When α is lowered, more correlated spectral bins are added to the visualization that contribute less to a principal component but could contain some correlated spatial or spectral characteristics. For example an $\alpha = 0.3$ in the spectral profile of Figure 2a selects those three peaks with the largest contribution in this second principal component as shown in Figure 3. In this case we have two negative contributions and one positive contribution. A user can lower α and add more contributing correlated spectral windows to the resulting visualization.

3.2. Convolution

The datacube was binned before applying PCA and with α only the highest contributions in a principal component were selected. These selected spectral profiles will be used to extract the feature data from the original unbinned datacube. The resulting 3D clouds with the high-resolution ion-counts do not reveal a clear structure. Most ($\sim 99\%$) intensities have either value one ($\sim 9\%$) or zero ($\sim 90\%$). To be able to visualize more structural details from the cloud, a 3D convolution filter transforms the datacube into a scalar field with

continuous values. This low-pass frequency filter blurs the volume in such a way that those regions with a high concentration of data values can be represented by an iso-surface. When smaller kernel sizes are chosen, more fine-scaled anomalies will appear in a scale-space (see [Wit83, Koe84]) representation of the extracted features.

A second parameter β controls the size of the kernel and therefore the level of detail of the smoothed 3D feature. A standard Gaussian isotropic convolution kernel h_β is chosen for the smoothing as in Equation 5 in such a way that β is the variance of the Gaussian. $\|x\|$ is defined as the length of vector multidimensional x where n is dimensionality of vector x.

$$h_\beta(x) = (2\pi\beta)^{-n/2} \cdot e^{-\frac{\|x\|^2}{2\beta}} \tag{5}$$

For the practical applicability of the filter on the 3D datacubes in an interactive visualization, the discrete Fourier transform can be used to implement this convolution as in Geusebroek et al. [GSvdW02]. According to the convolution theorem, a convolution in spatial domain is equivalent to multiplication in frequency domain. The 3D convolution filter can now be defined in the frequency domain with the discrete Fourier transform in Equation 6, where $n = (n_x, n_y, n_m)$ and $k = (k_x, k_y, k_m)$ are defined as the three-dimensional vectors of indices of the selected datacube $N = (N_1, N_2, N_3)$ to simplify the equation.

$$F(k) = \sum_{n=0}^{N-1} f(n) \cdot e^{-2\pi i k n/N} \tag{6}$$

After the Fourier transformations of datacube $f(n)$ and filter $h_\beta(x)$, they are multiplied as in Equation 7 after which the inverse discrete Fourier transform in Equation 8 results in a convolved datacube.

$$G(k) = F(k) \star H_\beta(k) \tag{7}$$

$$g(n) = \frac{1}{\prod_{l=1}^{3} N_l} \sum_{k=0}^{N-1} G(k) \cdot e^{2\pi i n k/N} \tag{8}$$

It is now possible to extract iso-surfaces from the high density regions representing a high concentration of a specific element on a certain location without losing the spectral information on the highest resolution. Figure 4 shows how β influences two extracted features from the enlarged the top right part of Figure 1b.

3.3. Correlated geometric shapes

Each extracted iso-surface represents the spectral and spatial distribution of elements or molecules in the datacube. These iso-surfaces can be visualized as different geometric shapes, but PCA enables us to add even more information to these shapes. Information about correlation between peaks and regions from the PCA could be used in the visualization of 3D shapes. A fractional weight function is derived for each spectral window emphasizes on the contribution of that element relative to the contribution of other elements in one principal component. The fractional weight function for a positive correlated shape is derived using Y_α^+ from Equation 3. Using Equation 2 the spatial contribution of the positive part thresholded by α can be calculated as shown in Equation 9. This can be done in a similar fashion for the negative spatial contribution I_α^-.

$$I_\alpha^+ = X^T \cdot Y_\alpha^+ \tag{9}$$

Now the fractional spatial weight w^{m+} for each positive correlated peak λ can be calculated with Equation 10.

$$w^{m+} = \frac{\Sigma I_\alpha^+ - \Sigma I_\alpha^-}{i_\alpha^{m+}} \tag{10}$$

This weight should be used on the convolved *raw* data from one bin m in order to highlight its specific contribution to all the selected elements in one component. Similar weights can be derived for each negative correlated peak. There are areas in a feature with a high density of intensity values and areas with lower densities. The third parameter γ can now be set to show iso-surfaces on different values of the selected features in the datacube in a 3D space. Each shape can have a different γ which enables visualization of the features according to different 3D density distributions as shown in Figure 5.

4. Results

The data used in this example was measured using a time-of-flight secondary ion mass spectrometer (ToF-SIMS). The sample is a thin cross-section of a chicken embryo. The cross-section is $8 \times 8mm$ in size and contains a spectral mass window from $\sim 1 - 2000m/z$. The spatial dimensions of a dataset can be 512×512 or higher and the spectral dimension can have $\sim 2 \cdot 10^6$ intensity values. The Matlab environment is used for the implementation of this example.

All spectral intensities are summed in Figure 1b which makes it impossible to distinguish between different values in a spectral profile and their corresponding specific spatial contribution. Interesting features like the heart, blood vessels, bone structures or the distribution of cholesterol remain hidden or poorly visible in these representations. It is hard to make the distinction between the cross-section itself and the material in which it is embedded. In our approach, α is used to reduce the amount of spectral noise in the selection of spectral windows. The second parameter β enables a user to view the resulting features on different levels of detail. The highest level of detail shows the original cloud of points from one particular spectral window, but the iso-surface of

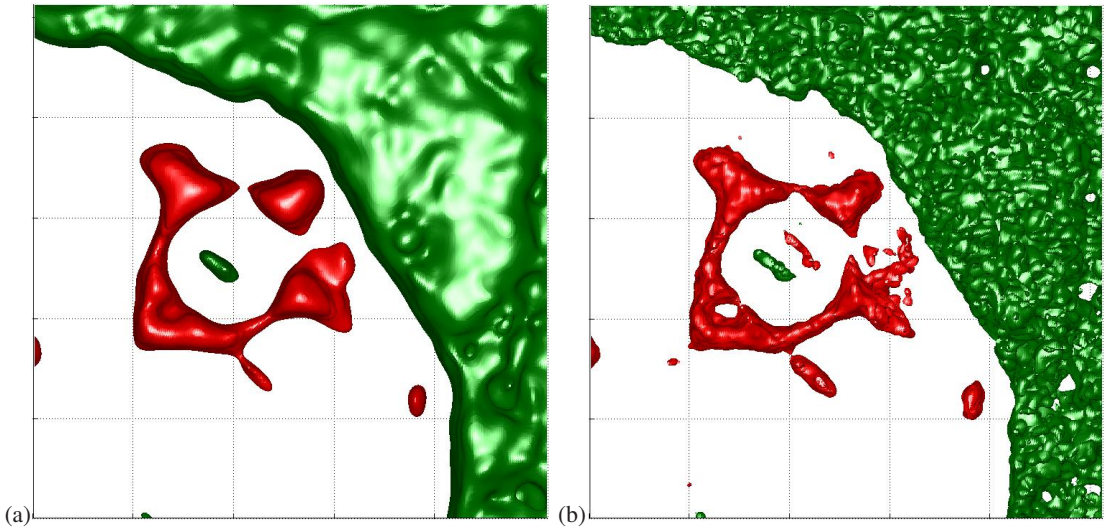

Figure 4: *the iso-surfaces of two anti-correlated features within the second principal component with (a)* $\beta = 16$ *and (b)* $\beta = 32$

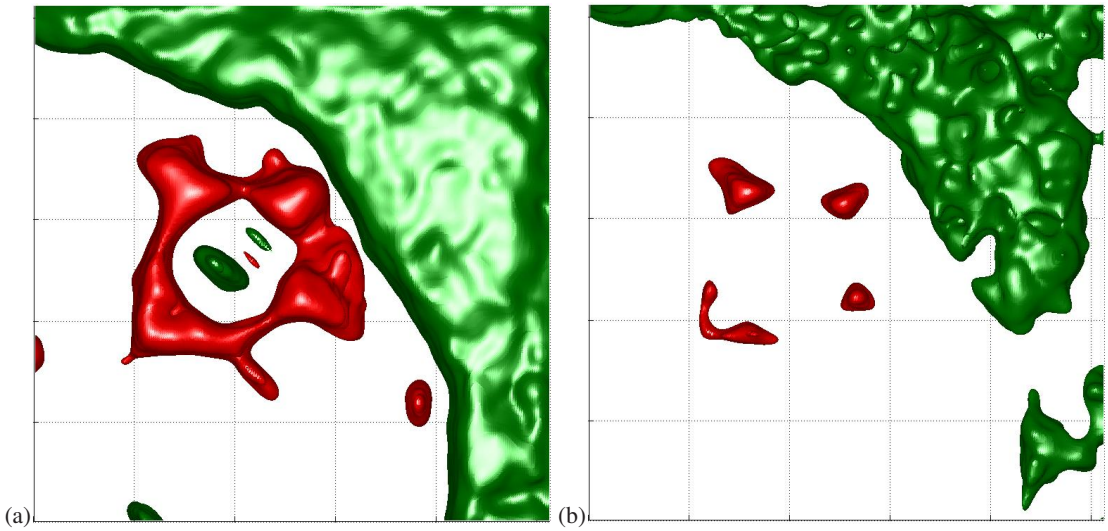

Figure 5: *the iso-surfaces of two anti-correlated features within the second principal component with (a) the red shape on* $\gamma = 0.4$*, a green shape on* $\gamma = 0.2$ *and (b) the red shape on* $\gamma = 0.6$*, a green shape on* $\gamma = 0.4$

the unconvolved data does not reveal clear coherent information in the cloud. The information of the spectral structure becomes more apparent when smaller values are chosen for β. Figure 4 shows the iso-surfaces of two anti-correlated features within in the second principal component. Both show the same part of the cross-section of the backbone with the red shape representing sodium and the green shape representing indium. Clearly, more structural details can be seen in the right view compared to the one on the left. For instance small red regions appear beside the backbone that could be identified as blood vessels on a higher level of detail. The neural tube represented by the hole on the left-bottom part of the cross-section of the backbone is visible in the image on the right but is closed on the left. Different values for β can be used to find a balance in the complexity of the structure of the iso-surfaces and the desired level of detail. The size of the extracted shapes can be controlled with the third parameter γ. In Figure 5 is shown that different values for γ can be used to find an appropriate density on which the component is being displayed. Those regions that have the highest data density are selected with higher values for γ.

A final visualization of the cross-section can be made with

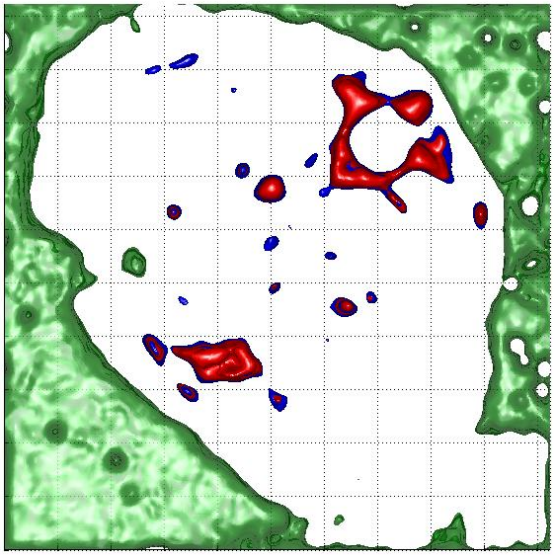

Figure 6: *the extracted correlated shapes in the second principal component component with* $\alpha = 0.3$, $\beta = 16$ *and* $\gamma = 0.35$ *for the blue and red shape,* $\gamma = 0.15$ *for the green shape*

all three parameters in Figure 6. In this example we used only the second principal component as it creates a clear distinction between bone tissue and the material in which the cross-section is embedded. Again, the red and blue shapes are elements that are correlated and the green represents the anti-correlated material outside the embryo. The holes in the green shape contain the fragments of other elements that are deselected and the irregularities on the green surface are due to noise artifacts in the sample itself. An expert is able to interpret the distribution of elements this visualization. For instance the blue element (potassium) shows a similar distribution as the red element which represents sodium. Both are present in the bone-tissue and blood so the large red shape on top of the figure can be identified as the cross-section of the backbone and the large red shape on the bottom can be identified as the heart. Different principal components can be used to create multiple views of the distribution of correlated features within the same datacube. For instance if other components contain elements or molecules present in the heart and not in the bone (or vice versa) they can be classified and separated as different types of tissue.

5. Discussion and future work

The proposed method of parametric visualization of high-resolution correlated features has a number of advantages compared with the classical method of manual exploration. First, the extracted correlated and anti-correlated patterns are made distinctive through different colors in this new visualization. The threshold α is used to reduce the spectral noise present in one feature. Second, the full spectral resolution of the dataset is utilized instead of a spectrally binned one, what lacks in other tools for visualization. Third, our method will display the extracted features with better defined boundaries extracted in three dimensions instead of a traditional 2D view of the summed image. Therefore more spectral noise is left out in our shape extraction. The shapes contain only those 3D regions with the highest density in contrast with the low-density regions with a lower signal to noise ratio. These geometric shapes can be used to conveniently select one or more spectral windows or even just a spatial region of interest. Finally our visualization is parametrically controlled in such a way that an analyst is in control of the feature extraction and can set a desired level of detail.

Some problems remain present in this approach. We chose PCA for the feature extraction because it has already proven itself in this field of application. Other methods for decomposition can be used as well, but due to the enormous sizes of the datacubes and distribution of peaks among multiple spectral levels it is not yet possible to apply the algorithm on a full resolution dataset. Part of the problem is that the results in our visualization still depend on the effectiveness of the PCA. However, even with this limitation, our method is still able to create a better representation of the distribution of mass spectral components than examining the results of PCA on a traditional way. The problem of not being able to give an intuitive interpretation to the negative scores that result from PCA ([SOKK04]) is solved by our method. The presence of negative scores even contributes to our visualization because they enable the display of anti-correlated features. It is also easy to experiment with different 3D filtering techniques, for instance anisotropic convolution that may provide less smoothed boundaries in the extracted shapes. Eventually we would like to add functionality to this parametric visualization method that is able to select the most appropriate values for the three parameters we introduced. These values have to be independent from the methods for decomposition, use of different convolution kernels or most important, different datasets.

6. Acknowledgements

We would like to thank Els Bon from the FOM Institute for Atomic and Molecular Physics (AMOLF) who prepared and supplied the spectral datasets that were used for the given examples.

This work was carried out in the context of the Virtual Laboratory for e-Science project (`www.vl-e.nl`). This project is supported by a BSIK grant from the Dutch Ministry of Education, Culture and Science (OC&W) and is part of the ICT innovation program of the Ministry of Economic Affairs (EZ).

References

[BvL05] BROERSEN A., VAN LIERE R.: Transfer functions for imaging spectroscopy data using principle component analysis. In *Eurographics / IEEE VGTC Symposium on Visualization* (June 2005), Brodlie K., Duke D., Joy K., (Eds.), pp. 117–123.

[GSvdW02] GEUSEBROEK J. M., SMEULDERS A. W. M., VAN DE WEIJER J.: Fast anisotropic gauss filtering. In *European Conference on Computer Vision, Copenhagen, Denmark, Part I* (May 2002), vol. 2350, Springer Berlin / Heidelberg, pp. 99–112.

[HKR*97] HAIGH S., KENNY P. G., ROBERTS R. H., BARKSHIRE I. R., PRUTTON M., SKINNER D. K., PEARSON P., STRIBLEY K.: Automatic and interactive correlation partitioning compared : Application to tin/ ti/ sio. *Surface and Interface Analysis 25* (1997), 335–340.

[JG05] JACOBSON N. P., GUPTA M. R.: Design goals and solutions for display of hyperspectral images. In *IEEE Transactions on Geoscience and Remote Sensing* (2005), vol. 43, pp. 2684–2692.

[Jol02] JOLLIFFE I. T.: *Principal Component Analysis*, second ed. Springer series in statistics. Springer-Verlag, 2002.

[KBF*07] KLERK L. A., BROERSEN A., FLETCHER I. W., VAN LIERE R., HEEREN R. M. A.: Extended data analysis strategies for high resolution imaging ms:new methods to deal with extremely large image hyperspectral datasets. *International Journal of Mass Spectrometry 260*, 2-3 (Feb. 2007), 222–236.

[Kee05] KEENAN M. R.: Maximum likelihood principal component analysis of time-of-flight secondary ion mass spectrometry spectral images. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films 23*, 4 (July 2005), 746–750.

[KK04] KEENAN M. R., KOTULA P. G.: Optimal scaling of tof-sims spectrum-images prior to multivariate statistical analysis. *Applied Surface Science 231-232* (2004), 240–244.

[KNM*97] KENNY P. G., NEWBURY P. F., MOUNTAIN D. L., WHITEHOUSE D., HAIGH S. A., PRUTTON M., ROBERTS R. H., BARKSHIRE I. R., WENHAM M. J. G.: Compression, visualization and segmentation techniques for 3d spectrum-images from multispectral analytical electron microscopy. In *International Conference on Imaging Science, Systems and Technology* (1997), Arabnia H. R., (Ed.), CSREA, pp. 355–363.

[Koe84] KOENDERINK J. J.: The structure of images. *Biological Cybernetics 50*, 5 (Aug. 1984), 363–370.

[MSSK05] MCCOMBIE G., STAAB D., STOECKLI M., KNOCHENMUSS R.: Spatial and spectral correlations in maldi mass spectrometry images by clustering and multivariate analysis. *Analytical chemistry 77*, 19 (Oct. 2005), 6118–6124.

[Pac04] PACHOLSKI M. L.: Principal component analysis of tof-sims spectra, images and depth profiles: an industrial perspective. *Applied Surface Science 231-232* (2004), 235–239.

[PvdH01] POLDER G., VAN DER HEIJDEN G. W. A. M.: Visualization of spectral images. In *Visualization and Optimization Techniques* (2001), Censor Y., Ding M., (Eds.), vol. 4553, pp. 132–137.

[SOKK04] SMENTKOWSKIA V. S., OHLHAUSENB J. A., KOTULAB P. G., KEENAN M.: Multivariate statistical analysis of time-of-flight secondary ion mass spectrometry images: looking beyond the obvious. *Applied Surface Science 231/232* (2004), 245–249.

[Wit83] WITKIN A. P.: Scale-space filtering. In *International Joint Conference on Artificial Intelligence* (1983), pp. 1019–1021.

[WKC03] WICKES B. T., KIM Y., CASTNE D. G.: Denoising and multivariate analysis of time-of-flight sims images. *Surface and Interface Analysis 35* (2003), 640–648.

[WSH99] WOLKENSTEIN M., STUBBINGS T., HUTTER H.: Robust automated three-dimensional segmentation of secondary ion mass spectrometry image sets. *Fresenius' Journal of Analytical Chemistry 365*, 1-3 (Sept. 1999), 63–69.