



Automatic Detection and Visualisation of Metastatic Bone Disease

Nathan Sjoquist¹(✉), Tristan Barrett², David Thurtle³,
Vincent J. Gnanapragasam³, Ken Poole⁴, and Graham Treece¹

¹ Department of Engineering, University of Cambridge, Cambridge, UK
ns706@cam.ac.uk

² Department of Radiology, University of Cambridge, Cambridge, UK

³ Department of Urology, University of Cambridge, Cambridge, UK

⁴ Department of Medicine, University of Cambridge, Cambridge, UK

Abstract. This paper presents a novel method of finding and visualising metastatic bone disease in computed tomography (CT). The approach we suggest locates disease by comparing trabecular bone density between symmetric bony regions. Areas of strong difference could indicate metastatic bone disease as bone lesions either increase or decrease bone density. Our detection method is completely automatic and only requires raw CT data as input. Results are visualised in an interactive 3-dimensional viewer which displays a polygonal mesh of the bone structure overlaid with colour combined with resliced CT data. Diseased regions are clearly highlighted in both the mesh and in the resliced CT data. We test our method on both healthy and diseased CT data to demonstrate the validity of the technique. Experimental results show that our method can detect metastatic bone disease, although further work is needed to improve the robustness of the technique and to decrease false positives.

Keywords: Metastatic bone disease · Disease asymmetry · CT analysis · CT segmentation · Articulated registration · Computer aided detection

1 Introduction

Metastatic bone disease (MBD) is a common secondary feature of prostate cancer, breast cancer and many other malignancies. Bone scintigraphy is the usual imaging detection technique but has limitations in detecting changes over time and is only 2-dimensional. With improvements in bone-targeting therapies, better evaluation of therapeutic response is now a Europe-wide research priority. Positron-emission tomography is promising, but the expense and inaccessibility means CT is increasingly used for serial assessment. This requires time-consuming reviews of many CT cross-sectional images looking for small and diffuse bony changes – the summary of which is difficult to demonstrate in a

Supported by the W. D. Armstrong Trust.

© Springer Nature Switzerland AG 2020

Y. Zheng et al. (Eds.): MIUA 2019, CCIS 1065, pp. 327–338, 2020.

https://doi.org/10.1007/978-3-030-39343-4_28

multidisciplinary review setting as well as to demonstrate to the patient. To address these challenges, we have designed and implemented a new method for the detection and visualisation of MBD.

The human skeleton is nearly symmetric in both shape and bone density. When two healthy symmetric bony regions are compared with each other, the regions will hence contain little difference. However, metastatic bone disease causes sclerotic and lytic bone lesions to form, increasing the density (sclerotic lesions) or decreasing the density (lytic lesions) of the bone. When the bone density of symmetric regions containing bone lesions is compared, the difference is greatly increased as the lesions do not generally form symmetrically across both regions.

Our algorithm can automatically find these bone lesions by symmetrically comparing the bone density in the human skeleton and grouping areas of large difference. The method displays these differing densities through a novel visualisation. We believe this method has the potential to improve accuracy and to reduce the time-consuming review of the individual inspection of hundreds of cross-sectional images. We envisage our technique being used to identify potential areas of concern, although the final reporting of a bone lesion will require multi-planar assessment of the source data by the radiologist, oncologist and surgical team. In this paper, we describe our method in detail and show that it can automatically locate and visualise MBD in CT.

2 Background

Bone is a common site for metastasis by malignant tumours such as the lung, breast, and prostate cancer. This is significant as MBD can cause more devastating effects than the primary cancer does. MBD substantially increases morbidity due to its complications which include pain, impaired mobility, pathological fracture, spinal cord compression, cranial nerve palsies, nerve root lesions, hypercalcaemia and suppression of bone marrow [6] making it an important area of research.

In normal bone, development and maintenance of bone tissues is sustained through a balance of osteoclasts and osteoblasts resorbing and depositing bone tissue. With MBD, this process is disrupted as cancerous cells from the primary site contribute to the establishment of bone metastatic lesions. Most metastases are osteoblastic, causing an increase in bone density, although some metastases are osteolytic causing a reduction of bone density. In both cases, MBD causes an abnormal change in the density of bone.

Publications describe methods that can semi-automatically or automatically detect bone lesions by comparing learned features (generated by hand or through an automated machine learning technique) to raw and segmented CT data [3–5, 10–12, 16, 18]. While these methods are promising, they have a limited spatial scope (only the vertebral column), offer poor visualisation of the results (only 2-dimensional) and often have a low accuracy. Machine learning in general also requires large quantities of accurately marked training data which can be

difficult to obtain in a medical imaging setting due to patient agreements and lack of experts to mark the data.

Instead of using a machine learning approach, our paper explores a completely new approach to the problem. Rather than using training data to learn features distinct to bone lesions, we compare each CT data set to itself through symmetry in order to find irregularities. This removes the need for training data and bypasses the problem often encountered with machine learning algorithms – poor results when the training data does not contain a specific, never before encountered pattern. Also, instead of searching directly for 2-dimensional bone regions, we compare the modelled bone density between each symmetric pair of points in our 3-dimensional mesh, which reduces the dimensionality of the problem (when compared to searching for all combinations of 2-dimensional slices through an entire CT data set).

Using symmetry to locate disease has been shown to be a successful approach in a variety of situations [13, 17]. Although results from these papers could be improved, they demonstrate the feasibility of the approach. We have extended these ideas and have applied them to the problem of locating MBD. When compared to the previous approaches [13, 17], the rigidity of the 3-dimensional bone improves symmetric matching as the bone surface is not malleable like soft tissue. We do not compare shape differences directly (which is a harder problem) but only the bone density between each pair of points. This leads to improved accuracy. We chose to compare the less dense inner trabecular bone density instead of the more dense outer cortical bone, as the trabecular density measurements are more often affected by the disease.

3 Methodology

Our method first extracts regions containing bone in the CT data. It converts these 2-dimensional segmentations into a 3D polygonal (target) mesh that encapsulates the bone structure. Using the Cortical Bone Mapping (CBM) method [15], we measure the trabecular bone density. Our method then solves for symmetric point correspondence between left and right sides by registering the segmented target mesh with an atlas mesh. Finally, we display differences in trabecular bone across symmetric regions in an interactive 3-dimensional viewer.

3.1 Segmentation

In order to measure trabecular bone density and to find bone symmetry, the contours that surround regions of bone are first located within the CT images through the process of segmentation. In this process, pixels are classified as being either of a bone type or of a non-bone type.

Pixels labelled as bone generally contain a higher intensity than that of non-bone (unless the CT data contains a metal implant or a contrast agent). Bone has a radiodensity typically greater than 1200 Hounsfield units (HU) while soft

tissue and fat have a radiodensity typically between -70 HU and 100 HU. However, in the region between 100 HU and 1200 HU, there is much overlap between the radiodensity of trabecular bone, fat and soft tissue. Because of this, a single intensity threshold cannot be used to separate bone from non-bone, as a threshold set low enough to include all bone will include also fat and soft tissue. A threshold set high enough to exclude fat and soft tissue will also exclude bone. To overcome this issue, context from surrounding pixels must be combined with pixel intensity to correctly classify a pixel as being bone or non-bone.

Many approaches have been developed that can segment bone in CT. However, the problem has yet to be fully and accurately solved. Approaches can be placed into general groups including adaptive thresholding, hysteresis thresholding, region growing, watershed, active contour, edge based, level set, graph cuts, statistical shape models, articulated registration based and machine learning based.

We chose to use a hysteresis thresholding technique as regions of bone and non-bone can be generally divided into high intensity (bone) and low intensity (non-bone) regions making it a natural thresholding type problem. Hysteresis thresholding addresses the overlapping regions of bone and non-bone by provided context which greatly improves the accuracy when compared to global thresholding. Furthermore, the Cortical Bone Mapping (CBM) method used for measuring bone density, produces the best results when the contour lines surrounding the bone lie exactly between the bone and non-bone pixels. Because we use a thresholding technique, we can produce contour lines with sub-pixel accuracy by interpolating between threshold pixels, improving both the density calculation and the segmentation accuracy.

First, we select all pixels that are very likely bone (dense outer cortical bone). Then, a second lower threshold is used to select all pixels that may be bone but also may contain some soft tissue and fat. The pixels segmented at this lower level are only kept if they are connected to a pixel segmented at the higher threshold.

Additionally, we developed a new method to improve accuracy based on the gap filling ideas from Gelaude [9]. We found that using a third threshold slightly lower than the second was useful in filling gaps. These gaps occur as the hard cortical bone is sometimes non-existent in the CT. They can severely decrease segmentation accuracy as the contour traces the inside of the bone structure instead of the outside. Figure 1a shows an example.

To address this, we first generate a contour line using a third slightly lower threshold. We then compare this contour line with the contour line generated by the hysteresis thresholding technique. As the third threshold is lower than both the hysteresis thresholds, its contour line lies on the outside of the hysteresis contour at all points. By following these contour lines in counter clockwise direction, each time these two contours diverge, the path lengths are compared. The contour lines lie relatively close to each other until gaps in the cortical bone are reached. In these locations, the contours often diverge as the hysteresis contour traces the gap and then re-joins with the lower contour. The hole is filled if the

path length of the hysteresis threshold contour is more than three times longer than that of the lower contour. Figure 1b shows an example of a hole being filled. The green (hysteresis) contour follows the blue (lower threshold) contour until they diverge at the gap. The yellow contour shows the divergence. Its length is compared to the length of the blue line during the divergence, causing the final hysteresis contour to skip the gap.




Fig. 1. Generic segmentation of bone from CT data. (Color figure online)

3.2 Bone Mesh Creation and Cortical Bone Mapping (CBM)

Once contours are found in each original CT image, the marching tetrahedra algorithm is used to create a polygonal mesh of the bone structure. We used the Stradwin software implementation [14] for this. Figure 2a shows the results of our segmentation and mesh

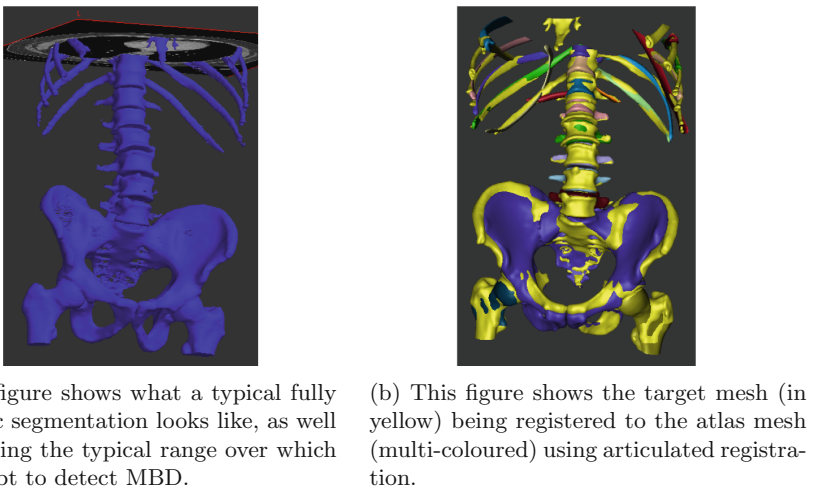




Fig. 2. Segmentation and registration. (Color figure online)

Once the 3D polygonal meshes are formed, the Cortical Bone Mapping (CBM) technique is used to accurately estimate trabecular bone density [15]. In CBM, the trabecular bone density is estimated by following the vector normal at each triangle in the mesh into the bone. The distribution of bone density along this line is modelled, which crucially allows average trabecular density to be estimated without any bias from the nearby, and much denser, cortical bone. Figure 3a shows the outline of the segmented bone in red running parallel to the cortical bone and a line normal to the cortical bone running through the bone surface in cyan while Fig. 3b shows the density being measured.



(a) Measuring the cortical thickness along the cyan line through the segmented bone.



(b) Measuring the bone density by modelling the cortical thickness.

Fig. 3. CBM can measure cortical bone mineral density (CBMD), trabecular bone mineral density (ECTD), cortical thickness (CTh), and cortical mass surface density (CMSD). (Color figure online)

3.3 Finding Symmetry

In order to compare the density of symmetric regions, we search for a symmetric point correspondence for each point in the segmented target mesh. To do this we register the target segmentation to an atlas containing a polygonal mesh of a healthy skeleton where each point has been symmetrically paired with another.

3.4 Atlas Creation

To create our atlas, a healthy CT data set was initially segmented using the threshold feature in the Stradwin software [14]. The segmentation was then thoroughly corrected by hand again using the Stradwin software. The full segmentation of the skeleton was also split into multiple regions by labelling general bone regions (pelvis, left femur, right femur, individual ribs and vertebrae).

Each segmentation was then paired with a symmetric region by mirroring one to match the other (e.g. left femur matched to right femur). In the case of a bone region without a corresponding symmetry match (vertebrae and pelvis), the region is matched with an x-axis inverted copy of itself. Similarity-based rigid registration was first used to register each pair of symmetric regions using the wxRegSurf software [8]. Where necessary, deformable registration was used to improve accuracy.

3.5 Articulated Registration

The atlas is registered to the target segmentation using an articulated registration technique. Our registration method was inspired by the techniques used to segment skeletons in mouse CT [1] and human CT [7]. Registration in general produces the best results when initial placement is as accurate as possible. Articulated registration takes advantage of this through a series of steps. The entire skeleton as a whole is first registered for initial placement. Then smaller and smaller sub-regions of the skeleton are registered following a hierarchical anatomical tree atlas. This guides the small subsections of the skeleton into a generally correct location before final registration. This greatly improves accuracy as the whole skeleton guides the registration but each bony region can move independently of each other. For each registration step, a rotation, translation and scale is found that minimizes the distance between each point on the atlas (or bony sub-region) and the target mesh. We implemented the classic registration method described by Besl [2]. Figure 2b shows the atlas mesh being registered to the target mesh using articulated registration.

Our method differs from the papers above [1, 7] as their purpose is to segment the skeleton. Instead, we use this registration technique to find the symmetry between each point. At the end of the registration, for each point in the target mesh, we look up the closest point in the atlas mesh. Each point in the atlas has a symmetric point, so we use that symmetric point to look up the closest point in the target mesh. In this way we can find the symmetric mappings between all points in the target mesh. We compare vector normals and point distances (between atlas and target) and only set symmetric mappings if these are within a reasonable range. Our current registration method is only rigid with scaling—in future we plan to implement a non-rigid registration to improve upon the performance shown in Fig. 2b.

4 Visualisation of Disease

Once every point on the left side has been mapped to a corresponding point on the right side, the trabecular bone density at every point on one side of the skeleton is compared with the other side. As long as one side is diseased, the difference in trabecular bone density between left and right is significant. We wrote a specialized viewer to visualise the results by displaying a polygonal mesh of the skeleton and uniquely colouring each vertex. Figure 4a shows the trabecular bone density (i.e. not using symmetrical difference) overlaid on the mesh. A darker blue indicates a more dense bone. Figure 4b shows the segmented mesh with a raw CT slice.

Areas of large difference are grouped together, and each group has the potential to be either healthy or diseased. Healthy regions are still included with diseased regions at this point, as all groups with large differences in trabecular bone density between left and right sides are selected without taking into account which side caused the differences. Intuitively, trabecular bone density in healthy tissue will be relatively constant, while diseased bone that has increased

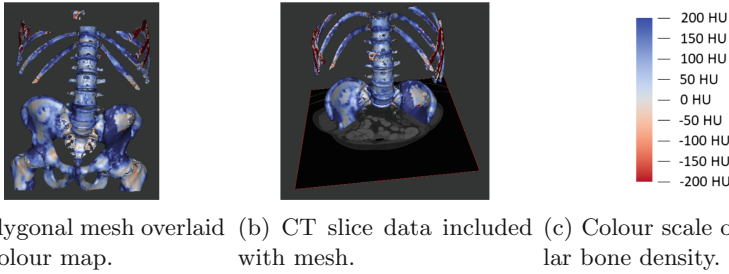



Fig. 4. The segmentation result is overlaid with a colour map representing the trabecular bone density measurements. (Color figure online)


or decreased bone density will have certain intense changes in trabecular bone density. When the gradient of trabecular bone density of healthy bone is compared to the gradient of its difference in trabecular bone density, the differences are generally high. However, when the gradient of trabecular bone density of diseased bone is compared to the gradient of its difference in trabecular bone density, the differences are generally low. In this way, regions are determined to be either diseased or healthy. Figure 5a shows the trabecular bone being subtracted from each side. In Fig. 5b, the density differences greater or less than 200 HU are grouped together. In Fig. 5c, only the diseased areas are shown.




(a) The density of each side is subtracted from the other side. (b) Regions of strong density differences are grouped together. (c) Healthy regions are removed by comparing gradients.

Fig. 5. Disease is visualized by comparing trabecular bone density differences and gradients of symmetric regions.

A user can click on any point in the mesh and find its symmetric point pair. A green line is drawn between the points, and their surface normals are drawn in red and blue. A new, re-slice image plane is defined using the vector pointing from first point to the second, and the vector created by adding the two normals together. In Fig. 6a, a user has clicked on the diseased right side. The CT is automatically re-sliced to show the diseased region. Figure 6b shows MBD on the right side of the pelvis. In Fig. 6b, it can be seen that the right side contains a much greater trabecular bone density than the left side does.



(a) Re-slicing the CT data at the user selected location.




(b) The diseased right side can be seen in the re-sliced CT image.


Fig. 6. A user can click on the mesh to calculate a re-sliced CT image at that location.

5 Results and Discussion

We tested our method using eleven diseased CT data sets and one healthy CT data set. All of our CT data sets are under an ethics agreement enabling us to use them for research and publication. We compared the diseased areas found in our viewer (areas of high colour contrast) to metastatic scoring sheets from conventional CT review marked by experts. We also compared our results to 3-dimensional bone lesions marked by an expert which can be seen in Fig. 7a. Areas were marked as being either malignant or suspicious. Malignant areas are diseased regions and suspicious areas are abnormal regions that could be disease. Table 1 shows our results.



(a) 3-dimensional lesions marked by an expert are displayed in red inside the mesh and compared to our results displayed on the surface of the mesh.



(b) Results correctly show our results displayed on the surface of the mesh. Results correctly show our results displayed on the surface of the mesh to be mostly disease free.

Fig. 7. Results are visually compared to 3-dimensional lesions marked by an expert to determine accuracy.

5.1 Evaluation

Our method found 73.1% of suspicious areas and 79.2% of malignant areas correctly. It found 77.0% of all diseased areas. Our results also contained a number of false positives in each data set - most of which are small in size. This can be

seen in Fig. 8. However, there are limitations in finding the true accuracy of our results as the clinician-defined areas of suspicion may not be perfect. This may be the cause of some of the false positives as they may actually be abnormal regions missed by the expert.

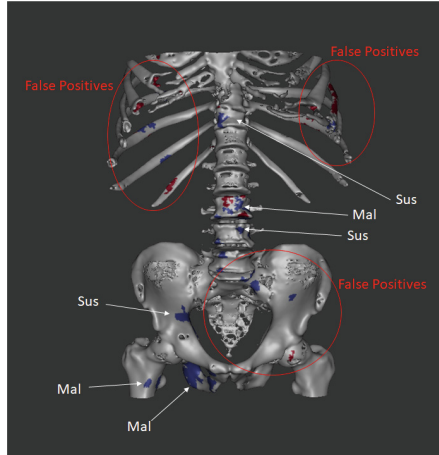


Fig. 8. A typical result showing suspicious areas (Sus), malignant areas (Mal) and false positives.

Table 1. Sensitivity of the viewer finding diseased regions.

Name	Suspicious areas	Malignant areas
Case 1	3/4	1/1
Case 2	4/4	1/1
Case 3	2/3	5/8
Case 4	1/2	None
Case 5	1/2	None
Case 6	1/2	1/1
Case 7	None	4/4
Case 8	2/3	3/4
Case 9	2/3	3/4
Case 10	3/3	9/13
Case 11	None	11/12

We also tested our method on a healthy CT data set (our atlas) by processing its raw CT data. The method found the healthy data mostly disease free although it incorrectly displays small spots of disease in a few locations. This can be

seen in Fig. 7b. Some of these spots of disease are due to natural and healthy asymmetry. This can be seen in the sacroiliac joint which is shown as partially diseased. These areas of healthy asymmetry will need to be addressed in the future.

5.2 Discussion

These results demonstrate that our technique can successfully detect and visualise metastases in bone. Our method already has a relatively high sensitivity in finding areas of diseased bone, but can be improved. It can find bone metastases in the vertebra, pelvis and upper femurs, which is an improvement over existing methods which are mostly limited to the vertebra.

Our results currently contain many false positives. These are due to errors in the automatic segmentation, registration and symmetry point matching although in some cases, our method may be correctly identifying abnormal regions missed by the expert. One important next step to reduce false positives, is to use a deformable registration technique that can both preserve atlas symmetry but also provide more accurate matching with the target mesh. After this step, the deformed atlas mesh can be used to measure the trabecular bone density instead of using the target segmentation mesh. This should greatly improve the accuracy, as errors will be reduced both in the trabecular bone density measurement and in the symmetric point matching.

Our method produces poor results in the ribs. This is due to significant errors in the registration as there is much variability between rib shape and size. Poor registration produces poor symmetric point matching severely decreasing the accuracy of this method. Deformable registration as well as improvements to the rib components of the articulated registration should help with this problem.

6 Conclusions

We have introduced a new algorithm that can automatically detect and display metastasized cancer in bones. Preliminary results demonstrate that the algorithm can successfully find areas of bone disease and visualise them in a unique way. It already has a relatively high sensitivity in finding areas of diseased bone but it can still be improved. The technique currently suffers from a high false positive rate, which will need to be addressed. It is fully automatic, and successfully works as a diagnostic tool providing a unique way to visualise disease.

References

1. Baiker, M., et al.: Atlas-based whole-body segmentation of mice from low-contrast Micro-CT data. *Med. Image Anal.* **14**(6), 723–737 (2010)
2. Besl, P.J., McKay, N.D.: Method for registration of 3-D shapes. In: *Sensor Fusion IV: Control Paradigms and Data Structures*, vol. 1611, pp. 586–607. International Society for Optics and Photonics (1992)

3. Burns, J.E., Yao, J., Wiese, T.S., Muñoz, H.E., Jones, E.C., Summers, R.M.: Automated detection of sclerotic metastases in the thoracolumbar spine at CT. *Radiology* **268**(1), 69–78 (2013)
4. Chmelik, J., Jakubicek, R., Jan, J.: Tumorous spinal lesions: computer aided diagnosis and evaluation based on CT data—a review. *Current Med. Imaging Rev.* **14**(5), 686–694 (2018)
5. Chmelik, J., et al.: Deep convolutional neural network-based segmentation and classification of difficult to define metastatic spinal lesions in 3D CT data. *Med. Image Anal.* **49**, 76–88 (2018)
6. Diel, I.J., Kaufmann, M., Bastert, G.: *Metastatic Bone Disease: Fundamental and Clinical Aspects*. Springer, Heidelberg (2012)
7. Fu, Y., Liu, S., Li, H.H., Yang, D.: Automatic and hierarchical segmentation of the human skeleton in CT images. *Phys. Med. Biol.* **62**(7), 2812 (2017)
8. Gee, A.: wxregsurf software (2019). <http://mi.eng.cam.ac.uk/~ahg/wxRegSurf/>
9. Gelaude, F., Vander Sloten, J., Lauwers, B.: Semi-automated segmentation and visualisation of outer bone cortex from medical images. *Comput. Methods Biomech. Biomed. Eng.* **9**(1), 65–77 (2006)
10. Jan, J., Novosadova, M., Demel, J., Ouředníček, P., Chmelik, J., Jakubíček, R.: Combined bone lesion analysis in 3D CT data of vertebrae. In: 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pp. 6374–6377. IEEE (2015)
11. Peter, R., Malinsky, M., Ourednicek, P., Jan, J.: 3D CT spine data segmentation and analysis of vertebrae bone lesions. In: 2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pp. 2376–2379. IEEE (2013)
12. Roth, H.R., Yao, J., Lu, L., Stieger, J., Burns, J.E., Summers, R.M.: Detection of sclerotic spine metastases via random aggregation of deep convolutional neural network classifications. In: Yao, J., Glocker, B., Klinder, T., Li, S. (eds.) *Recent Advances in Computational Methods and Clinical Applications for Spine Imaging*. LNCVB, pp. 3–12. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-14148-0_1
13. Tahmouh, D., Samet, H.: Using image similarity and asymmetry to detect breast cancer. In: *Medical Imaging 2006: Image Processing*, vol. 6144, p. 61441S. International Society for Optics and Photonics (2006)
14. Treece, G.M.: Stradwin software (2019). <http://mi.eng.cam.ac.uk/~rwp/stradwin/>
15. Treece, G.M., Poole, K.E., Gee, A.H.: Imaging the femoral cortex: thickness, density and mass from clinical CT. *Med. Image Anal.* **16**(5), 952–965 (2012)
16. Wiese, T., Yao, J., Burns, J.E., Summers, R.M.: Detection of sclerotic bone metastases in the spine using watershed algorithm and graph cut. In: *Medical Imaging 2012: Computer-Aided Diagnosis*, vol. 8315, p. 831512. International Society for Optics and Photonics (2012)
17. Yang, L., Xie, Y., Li, B., Xie, M., Wang, X., Zhang, J.: Symmetry based prostate cancer detection. *Br. J. Radiol.* **88**(1050), 20150132 (2015)
18. Yao, J., Burns, J.E., Summers, R.M.: Computer aided detection of bone metastases in the thoracolumbar spine. In: Li, S., Yao, J. (eds.) *Spinal Imaging and Image Analysis*. LNCVB, vol. 18, pp. 97–130. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-12508-4_4