

# Creating Graph Abstractions for the Interpretation of Combined Functional and Anatomical Medical Images

Ashnil Kumar<sup>1</sup>, Jinman Kim<sup>1</sup>, Michael Fulham<sup>1,2,3</sup>, and Dagan Feng<sup>1,4</sup>

<sup>1</sup> School of Information Technologies, University of Sydney, Australia

<sup>2</sup> Sydney Medical School, University of Sydney, Australia

<sup>3</sup> Department of Molecular Imaging, Royal Prince Alfred Hospital, Sydney, Australia

<sup>4</sup> Med-X Institute, Shanghai Jiao Tong University, China

{ashnil.kumar, jinman.kim, michael.fulham, dagan.feng}@sydney.edu.au

**Abstract.** The characteristics of the images produced by advanced scanning technologies has led to medical imaging playing a critical role in modern healthcare. The most advanced medical scanners combine different modalities to produce multi-dimensional (3D/4D) complex data that is time-consuming and challenging interpret. The assimilation of these data is further compounded when multiple such images have to be compared, e.g., when assessing a patient’s response to treatment or results from a clinical search engine. Abstract representations that present the important discriminating characteristics of the data have the potential to prioritise the critical information in images and provide a more intuitive overview of the data, thereby increasing productivity when interpreting multiple complex medical images. Such abstractions act as a preview of the overall information and allow humans to decide when detailed inspection is necessary. Graphs are a natural method for abstracting medical images as they can represent the relationships between any pathology and the anatomical structures they affect. In this paper, we present a scheme for creating abstract graph visualisations that facilitate an intuitive comparison of the anatomy-pathology relationships within complex medical images. The properties of our abstractions are derived from the characteristics of regions of interest (ROIs) within the images. We demonstrate how our scheme is used to preview, interpret, and compare the location of tumours within volumetric (3D) functional and anatomical images.

**Key words:** graph abstractions, medical imaging, image interpretation, image comparison

## 1 Introduction

Medical imaging plays an indispensable role in modern healthcare for diagnosis and in the assessment of a patient’s response to treatment. Technological advancements have led to the creation of scanners that combine different imaging

modalities into a single device and are capable of producing high resolution and multi-dimensional (3D/4D) images. The first mainstream device was the combination of positron emission tomography (PET) and computed tomography (CT) to produce a PET-CT scanner that provides volumetric (3D) anatomical (CT) and functional (PET) data, and enables clinicians to visualise the spatial relationships between activity in a tumour with PET and the underlying location from CT [1].

The interpretation of PET-CT images involves the assimilation of information from both modalities simultaneously. This entails traversing the 3D image as an ordered set of 2D slices and mentally reconstructing a spatial understanding of the relationships between the anatomy and any pathology (disease); these relationships are important for accurate diagnosis, for staging cancer, and classifying different conditions [2]. This interpretation process is time consuming since most modern scanners produce hundreds (sometimes thousands) of slices per image volume. Alternatively, the images can be fused into a 3D rendering but this requires several manual image-specific adjustments, e.g., visibility transfer functions [3]. Interpretation is more problematic when clinicians need to interpret and compare multiple image volumes at the same time, e.g., when comparing multiple images to assess a patient’s response to treatment or when analysing the results of a clinical image search engine.

Existing methods for comparing images can be found as part of medical image retrieval engines [4–6]. Similar to Google Image Search<sup>1</sup>, most of these search engines [4, 5] present their retrieved images as a grid. Users then must inspect and compare the pixel information of these images manually to select the image most relevant to their query. However, such an approach is not feasible for volumetric (3D) and multi-modality medical images due to the effort required during interpretation (described above).

Tory and Moller [7] recommended several techniques that could assist humans in interpreting and analysing visualisations. They suggested that enhanced recognition of higher level patterns in complex information could be achieved by creating abstractions from the selective omission and aggregation of the original data. Since graphs are a natural and powerful way of representing relational information [8], we propose that medical images could be visualised using graph abstractions of the complex spatial relationships between pathology and anatomy. In our prior work [6], we investigated this possibility by integrating 2D graph abstractions as part of a medical image retrieval engine. A user study revealed that the users found the abstractions helpful in determining which images were relevant to their query [6]. Users were able to eliminate dissimilar images based on the tumour locations revealed in the abstraction without needing to inspect the irrelevant images in detail but the 2D nature of the visualisation meant that some information was obscured.

In this paper we present a scheme for constructing graph abstractions of relational information derived from medical images. In our method the properties of the graph abstractions are derived from the visual characteristics of the

<sup>1</sup> <http://images.google.com>

images they represent. Our aim is to create an abstraction that will act as a summary or preview of the main content of an image and thus enable users to decide when detailed inspection is necessary, especially in the context of determining image similarity [9]. The critical element is to preserve the spatial layout of the important visual information while simplifying the overall visualisation. We demonstrate our scheme on PET-CT images of patients with lung cancer. PET-CT (as given earlier) is representative of modern complex medical images that stand to benefit most from such abstractions. Our evaluation compares 2D and 3D graph abstractions of medical images. We also compare the different information from abstractions showing large objects (organs and tumours) and smaller key points (landmarks).

## 2 Methods

### 2.1 Scheme for Creating Graph Abstractions

Figure 1 shows the overview of our abstraction scheme. As a first stage we detect regions of interest (ROIs) within the images. These regions can be single pixels (e.g., key points) or they can be a collection of pixels that represent a particular object (e.g., pixels representing an organ). We then extract features from individual ROIs. Spatial relationships are calculated using location of the ROIs and their proximity to others. The ROI features and relationships are used to construct a graph abstraction which is then visualised.

In the following subsections, we describe two realisations of our scheme. The first realisation is an abstraction of the relationships between objects within the images. The second realisation is an abstraction of the relationships between key points within the images. Both realisations follow the same overall process but use different techniques to detect the ROIs, extract the features, and construct the graphs. We currently do not perform any post-processing in our two realisations to allow us to examine the base visualisations so we can determine what optimisations are necessary.

We chose to abstract objects for the first realisation because these objects represent the structures defined in cancer staging literature [2, 10]. However, accurate volumetric segmentation techniques are needed for object delineation. Such techniques are currently not available for all image modalities or objects. For this reason, we chose key points for the second realisation because they are powerful tools for recognising similar structures in different images regardless of scale and orientation transformations [11]. This is useful when comparing medical images due to the naturally occurring variation in patients (e.g., different organ sizes).

### 2.2 Abstraction of Objects

Given that our examples are patients with lung cancer we chose as ROIs the tumours and major anatomical structures above the diaphragm. We extracted

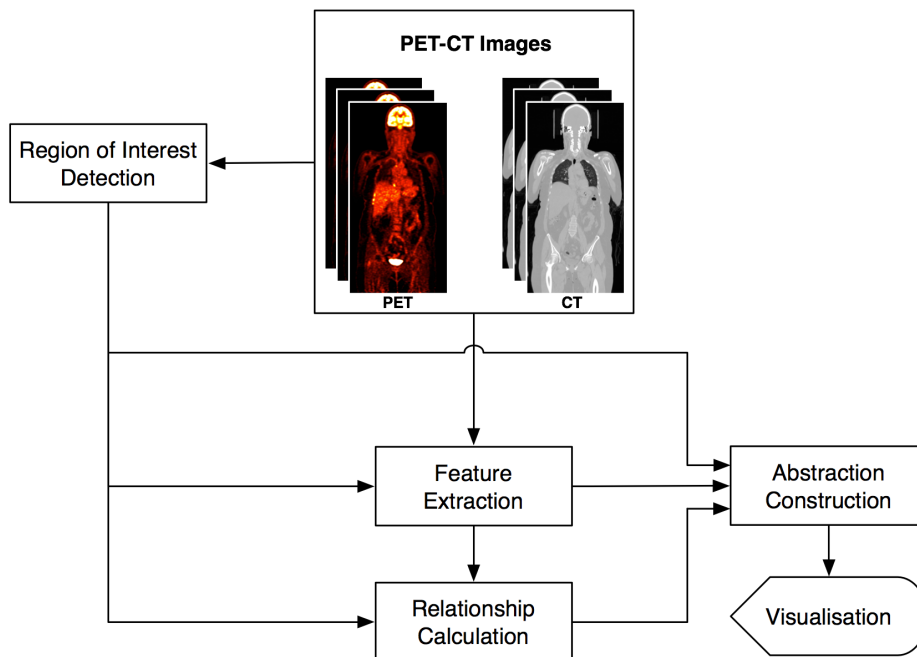


Fig. 1: Overview of the process to create graph abstractions of medical images.

the left and right lungs from the CT images using a well-established adaptive thresholding segmentation algorithm [12]. We also extracted the brain and mediastinum from the CT images using manual connected thresholding. The tumours in the PET images were segmented by first detecting the locations with a local peak radiotracer uptake (high image intensity) and then performing 40% connected thresholding in the neighbourhood of the peaks [13].

We analysed the 3D ROIs and extracted from each the volume (size), centroid (absolute location), and distance to other ROIs. It is possible to extract more features from these ROIs, e.g., as described by Kumar et al. [14]. Here, we list only those features used specifically in this paper for creating our graph abstractions.

The visualisation was created as follows:

1. Each object (segmented ROI) was represented by a single node on the graph.
2. The position of each graph node was derived from the coordinates of the centroid of the ROI.
3. The proximity of the ROIs was used to determine the edge links [14].
4. The size of each graph node was based upon the volume of the ROIs.
5. The colour of each graph node was determined according to the structure it represented (e.g., all tumours were given the same colour).
6. The final position and size of each node (and the lengths of the edge links) were adjusted according to the size of the rendering.

### 2.3 Abstraction of Key Points

We used a 3D equivalent of the Gaussian pyramid method of Lowe [11] to detect key points in the form of Difference-of-Gaussian extrema in each of the image volumes. Each key point was represented by a 3D coordinate, a scale factor, and orientation parameters. We filtered the key points to retain only those CT points that were in proximity to PET key points, and vice versa. Two key points were determined to be in close proximity if the 3D distance between the coordinates of the two points was less than or equal to any of their scale factors. This filtering step eliminated key points that did not contribute to any relationship between tumour and anatomy.

We then extracted scale-invariant feature transform (SIFT) descriptors using a 3D SIFT feature extractor [15] on these key points. We used  $k$ -means clustering separately on the PET and CT descriptors to divide these descriptors into 200 groups (100 for CT and 100 for PET) [16].

The visualisation was created as follows:

1. Each key point was represented by a single node on the graph.
2. The position of each graph node was derived from the key point's 3D coordinates.
3. Two nodes were linked by an edge if they were in close spatial proximity to each other. Proximity was determined in the same way as the filtering step described above.
4. The colour of each node was determined by the group to which its descriptor belonged. As such, two nodes with descriptors in the same group would have the same colour.
5. The final position and size of each node (and the lengths of the edge links) were scaled according to the size of the rendering.

### 2.4 Implementation

We produced 2D and 3D visualisations of the graphs derived from our abstraction scheme. The 2D visualisation of our abstraction was implemented using the Java Universal Network/Graph (JUNG) library [17]. The 3D graph abstraction was implemented using WebCoLa [18].

## 3 Results and Discussion

Figure 2 shows four coronal PET-CT slices from a patient study. The PET and CT images have been fused (overlaid) and a colour table is used to highlight the PET functional information. The areas of high activity (bright yellow spots) in the thorax indicate the presence of lung tumours in all the slices. There were seven tumours segmented from this volume; several are not shown in any of these slices. Note that these images show the patient 'facing forward' and as such the patient's left side appears on the right side of the images, and vice versa.




Fig. 2: Four coronal slices from the PET-CT study to be abstracted. The PET and CT images have been ‘fused’. The bright yellow spots in the chest are potential tumour sites.

Figure 3 shows a 2D graph abstraction of the objects inside the PET-CT volume shown in Figure 2. The grey nodes represent tumours while the coloured nodes represent different anatomical structures. The abstraction shows that all the tumours were identified in the right lung (brown node) and that several of these tumours invaded the mediastinum (red node). This was the abstraction that we integrated into a PET-CT image retrieval engine in our previous work [6].

The 3D graph abstraction of the objects in the same PET-CT image is shown in Figure 4. The 3D abstraction reduces the occlusion of the tumour nodes in Figure 3. It is also easier to see where a tumour occurs in 3D space (above or below an organ, in the anterior or posterior parts of the body, etc.). The 3D abstraction had an improved ability, over the 2D abstraction, when discriminating between images based upon the anatomical location of tumours.

Figure 5 shows a 3D graph abstraction of the key points inside the PET-CT volume. There were 518 nodes in the abstraction corresponding to 518 key points from an image containing almost 6 million pixels. Due to the number of key points, this abstraction is filled with large groups of interconnected nodes thus leading to a large number of edge crossings. However, it is also possible to see nodes (key points) that form components of the graph that are connected to the central component by a small number of vertices. These components, marked with blue arrows, indicate interrelated points of interest within the image that are isolated from other areas.

Figure 6 shows the 3D graph abstraction of a different PET-CT image. The abstraction was generated using the same camera position in Figure 5. There were 360 nodes in the abstraction. The blue arrows indicate components of the graph that are connected to the central part of the visualisation by a small number of vertices. The low level of correspondence among these structures suggests a low degree of similarity between the two PET-CT images abstracted by these graphs. This conclusion was supported by the clinical reports. This abstraction

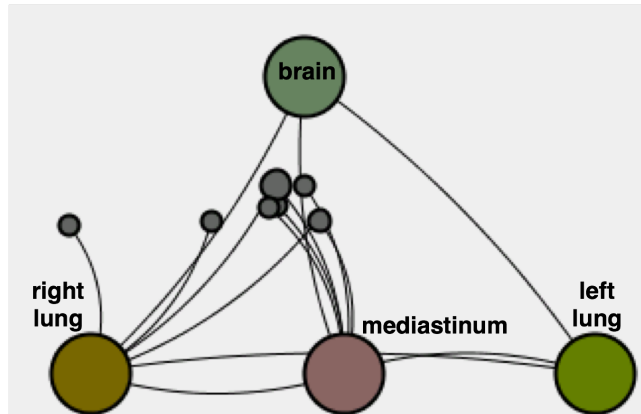


Fig. 3: A 2D graph abstraction of the objects within Figure 2.




Fig. 4: A 3D graph abstraction of the objects within Figure 2.

therefore made it possible to compare images based on the arrangement and groupings of the key points within the images.

Both the object and key point graph abstractions provide new views of the content of the PET-CT image. The object abstractions depict the location of the tumours and the structures that they affect. These abstractions summarise information that a clinician could potentially use when staging a cancer or when determining if a patient's prognosis is improving [2]. The key point abstractions can show the complex interrelations within the image. An important property of these images is when a portion of the graph is not highly connected with other parts of the graph; then depending on the image features, such components may be areas of further investigation, e.g. sites of new disease, tumour necrosis, etc. The 3D abstraction of these key points offers the opportunity for node merging or clustering to identify large objects of interest in different images because key points were originally used for object recognition [11].




Fig. 5: A 3D graph abstraction of the key points within Figure 2. The blue arrows indicate components of the graph representing isolated areas of interest in the image.

A limitation of our abstractions is the densely connected graph (see central parts of Figures 5 and 6). The densely connected graph is a result of the vertex layout being dependent upon the physical location of the image key points. Important information could then be hidden within these dense graphs, e.g., tumours within the mediastinum (the mediastinum is the central part of the chest). Grouping vertices that represent similar image features and that are pairwise adjacent into a single “super-vertex” would improve clarity in the visualisation of key points. This is a clique detection problem, which is computationally expensive.

Another limitation is that our abstractions may obscure detailed pixel information. However, since each node corresponds to a physical location within the image, it is possible to create links between the abstractions and the pixel data. In this manner, the abstraction can be used as a map of the important areas in the image. We implemented such a map for a PET-CT retrieval engine [6];






Fig. 6: A 3D graph abstraction of the key points within a different PET-CT image. The blue arrows point to isolated components. A comparison with Figure 5 shows that the abstractions (and thus their corresponding PET-CT images) are not similar.

clicking on an abstract node would outline the ROI in multiple views of the corresponding PET-CT image. Linking nodes on the abstraction to the pixel data is an interaction that can facilitate a more detailed understanding of the complex pixel data [7].

## 4 Conclusions

We present a scheme for creating graph abstractions that summarise the content within medical images. We provided 2D and 3D examples of our abstractions applied to 3D PET-CT images of patients with lung cancer. Our abstractions showed how complex image content could be summarised and interpreted.

In future work, we will adapt our abstractions to more complex diseases, e.g., lymphoma, which can have multiple clusters of tumours throughout the body. The abstraction of these images will be more complex and will require further optimisation of the properties and graph layout. We will investigate enhancements to our abstraction by hierarchically grouping related nodes based on the

spatial location of nodes in relation to body regions (head, thorax, abdomen, limbs) and by clustering cliques into a single representative node.

## References

1. Townsend, D.W., Beyer, T., Blodgett, T.M.: PET/CT scanners: A hardware approach to image fusion. *Semin Nucl Med* **33**(3) (2003) 193 – 204
2. Edge, S.B., Byrd, D.R., Compton, C.C., Fritz, A.G., Greene, F.L., Trotti, A., eds.: *AJCC Cancer Staging Manual*. Springer New York (2010)
3. Jung, Y., Kim, J., Eberl, S., Fulham, M., Feng, D.D.: Visibility-driven PET-CT visualisation with region of interest (ROI) segmentation. *Visual Comput* **29**(6-8) (2013) 805–815
4. Deserno, T., Güld, M., Plodowski, B., Spitzer, K., Wein, B., Schubert, H., Ney, H., Seidl, T.: Extended query refinement for medical image retrieval. *J Digit Imaging* **21**(3) (2008) 280–289
5. Hsu, W., Antani, S., Long, L.R., Neve, L., Thoma, G.R.: SPIRS: a web-based image retrieval system for large biomedical databases. *Int J Med Inform* **78**(Supplement 1) (2009) S13–S24
6. Kumar, A., Kim, J., Bi, L., Fulham, M., Feng, D.: Designing user interfaces to enhance human interpretation of medical content-based image retrieval: application to PET-CT images. *Int J Comput Assist Rad Surg* **8**(6) (2013) 1003–1014
7. Tory, M., Moller, T.: Human factors in visualization research. *IEEE T Vis Comput Gr* **10**(1) (2004) 72 –84
8. Bunke, H., Riesen, K.: Towards the unification of structural and statistical pattern recognition. *Pattern Recogn Lett* **33**(7) (2012) 811 – 825
9. Wilson, M.L.: Search user interface design. *Synthesis Lectures on Information Concepts, Retrieval, and Services* **3**(3) (2011) 1–143
10. Detterbeck, F.C., Boffa, D.J., Tanoue, L.T.: The new lung cancer staging system. *Chest* **136**(1) (2009) 260–271
11. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *Int J Comput Vision* **60** (2004) 91–110
12. Hu, S., Hoffman, E., Reinhardt, J.: Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images. *IEEE T Med Imaging* **20**(6) (2001) 490–498
13. Bradley, J., Thorstad, W.L., Mutic, S., Miller, T.R., Dehdashti, F., Siegel, B.A., Bosch, W., Bertrand, R.J.: Impact of FDG-PET on radiation therapy volume delineation in non-small-cell lung cancer. *Int J Radiat Oncol Biol Phys* **59**(1) (2004) 78–86
14. Kumar, A., Kim, J., Wen, L., Fulham, M., Feng, D.: A graph-based approach for the retrieval of multi-modality medical images. *Med Image Anal* **18**(2) (2014) 330–342
15. Toews, M., III, W.M.W.: Efficient and robust model-to-image alignment using 3d scale-invariant features. *Med Image Anal* **17**(3) (2013) 271 – 282
16. Zhou, X., Stern, R., Müller, H.: Case-based fracture image retrieval. *Int J Comput Assist Rad Surg* **7** (2012) 401–411
17. O'Madadhain, J., Fisher, D., White, S., Boey, Y.: The JUNG (Java Universal Network/Graph) framework (2003) <http://jung.sourceforge.net/>, Last Checked: 30/05/2014.
18. Dwyer, T.: *cola.js: Constraint-Based Layout in the Browser* (2003) <http://marvl.infotech.monash.edu/webcola/>, Last Checked: 28/05/2014.