Vision, Modeling, and Visualization (2010)

An Interactive, Multi-Modal Approach to Analysing High-Resolution Image Mass Spectrometry Data

F.A. Smit[1] and R. van Liere[1] and L. Fornai[2] and R. Heeren[2]

[1]CWI Amsterdam, The Netherlands
[2]AMOLF, Amsterdam, The Netherlands

Abstract

The output resolution of imaging mass spectrometers is increasing rapidly due to advances in engineering and the use of tiling. Imaging-MS data is often displayed as a total-ion-count (TIC) image; however, anatomical structures are not easily identifiable from TIC images. For this purpose, additional high-resolution images that originate from different imaging modalities, such as stained histological data, are preferred.
These modalities are most useful when fused; i.e., when the corresponding images are spatially aligned with respect to each other. The viewing and analysis of such data is ideally performed in real-time and at the highest possible resolution, allowing users to interactively query the combination of all fused data at the highest detail. However, proper alignment between modalities and interactively presenting large volumes of data is as of yet a challenge.
We present a system for the simultaneous viewing and analysis of high-resolution data from different imaging modalities. Fusion is provided in such a way that interaction in one modality can be mapped to different modalities. For example, anatomical structures can be identified from histological data and their spatial extent mapped to a corresponding region-of-interest in the image MS data, allowing the analysis of its chemical compounds. In turn, the MS data can be analysed and filtered, for example using multi-variate analysis such as PCA, and the result mapped back to structures in other modalities. Level-of-detail, region-of-interest and asynchronous data processing algorithms ensure that the system can be operated interactively at the highest resolution.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Methodology and Techniques I.3.8 [Computer Graphics]: Applications

1. Introduction

The output of a mass spectrometer consists of a set of mass-to-charge ratios (m/z) of detected ions. A histogram of this data is used to identify the elemental composition of the sample. Image mass spectrometry produces such a set of mass-to-charge ratios for each pixel in a 2D grid, allowing one to analyse the chemical structure of the sample. The resulting grid of m/z histograms is commonly visualized in the form of a total-ion-count (TIC) image. For each pixel, a TIC image maps the number of items in the corresponding set of m/z values to an intensity value. Other methods of visualization are based on various forms of multi-variate analysis (MVA) on the m/z data, producing maps of chemically similar regions.

Pathologists often want to analyse known anatomical structures in a sample; however, these can not easily be identified from MS data alone. While TIC images give a rough overview of the data, the only information conveyed is the number of detected ions, which vary considerably due to detector sensitivity and material hardness; therefore, it is hard to reliably identify anatomical structures from TIC images. An MVA analysis allows one to identify chemical but not anatomical structure. Therefore, different imaging modalities for the same sample, such as stained histology, are preferred to identify anatomical structures in the data. For these different modalities to be most useful, they should be aligned spatially.

The output resolution of image mass spectrometers is in-

© The Eurographics Association 2010.

DOI: 10.2312/PE/VMV/VMV10/283-289

delivered by EUROGRAPHICS DIGITAL LIBRARY www.eg.org diglib.eg.org

creasing rapidly due to advances in engineering and the use of tiling. For example, the MS dataset that is used in this paper consists of 128×128 tiles at a resolution of 256×256 pixels with a spectral resolution of 22 bit mass channels, for a total of about one billion pixels. Since the sample size is approximately 2cm×2cm, the resulting image fidelity of such TOFSIMS experiments is close to 0.6 micrometer per pixel. Similarly, the output of other imaging modalities often exceeds the one gigapixel mark. This high image fidelity is required to identify small molecular structures in the order of micrometers, which would not be visible at lower resolutions. Therefore, to avoid a significant loss of spatial information, it is essential that all the data be viewed and analysed at the highest possible resolution.

Furthermore, for effective exploration of the data, an interactive system is necessary so that the user can interactively query the combination of all fused data at the highest detail. It should be noted that at these high data resolutions, it is no longer computationally feasible to perform classic multivariate analysis on the entire data set. For the system to remain interactive, different methods of analysis with reduced computational cost need to be developed.

The combination of these observations leads to the following requirements from a user's point of view for a system to view and analyse MS data:

- Operating at the highest resolution possible
- Interactive viewing and analysis
- Support for multiple fused modalities

We present an interactive multi-modal system satisfying these requirements. Fusion between modalities is provided in such a way that interaction in one modality can be mapped to different modalities. A region-of-interest (ROI) can be selected using brushing techniques and used in different modalities. For example, anatomical structures can be identified from histological data and their spatial extent mapped to a corresponding ROI in the image MS data, allowing the analysis of its chemical compounds. In turn, the MS data can be analysed and filtered, for example using multi-variate analysis such as PCA, and the result mapped back to structures in other modalities. Level-of-detail, ROI and asynchronous data processing algorithms ensure that the system can be operated interactively at the highest resolution. A sample of the system's interface is shown in Figure 1.

2. Related Work

Many different methods of analysing MS data have been proposed (e.g. [BGC03, GWC06]). Most of these methods are based on multi-variate analysis, such as principle component analysis (PCA). A disadvantage of these methods is that they are computationally expensive, especially for modern high-resolution MS datasets exceeding a billion pixels. Furthermore, MVA analysis does not directly take into account the spatial configuration of pixels. All m/z values are treated equally in statistical fashion, irrespective of their position in the MS image. Furthermore, it is difficult to reliably identify anatomical structure from MS data and its analysis.

The need to identify anatomical structure alongside chemical molecular structure was previously noted by McDonnel et al. [MvARvZD08]. They state that one of the principal motivations for spatially resolved mass spectrometry is the ability to correlate anatomical information provided by histology with biochemical information provided by imaging mass spectrometry. In this way, anatomical changes associated with a pathology and corresponding chemical molecular changes can be identified, advancing the understanding of such processes.

Fornai et al. constructed a high-resolution molecular atlas of successive tissue section of a rat heart [FKA*10]. They state that since the rat heart exhibits physiological characteristic similar to those of humans, it is a key experimental model in biomedicine. A future goal is to identify changing tissue regions that are indicative of human heart disease, and thus correlate the anatomy of the rat heart with that of a human. Therefore, this dataset is an ideal candidate for the testing and development of the multi-modal system proposed in this paper, which fuses anatomical data from other modalities with chemical molecular data from MS imaging.

Interactively visualizing the large volumes of data originating from imaging mass spectrometers and microscopes is still a challenging problem. With datsets consisting of several billion pixel layers, the problem is closely related to the visualization for geographic information systems (GIS). Similar tiling and level-of-detail techniques have been used here to implement a fully interactive system.

3. Methods

3.1. Basic system design

To meet the requirements proposed in Section 1 – interactive, multi-modal and high-resolution – the technical system design follows a number of guidelines. First, since the viewing and analysis should be performed at the highest resolution for a multitude of potentially ultra-high resolution images, the system should never attempt to load an entire data model at the same time. To this end, a model-view architecture is used where the view determines which parts of the model are loaded. The model is responsible for loading data from various modalities and caching it in memory as efficiently as possible. Second, to support many different modalities, as few assumptions as possible are made about the type of data and how to display it, resulting in a system that is easily extensible to new modalities. Third, since different modalities have different resolutions, the system should not be constrained by a single, uniform notion of pixel size.

The view on the data can be translated and zoomed freely,

© The Eurographics Association 2010.

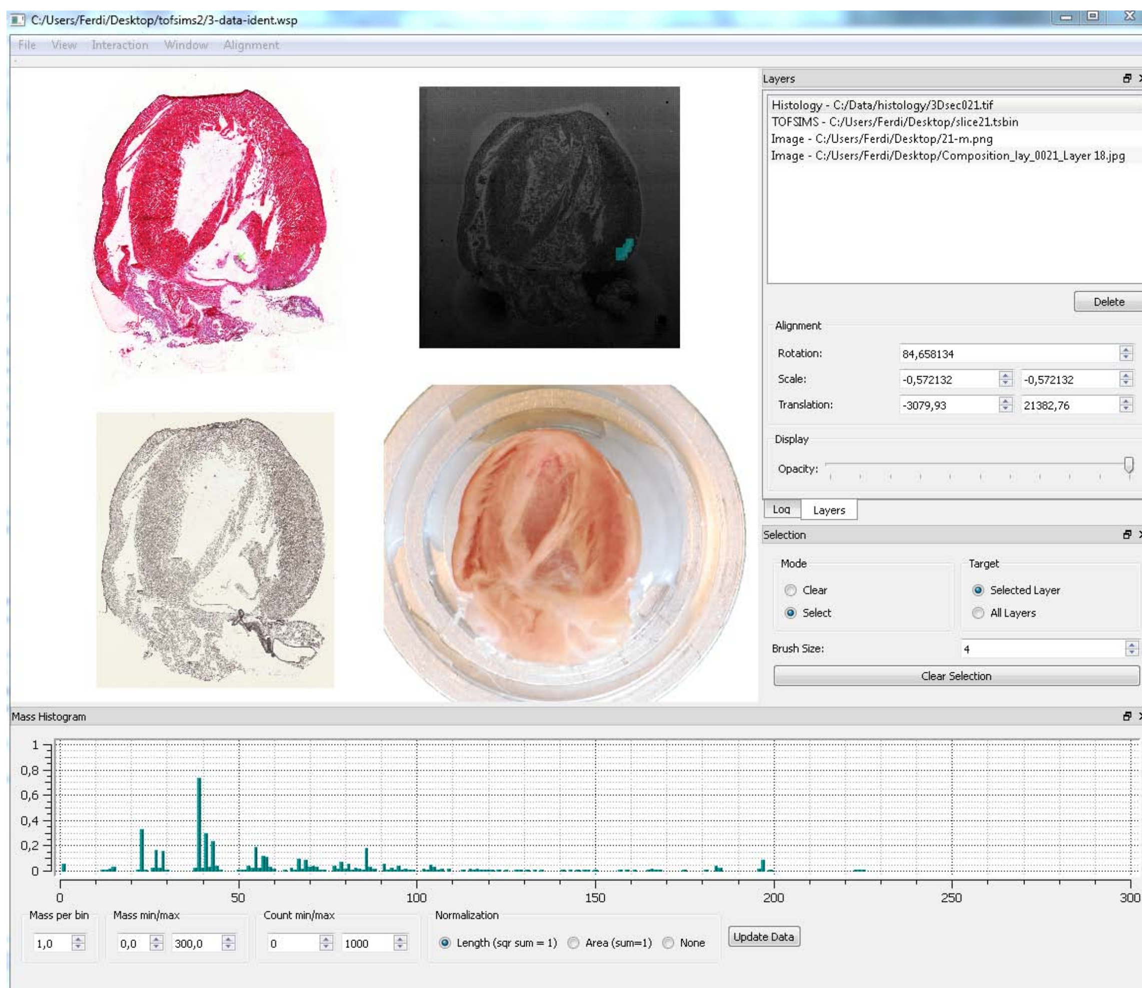

Figure 1: *Sample overview of the presented system. Four different imaging modalities have been loaded into layers and are displayed side-by-side. From left-to-right, top-to-bottom the images represent stained histology, image TOFSIMS, microscopy and a regular photograph of the same mouse-heart sample. At the bottom an m/z histogram is displayed for a selected ROI. The MS data resolution is sub-sampled to 8k×8k for an image fidelity of about 2.5 micrometer per pixel. The histological data has a resolution of 20k×20k with an image fidelity of 1 micrometer per pixel. Such high resolutions are essential for the identification of small chemical structures.*

which is realized by dragging with the mouse for translation and using the mouse's scroll wheel to zoom. Finally, rotation is achieved by dragging the mouse and pressing the right mouse button. Zooming and rotation are performed with respect to the current cursor position, allowing one to intuitively zoom in on any part of the data. The data are organized in a set of 2D blocks, and the view determines which parts of the data are visible and the size of their projected area. Subsequently, the view requests the model to load these blocks of data from disk. This loading is performed asynchronously in a separate thread, so the view thread does not become blocked after a request to load data. In the meantime, while the data is being loaded, the view uses a low-resolution approximation of the data until the actual data becomes available. This ensure an interactive system, similar to the mode of operation in Google Maps.

Currently three data formats are supported by our system: regular image files in various formats, high-resolution tiled TIFF images and mass spectrometry data. Regular images are assumed to be of sufficiently low resolution that they can be loaded in their entirety and require no special handling. Tiled TIFF is a popular output format for ultra-high resolution images, for example originating from microscopes. This format allows the loading of individual blocks of data as re-

© The Eurographics Association 2010.

quested by the view. For the MS data we make use of a special, in-house binary format. For every pixel, the number of ion events for that pixel is stored. Following this data is a list of all m/z values corresponding to ion events. The pixels are stored in a space-filling curve according to their Morton- or Z-order. The advantage of using Z-order for storage is that scaling can be achieved efficiently by adding up consecutive values on disk. The file also contains an index block that points to the starting positions of blocks of data, so these can be individually loaded.

The view determines which data are visible, but the actual rendering is off-loaded to different display objects. These objects make use of OpenGL and GLSL to render the data. A texture format is acquired from the data model, after which custom, programmable shaders are used to visualize the data. This allows many different types of data to be displayed in a custom manner.

3.2. Features and system operation

The presented system consist of a number of key features, which will be described next. All of these are based on the basic system design as described in the previous section. The view on the data is layer-based, meaning that different modalities can be loaded into different layers, which can then be displayed either side-by-side or in overlay mode with configurable opacity. When the camera is moved, all viewports are updated accordingly simultaneously. Interaction, such as transformation and selection, can either be performed on all layers or on a selected layer of interest. Side-by-side viewing is especially useful for selection, where corresponding regions in both modalities are visible at the same time. Viewing in opacity overlay mode is useful for checking the alignment between modalities. This also allows one to manually align different modalities.

Semi-automatic alignment is also provided in the form of having to manually select corresponding feature points in two modalities. The user achieves this by simply double-clicking the mouse on corresponding points in different viewports using side-by-side viewing. The system then determines the transformation that best maps one cloud of feature points into the other and applies this transformation to the layer, effectively aligning the modalities. There are several ways to calculate this transform, but the method we have implemented is based on the Amoeba algorithm for the unconstrained optimization of a non-linear function [NM65]. Other methods, for example ones to solve the mean-centred, orthogonal Procrustes problem [Eve97], could be used as well.

Selection is performed by means of interactive brushing; the user can drag a virtual paintbrush to select a region-of-interest for analysis. This selection is either performed on the currently selected layer or on all layers at the same time. Especially the latter mode of operation is useful in side-by-side viewing to interact with the paintbrush in one modality and automatically select a corresponding ROI in another, previously aligned modality. On the user's request, the m/z data for the selected ROI in MS data is fetched from disk by the data model and then displayed in a histogram of m/z ranges. The histogram either displays raw ion-event counts per bin or, alternatively, the histogram can be normalized in such a way that the area under the histogram or the length of the histogram vector equals one. Length normalization can be used to subsequently compare m/z spectra by means of an inner product [Alf04]. This allows the identification of chemical compounds in the ROI.

Finally, a simple filtering mechanism on m/z ranges is provided. The user can select a range of m/z values, and for each pixel in the ROI, the percentage of m/z values in the pixel's histogram that fall within this range is mapped to an intensity value. This is equivalent to integrating the are normalized histogram over a user-selected domain. The resulting intensity map is overlaid on the data and gives a quick spatial overview of the concentration of certain chemical compounds corresponding to the selected m/z range. Extensions to this filtering mechanism are described at the end of Section 5.

4. Results

For system testing, we used an existing high-resolution data set of a flash-frozen rat heart [FKA*10]. The dataset consist of several 12 µm thick tissue sections that were cut with a cryomicrotome; however, we only used a single sample slice in this paper. The samples' dimensions are 2cm×2cm. The MS data was produced by a Physical Electronics (Eden Praire, MN) TRIFT-II time-of-flight SIMS (TOF-SIMS) using Au+ primary ions at a resolution of 32k×32k pixels. To increase the ion counts per pixel of the MS data to acceptable levels, the data was subsampled on-the-fly to 8k×8k by the model loader, resulting in a spatial resolution of approximately 2.5 micrometer per pixel. The spectral resolution for the m/z values is 22 bits. Haematoxylin and Eosin stained histological data with a resolution of 20k×20k pixels (1 micrometer per pixel) was also acquired for these sections after TOFSIMS analysis. Finally, 2346×2720 resolution optical images in JPEG format as well as low-resolution photos of the frozen sample are available. Full overview of the histology and MS data are shown in Figure 2.

From Figure 2 it can be seen that it is difficult to identify anatomical features from MS data. The problem becomes especially apparent when zoomed in at higher resolutions, as is shown in Figure 3. This part of the data is a zoomed-in version of the outline shown in Figure 2. Features such as cells can be identified in the histology data but not in the MS data. Without a multi-modal, aligned viewing system, it would be very difficult to determine cell distribution from the MS data alone.

As a proof of concept, we describe a sample operation of

© The Eurographics Association 2010.

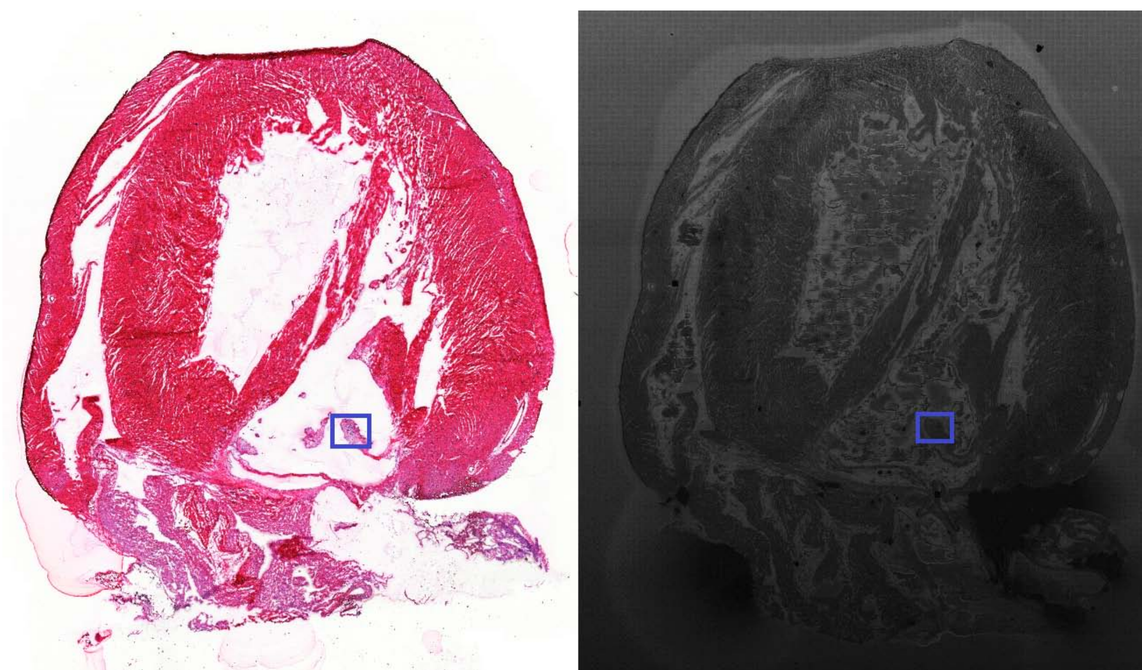

Figure 2: *A full overview of the used dataset. Stained histology data is shown on the left, while a TIC image of the MS data is shown on the right. It is difficult to reliably identify anatomical features in the MS data. The non-uniform distribution of ion counts is also apparent from the figure, which darkens towards the bottom of the image. The blue inset is shown magnified in Figure 3.*

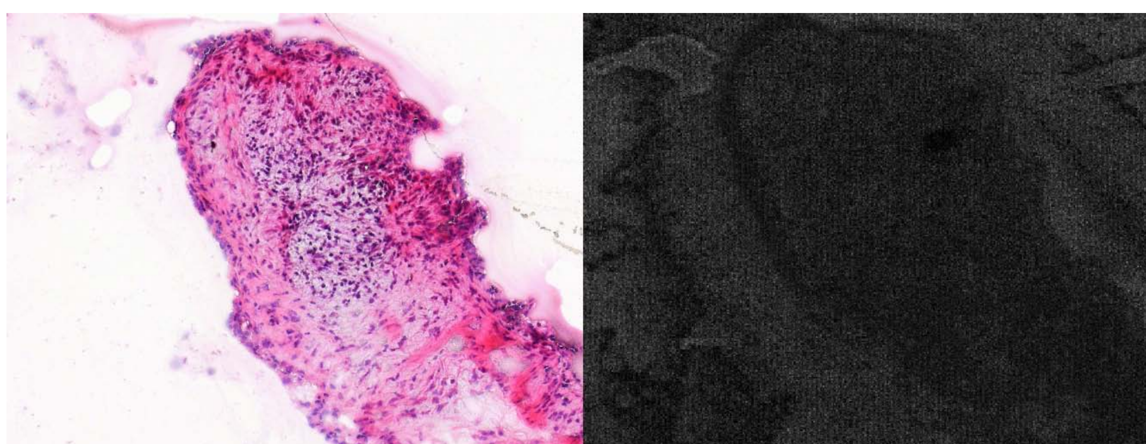

Figure 3: *A close-up view of the boxed inset from Figure 2. Stained histology data is shown on the left, while spatially corresponding TOFSIMS data is shown on the right in the form of a TIC image. The histology data clearly shows anatomical structure that would be hard to identify in the TIC image.*

© The Eurographics Association 2010.

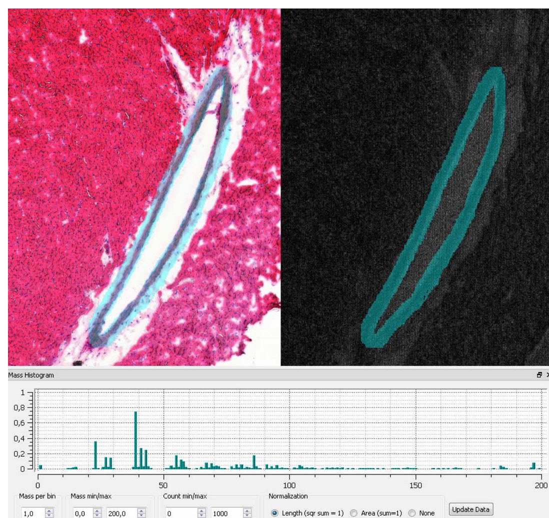

Figure 4: *Stained histology and TOFSIMS data are aligned and displayed at full resolution. Brushing has been used to select the artery wall in the histology data and is automatically mapped to the TOFSIMS data. The histogram at the bottom shows the m/z data for this selected ROI.*

selecting an artery wall and examining its chemical composition. The process is depicted in Figure 4. First, the artery is identified from the histological data and viewed at full resolution. Next, the histological and MS data are aligned using the earlier described methods. A side-by-side view is chosen to immediately see the effect of selection on both layers. The user now selects the artery wall by brushing over it in the histological data, causing it to be selected in both imaging modalities. Finally, the mass histogram is updated to match the selected ROI in the MS data. The whole process is performed interactively at maximum resolution with multiple layers loaded for a total viewed data size of over 7 gigabytes. When rapidly zooming in and out, the model loader occasionally took up to a second to load the high-resolution data, but overall no delays were noticeable during operation. Also, because of the level-of-detail processing, a lower resolution view of the data was shown during loading periods, so no perceptually disturbing artefacts were visible.

The display components for the images simply displayed texture data as RGBA images. For the MS data we used a component that displays single-channel TIC images. The number of ion-events per pixel was loaded from disk and mapped to a black and white intensity value. We noticed a large deviation in ion-counts across the image. Detector sensitivity proved not to be uniform within single blocks and also changed between blocks, causing an uneven distribution of ion-counts. This can be seen in Figure 2, where the intensity of the TIC image gradually degrades towards the bottom of the image. The edges of the individual blocks can also be distinguished. This shows once more that TIC images are not the best candidate to identify anatomical features in the data.

The alignment process proved to be non-trivial due to deformations in the samples for different modalities. We found it best to align the different modalities specifically for a smaller region of interest, resulting in high-quality alignments. Global alignment is still useful for a quick overview of the data, but when selecting detailed features a high-accuracy approach is required. Nonetheless, with some manual adjustment it was possible to align features quite well.

We have demonstrated the system to a number of physicists and pathologists in the department. Their first reaction was that they were impressed with the high resolution views of the data and the interactive speeds. Previously, they had mostly worked with 256×256 pixel output images of PCA analysis (78 micrometer per pixel), where detailed features of the data were impossible to see. They claimed that even though they knew the MS output data was of high resolution, they lacked the software to analyse and view the data at these resolutions. The interactivity and high detail of the system promotes exploration of the data. Finally, the pathologists were interested in the anatomy-based workflow; often, they wish to determine chemical properties of specific, known anatomical structures and are less interested in a global distribution of chemical compounds as found by MVA analysis of the MS data.

5. Discussion

While the presented system succeeds in meeting the requirements described in Section 1, there is much room for future improvement. First, the alignment between modalities is currently a manual process; the user either directly aligns the images or manually selects a set of points to be aligned. Ideally, this alignment should be performed automatically by the system. One possible way of achieving this is by using edge detection on both modalities and aligning the found edges automatically. Alternatively, feature points could be identified automatically and then be aligned. A problem with this and the edge detection approach is that it is not immediately obvious what constitutes a feature in different modalities. For example, MVA analysis can be used to automatically identify chemical feature points in the MS data; however, finding corresponding features in different types of data, such as stained histology, which only contain color data is a challenge. The problem is to determine which chemical features map to which optical features. Most likely, this process can never be performed fully automatically, but perhaps with some user guidance it can be achieved. Such a semi-automatic alignment algorithm would significantly ease the process of alignment.

The image data from all modalities is currently two dimensional. However, several images from consecutive slices

© The Eurographics Association 2010.

through a 3D object are available. It would be desirable to load such a stack of 2D slices and transform them into a 3D model. This transformation into 3D requires alignment between different slices. To perform this alignment automatically, similar considerations as described in the previous paragraph apply. An additional difficulty is that the slices are now subtly different from each other, and feature points found in one slice may not exist or be deformed in the next. This may be solved by placing physical alignment markers into the sample, but this does not help for data already produced without these markers. Furthermore, even with markers in the sample, deformations still pose a problem. Algorithms designed to perform alignment between modalities may be equally useful for constructing 3D models.

With multiple, ultra-high resolution images, the lack of sufficient display resolution and surface becomes an issue. The user interface should be designed in such a way that display space is not an overly limiting factor. To this end, the use of multi-display systems and multi-resolution system may prove beneficial. Software solutions to maximize the benefit of the available display space should also be explored.

Due to the introduction of tiling image mass spectrometers and further engineering advances, MS data resolutions are expected to increase still further over the upcoming years. Pathologists would ideally be able to discern individual cells and even parts of cells on a sub-micrometer scale. To achieve this, besides higher output resolution from mass spectrometers, high-resolution visualization and analysis of the data is essential. Low-resolution approximations for the sake of increased computing efficiency – such as 256×256 pixel PCA analysis – is not sufficient for this purpose. Since it is unlikely that such large datasets can be analysed interactively in their entirety, it is important to develop localized analysis methods that operate on regions-of-interest and take level-of-detail considerations into account.

We have used only a simple form of filtering for the MS data, but the system is easily extensible to different filtering and analysis methods. MS data is usually analysed using forms of MVA, such as principle component analysis (PCA). While these methods succeed in identifying similarities in the data, the analysis itself ignores all spatial information. Different methods that use the available spatial information may result in better analysis of the data. Another approach is to use custom domain filters over the m/z range, as opposed to a simple uniform domain of integration. In this way, the normalized histogram is convolved with a predefined histogram of interest, resulting in an intensity map that depicts how well the pixels' histograms match a template histogram. In this way, the filtering method is no longer restricted to identifying individual peaks in the m/z data, but can be used to determine the goodness-of-fit to complex histograms. Such methods can also be used for finding feature points for alignment.

6. Conclusions

We presented an interactive, multi-modal system for the viewing and analysis of MS data that operates at the highest possible resolution. The unique aspects of this system are that it manages to operate at very high resolutions while maintaining full interactivity and that the analysis of the data is linked between modalities. This was achieved using a level-of-detail, tiling and region-of-interest selection to minimize the amount of data to be loaded and processed. Data was loaded and processed asynchronously to avoid breaking the interactivity of the system. The multi-modal nature allows one to select anatomical features from different modalities and analyse their chemical structure from the corresponding MS data. These anatomical features would be hard to identify using MS data alone. Finally, the physicists and pathologists to whom we have demonstrated the system were impressed with the high resolution and speed of interaction.

References

[Alf04] ALFASSI Z.: On the normalization of a mass spectrum for comparison of two spectra. *Journal of American Society of Mass Spectrometry* (2004), 385–387. 4

[BGC03] BELU A. M., GRAHAM D. J., CASTNER D. G.: Time-of-flight secondary ion mass spectrometry: techniques and applications for the characterization of biomaterial surfaces. *Biomaterials 24*, 21 (2003), 3635–53. 2

[Eve97] EVERSON R.: Orthogonal but not orthonormal procrustes problems. In *Advances in Computational Mathematics* (1997). 4

[FKA*10] FORNAI L., KLINKERT I., ANGELINI A., GISKES F., KLERK L. A., FEDRIGO M., THIENE G., HEEREN R. M.: 3d imaging mass spectrometry of the heart. *American Society of Mass Spectrometry (ASMS)* (2010). 2, 4

[GWC06] GRAHAM D. J., WAGNERB M. S., CASTNERA D. G.: Information from complexity: Challenges of tof-sims data interpretation. *Applied Surface Science 252*, 19 (2006), 6860 – 6868. 2

[MvARvZD08] MCDONNEL L. A., VAN ALEXANDRA REMOORTERE, VAN ZEIJL R., DEELDER A. M.: Mass spectrometry image correlation: quantifying colocalization. *Journal of Proteome Research 7*, 8 (2008), 3619–27. 2

[NM65] NELDER J. A., MEAD R.: A simplex method for function minimization. *Computer Journal 7* (1965), 308–313. 4

© The Eurographics Association 2010.